The Effects of Short-Term Rental Regulations: Evidence from the city of Santa Monica

Cayruã Chaves Fonseca*

Abstract

This paper studies the impacts of a specific regulation restricting short-term rental activity, the Home-Sharing Ordinance, adopted in the city of Santa Monica in May of 2015. It mainly focuses on carefully estimating how the ordinance has affected the number of housing units operating on Airbnb's platform. Using a dataset of Airbnb listings in the area surrounding the city of Los Angeles, I find that the ordinance has reduced the number of entire homes listed on Airbnb in Santa Monica by approximately 61%. I also study the impacts of this regulation on the long-term rental market and I find no evidence of a significant effect of the ordinance on residential rents in Santa Monica. Lastly, I provide suggestive evidence of the extent to which the policy under study has had any effect on housing reallocation in the city.

Keywords: Short-term rentals; Airbnb; Housing Rents; Policy Evaluation

JEL Classification: R31, R52, Z30, M13

*Chaves Fonseca: CEMFI, Calle Casado del Alisal, 5, 28014, Madrid, Spain (e-mail: cayrua.chaves@cemfi.edu.es).
I would like to thank Gerard Llobet for his guidance and support. For helpful suggestions and comments I also want to thank Diego Puga, Manuel Arellano, Guillermo Caruana, Pedro Mira, and participants of CEMFI's PhD workshop series.

1 Introduction

Short-term rentals and home-sharing have experienced remarkable growth over the last years. It is true that this kind of activity has always existed, specially in touristic cities during peak season. However, the emergence of online peer to peer platforms such as Airbnb, which reduced the costs of short-term renting, both from the demand and the supply side, has drastically increased the number of short-term rental units available.

It is hard to deny that there are many benefits related to the expansion of short-term rentals. For example, an individual struggling to make ends meet, may rent out part of his residential unit to visitors and use the extra income to afford his own housing expenses. Moreover, short-term rental units tend to be a cheaper accommodation option when compared to hotels, which surely benefits travellers. Additionally, supporters of home-sharing argue that short-term rentals of the style of the ones offered by Airbnb provide an experience of "living like a local" that hotels cannot offer to tourists.

On the other hand, there is a growing concern about the potential costs associated with the rise in short-term rentals (STR) and home-sharing. Critics have pointed out that the ever increasing presence of tourists in residential neighborhoods may bring negative consequences, such as a lack of community feeling, increased noise and disturbances, and increased competition for rival public resources such as parking space (Edelman and Geradin 2015). Besides these issues, another major concern relates to the potential impacts that short-term rentals may have on housing costs. The intuition behind this idea is straightforward. Assuming that the total number of housing units is fixed in the short-run, the expansion of short-term rental units comes at the cost of a reduction in the units offered for long-term tenants, which could cause rents to increase for residents (Barron, Kung, and Proserpio 2018).

The combination of these factors - the impressive growth in short-term rentals and the controversy in the potential effects that it may impose on residents - has caused many cities to enact local regulations specifically designed to deal with STR and home-sharing. A primary example of such local laws is the Santa Monica Home-Sharing Ordinance, which is one the strictest home-sharing regulations currently in place and completely prohibited short-term rentals of entire home units in the city.[1]

In this paper, I use a panel data set containing information on Airbnb listings and local rents in cities within the Los Angeles County to empirically evaluate the main impacts of the Santa Monica Home-Sharing Ordinance. The main focus is on presenting a detailed estimation of the causal impact of the ordinance on the number of Airbnb entire home listings operating in Santa Monica. Using a Synthetic Control Method approach (as in Abadie, Diamond, and Hainmueller 2010), I find that the Home-Sharing Ordinance, roughly two years after its adoption, had reduced the number of entire homes listed on Airbnb in Santa Monica by 861 units, which represents a 61% reduction in listings.

Moreover, I also try to answer the question of whether the ordinance has had any significant impact on the rents faced by long-term tenants. Using a similar estimation strategy, I do not find any significant effects of the Home-Sharing Ordinance on long-term rents in Santa Monica. Finally, by investigating the effects of the ordinance on housing reallocation, I suggest that long-term residential rents were not affected by the regulation because owners who were prevented from using their housing units as Airbnb short-term rentals are not yet supplying these units for

[1] Other examples include Berlin, which completely banned short-term rental of entire home units, San Francisco, where short-term rental units cannot be rented for more than 90 days per year, and Barcelona, where vacation rentals require a special tourist license which is limited to some neighborhoods. Information obtained at http://blog.airdna.co/effects-airbnb-regulation.

long-term tenants.

The remaining of the paper is organized as follows. In Section 2, I explain the structure of the Santa Monica Home-Sharing Ordinance. In Section 3, I describe the data and show some basic statistics. In Section 4, I present in detail the estimation of the regulation's impact on the number of Airbnb entire home listings active in Santa Monica. Section 5 includes the estimated effects of the ordinance on residential rents. Section 6 presents suggestive evidence of the regulation's impact on housing reallocation and section 7 concludes.

2 The Santa Monica Home-Sharing Ordinance

Santa Monica, a beach-front city in western Los Angeles County, has been a famous touristic destination town since the early 20th century.[2] The attractiveness of the city to tourists makes it an ideal place for the proliferation of short-term rental units (hereafter also denoted vacation rentals). By April 2015, among all cities in the L.A. county, Santa Monica was the second largest Airbnb market, featuring 955 listings active on the world's largest home-sharing platform (only below the city of L.A. itself, which by its very large size naturally featured more listings).

The large number of short-term rental units in the city and a growing concern about housing affordability for residents led the Santa Monica City Council to adopt the Home-Sharing Ordinance on May 12th, 2015. When adopting the new law, the city's administration explicitly expressed its worry that the rise in the supply of vacation apartments was effectively reducing the number of residential rental units that would otherwise be available for long-term tenants. Given this priority, the Home-Sharing Ordinance, which became effective on June 12th, 2015, separated into two the types of listings offered for short-term rental. The city administration would take a very different regulatory approach towards what was defined as "vacation rentals" as compared to what was defined as "home-sharing".

By "home-sharing", the City Council refers to the activity whereby individuals rent out part of their home to short-term guests, who will share the housing unit with at least one of the primary residents. That is, home-sharing as defined by Santa Monica's legislation requires at least one of the primary residents to be on-site sharing the home unit with the short-term guest. This kind of activity is considered legal within the framework of the Home-Sharing Ordinance, which only requires hosts to register with the city, get a business license (issued for free), and pay a 14% transient occupancy tax (the same taxation that hotels are subjected to). On the other hand, by "vacation rentals", the regulation refers to the activity whereby an individual rents out an entire home unit to a short-term guest, who makes exclusive use of the whole residential unit for a period of less than 30 days. Operating a vacation rental unit was deemed illegal by the ordinance and violators could face administrative fines, and even criminal prosecutions in case they refuse to cease operation.[3]

Besides the separation between home-sharing and vacation rentals, understanding the context in which the regulation was implemented and enforced is also of extreme importance to evaluate its impacts. Crucially, it is important to have in mind that the Home-Sharing Ordinance went through some modifications and amendments over time. The timing of the main aspects that constitute the regulation may be summarized as follows.[4] On May 12th, 2015 the ordinance was adopted by the Santa Monica City Council. One month later, on June 12th, 2015, the law became effective. Then, over the summer, staff conducted an educational campaign mainly focused on circulating

[2]Source: https://www.santamonica.com.

[3]Source: https://www.smgov.net/Departments/PCD/Permits/Short-Term-Rental-Home-Share-Ordinance.

[4]Information about the progress and modifications in the Home-Sharing Ordinance was obtained from direct e-mail exchanges with enforcement officers through the address code.enforcement@smgov.net.

information about the new law. Announcements were made via internet, local TV, newspapers, and even residential water wills contained information on the new prohibitions regarding vacation rentals. During this period, warning letters were sent to potential violators and enforcement action was taken only in cases arising from formal complaints.

Simultaneously, the City Council was setting up a team of three people, the "Short-Term Rental Team", exclusively dedicated to enforce the Home-Sharing Ordinance. While the educational campaign was under way, the Short-Term Rental Team designed the next actions and priorities of the enforcement plan. In accordance with the main goals of City Council when enacting the ordinance, the enforcement officers determined that the highest priority for enforcement would be professional hosts, defined as hosts who simultaneously operate multiple vacation rental units. The subsequent important milestone of the regulation came on November 1st, 2015, when the enforcement officers started what was defined as proactive enforcement. That meant that instead of conducting enforcement actions only when complaints arrived, the enforcement team would actively search for illegal activity and issue fines to violators. Lastly, on January 26th, 2016, the fine for operating a vacation rental unit (as well for operating a home-sharing unit without the proper license) was raised from $75 to $500 per day of illegal activity.

3 Data

3.1 Airbnb Listings

Airbnb is one of the most successful companies of the sharing economy. The company provides a convenient peer to peer on-line marketplace for short-term rentals, where hosts can advertise different types of accommodation units to potential renters. Airbnb, since its launch in 2008, has managed over 150 million guest arrivals.[5] The company profits from charging guests and hosts a percentage from each booking made through the platform.

All data regarding Airbnb listings was obtained from Tom Slee's webpage.[6] Using data provided by this source, I build a panel dataset containing all Airbnb listings in the area of Los Angeles County and some of its basic characteristics. The panel goes from September 2014 to July 2017 and contains data collected in 29 different points in time (hereafter denoted as "waves"). Although the initial waves of the panel are not evenly spaced, from August 2015 on, data is collected at (close to) monthly intervals.

The final dataset contains information on 77,502 unique listings and 47,142 unique hosts. It incluces basic variables faced by users of Airbnb's website such as a unique time-invariant identification number associated with each listing, another unique code associated with each host (who may own more than one listing), property location, the price charged per night, the number of bedrooms, the accommodation type (entire home or shared home), the average rating, and the number of reviews received up to that date.

3.2 Long-Term Rental Rates

Data regarding long-term rents was obtained from Zillow, one of the leading real estate and rental online marketplaces in the United States. Zillow's database contains information on more than 110 million U.S. homes as well as indexes that track the key housing market variables across a given region on a monthly basis.[7]

[5]Source: www.airbnbcitizen.com.

[6]Tom Slee is a blogger, sharing economy author, and software engineer, who scrapes data directly from Airbnb's website and makes it available at http://tomslee.net.

[7]Source: https://www.zillow.com/corp/About.htm.

To capture long-term market rent prices, I use the Zillow Rent Index (ZRI). The ZRI is a smoothed, seasonally adjusted measure that reflects the median market rent for the entire stock of homes in a given region at a given point in time. Importantly, since Zillow is a marketplace for landlords and tenants looking for long-term contracts, the ZRI reflects the rental rates faced by residents looking for a home to rent, which is the relevant rental price in terms of policy evaluation (Barron, Kung, and Proserpio 2018).

3.3 U.S. Census Bureau

Besides data on Airbnb listings and long-term rental rates, I also use location-year specific variables available from the U.S. Census Bureau.[8] From the American Community Survey (ACS) I obtain estimates of population, income, college education, and employment, which serve as the main control variables in most of the analysis. From the Business Patterns (BP), I get extra control variables associated with how attractive to tourists a given location is. To this purpose, I use the number of establishments in food and accommodation (NAICS code 72) as well as in arts and entertainment (NAICS code 71). And lastly, again from the ACS, I compile information on the proportion of housing units occupied by renters and the ratio of houses vacant in each region at different points in time. These two variables are used as outcomes in my analysis of whether the Home-Sharing Ordinance has had any significant impact on housing reallocation.

3.4 Summary Statistics

The Los Angeles County has seen persistent growth in the number of Airbnb listings throughout the period under study, Figure 1 depicts this trend.[9] In September 2014, the first month-year for which I have data on Airbnb listings, there were 10,027 active listings in the area under study. By July 2017, the last time period of my dataset, the same region featured a total of 32,146 active listings in the same home-sharing platform, an increase of 220.59% in less than three years.

Figure 1: Total Number of Airbnb Listings in the city of L.A. and neighboring cities

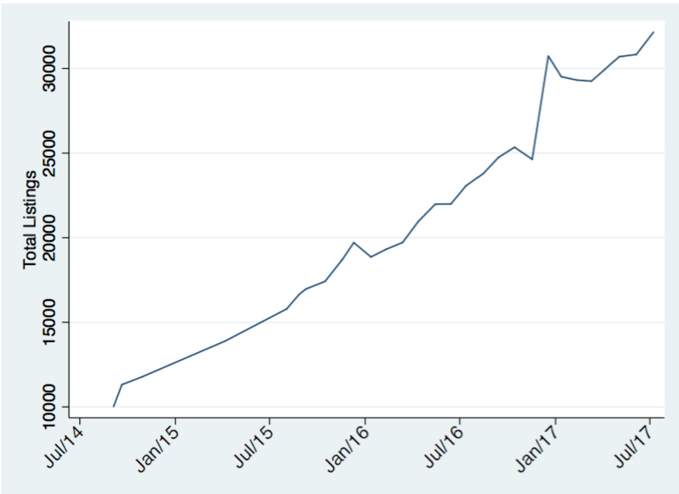

Note: Total listings include both entire home and shared home units

Table 1 shows the mean values of a few basic variables concerning Airbnb listings and long-term rents for the city of Santa Monica, for Venice, as well as the mean for all other cities in the

[8] https://factfinder.census.gov.

[9] All plots presented in this paper are constructed using data from http://tomslee.net/airbnb-data-collection-get-the-data.

L.A. county. Venice is an interesting comparison unit because it is adjacent to Santa Monica but for administrative purposes it belongs city of L.A., where no short-term regulation was explicitly imposed during the period of analysis. Table 1 only include numbers related to entire home listings, which are the ones mostly affected by the Home-Sharing Ordinance and constitute the focus of this paper. The table features the mean values across all listings active in a given location in four different points in time. In April 2015, right before the regulation was adopted. In October 2015, when the ordinance was already in effect but the proactive enforcement had not yet started. In February 2016, after the proactive enforcement was in place for three months and the fine had recently been increased from $75 to $500 per day. And lastly, in July 2017, the last month-year of my dataset.

Table 1: Mean characteristics by geographical region and period (entire home listings only)

City	Variable	Apr/2015	Oct/2015	Feb/2016	Jul/2017
Santa Monica	Entire home listings	695	622	401	570
	Price per night ($)	275.86	231.76	282.88	259.57
	Average rating	4.69	4.67	4.67	4.77
	Duration (months)	14.63	14.03	19.04	13.13
	Monthly Rent ($)	4380	4682	4887	4984
Venice	Entire home listings	921	1068	1117	1569
	Price per night ($)	306.16	243.29	239.70	236.29
	Average ratings	4.72	4.73	4.73	4.77
	Duration (months)	23.49	23.85	24.86	18.08
	Monthly Rent ($)	4966	5281	5209	5651
Mean of Other Cities	Entire home listings	60.03	76.84	85.72	145.99
	Price per night ($)	303.16	243.73	243.46	250.42
	Average ratings	4.74	4.73	4.73	4.76
	Duration (months)	21.25	21.19	21.60	14.18
	Monthly Rent ($)	2903	2981	3035	3251

Looking at the number of entire homes listed on Airbnb in Santa Monica in the earliest and latest time periods included in Table 1, respectively 695 units in April 2015 and 570 homes in July 2017, one may be led to think that the regulation was not very effective in reducing illegal vacation rental activity, reducing the number of listings in only 18%. However, when we compare with the same figures in Venice (arguably an area somewhat similar to Santa Monica) or with the mean of the rest of the cities, we realize that in both cases, the number of entire home listings advertised on Airbnb has increased significantly over the same time period, in Venice by around 70% and in the the rest of the cities by 143%. This fact indicates the importance of trying to use other cities as counterfactuals when estimating the impacts of the Home-Sharing Ordinance, since only looking at Santa Monica by itself may be misleading.

From the same table we can also note an interesting pattern in terms of the average price per night for an Airbnb entire home listing. In April 2015, data was scraped exactly in the Easter holiday weekend, which may (at least partially) explain the markedly higher prices comparing to other waves. If we focus on the other three periods included in Table 1 we note that both in Venice and in the remaining cities, the price did not fluctuate much, increasing in 2.86% for the latter and decreasing by 2.88% for the former. On the other hand, the average price per night in Santa Monica's Airbnb listings went through higher fluctuations. For example, from October 2015 to February 2016, the average nightly price in the city increased by approximately 22.06%. This price volatility relative to other cities is likely to be connected to the regulation. For example, the above mentioned price increase, not coincidentally, happens around the time in which the regulation starts to be actively enforced, leading many vacation units to exit Airbnb's platform, which according to a basic supply and demand framework, would lead to prices increases. We may note the same mechanism (although in the opposite direction) from February 2016 to July 2017, when the average price reduces as the number of entire home listings in Santa Monica increases. It is likely that the demand-supply mechanism explains only a small fraction of these price fluctuations, given that we also observe number of listings changing drastically in other cities, but only moderate price fluctuations. In fact, most of the price volatility in Santa Monica is probably connected to the changes induced by the regulation in the composition of listings active at a given point in time. For instance, just for the sake of argument, suppose listings owned by professional hosts (multiple-listing hosts) tend to be, on average, cheaper than home units owned by individual hosts (the ones who just own a single listing). Given that the regulation explicitly focus its enforcement efforts on professional hosts relative to individuals owning a single-listing, we should expect the average price in the city to go down, just as an result of the compositional change in the type of listings that remain active after the ordinance starts to be enforced.

The duration variable represents the average survival time of a listing (in months). Given that each listing is identified by a unique time-invariant number, I am able to determine when each listing appears and disappears from Airbnb. Thus, the duration is simply the difference between the last date in which I observe a listing and the first date in which that same listing was present on my dataset (in months). And the average duration displayed in Table 1 is the mean duration across all listings active in a certain location at a given point in time. And given the time span of my panel, the maximum duration a listing may have is around 34 months, from September 2014 to July 2017. By construction, it is very likely that the duration variable will be lower for periods very close to the last wave of the dataset. Take the extreme case of July 2017, the last year-month for which I have data. For all cities, the average duration among all listings active in that period is lower than the one computed for previous waves. This does not come from the fact that listings were indeed staying for a shorter period operating on Airbnb, but rather from the fact that there is constant entry of new listings, which by definition, if they first appear in my dataset t periods before the last wave, will have at most duration t. Alternatively, I could have considered for computing average duration only listings which are present in the dataset since the first wave. But this would not have a very interesting interpretation for the last wave either, since all listings surviving up to the last wave would have the same duration, 34 months. The interesting comparison here is for average duration in the same time period across Santa Monica and the other places. As expected, the average duration of entire home listings in Santa Monica is lower than the ones observed in Venice or in the rest of the cities.

When it comes to housing costs for residents, Table 1 shows that both Santa Monica and Venice are regions of high rents relative to the average across all the other cities. Interestingly, the rate of increase from April 2015 to July 2017 was approximately the same in the two neighboring cities. Long-term residential rents became 13.79% more expensive both in Santa Monica and Venice, which was higher than the percentage increase observed in the rest of the other cities, 11.99%.

4 Effect on Airbnb Listings

4.1 Alternatives for estimating the ordinance's causal impact

In this section I estimate the impact of Santa Monica's Home-Sharing Ordinance on the number of Airbnb listings active in the city. I focus on estimating the impact of the regulation on entire home listings only (I exclude shared home listings from all analysis). The reason for abstracting from shared homes is that, according to the regulation, as long as they are registered and pay taxes, they are legal home-sharing units, while an entire home being offered on Airbnb in Santa Monica after the passage of the Home-Sharing Ordinance, is definitely an illegal vacation rental.

Figure 2 plots the total number of entire home units listed on Airbnb in Santa Monica. Each vertical line represents one important action of the Home-Sharing Ordinance, namely adoption, proactive enforcement, and fine increase.[10] From the picture, it becomes clear that Airbnb entire home units in Santa Monica were increasing before May 2015, when the regulation was adopted and the educational campaign started. In April 2015, the last wave before the regulation's adoption, the number of entire home listings in Santa Monica was close to 700. Then, after the ordinance was adopted, there was a moderate decrease in the number of listings to approximately 600 homes. This was followed by a sharp drop in listings exactly in between the beginning of the proactive enforcement and the fine increase, when listings in the city were as low as 336. After that, the structure of the regulation remained unchanged and the number of listings in Santa Monica increased slightly until the end of 2016, when the number entire home listings jumped back to levels close to what they were in the summer of 2015.

Figure 2: Number of Airbnb Entire Home Listings in Santa Monica.

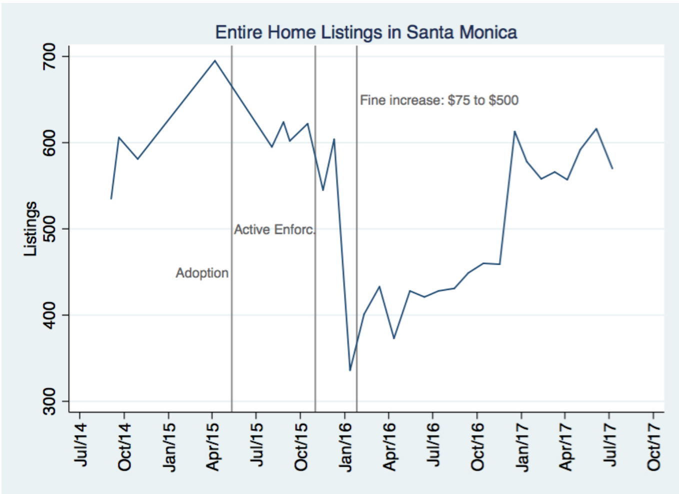

To accurately get the causal impact of the Home-Sharing Ordinance on the number of Airbnb homes operating in Santa Monica, the most important task is to get an estimate of the counterfactual. That is, we need to suggest a method to obtain a measure of what would have Santa Monica Airbnb market looked like in the absence of the regulation.

One way to approach this problem could be to look at the number of Airbnb listings in Santa Monica before and after the ordinance. However, this approach only gives unbiased estimates of the true impact of the regulation under strong assumptions about how the number of listings in the city would have increased over time. Another way to try to measure the causal impact of the regulation on the number of listings, would be to choose a control city comparable to Santa Monica and assume that, in the absence of the regulation, listings in the two cities would have followed a similar trend. The issue with this approach is that the choice of which city constitutes a good control is ambiguous. Moreover, there is a trade-off between geographical proximity (thus some

[10] Hereafter, in any plot where three vertical lines are depicted, they represent these three events in the regulation's structure: adoption, proactive enforcement, and fine increase respectively. I choose not to graphically depict the date in which the regulation became effective (June 12th, 2015) because of its proximity to the regulation's adoption and due to the fact that I have no wave of data collection in between these two events.

degree of similarity) versus danger of being exposed to general equilibrium effects. To illustrate this idea, consider Figure 3, which plots entire home units listed on Airbnb for Santa Monica and a nearby city, Venice.[11] Choosing Venice to be the control city has the advantage that it is arguably very much alike Santa Monica and therefore it is likely that Airbnb listings in the two cities would have evolved in a similar way if it was not for the regulation. On the other hand, it is also possible that a full prohibition on vacation rentals in Santa Monica reallocates some of the tourists that would otherwise want to stay in the city to other nearby areas. That would raise the demand for vacation rental units in Venice and thus the regulation would cause the number of Airbnb listings in this city to grow. In that sense, Venice is not a proper control unit, since it may be indirectly affected by the regulation. Particularly, to the extent that the ordinance causes Airbnb listings to increase in the immediate surroundings of Santa Monica, using one of these direct neighboring areas as a control unit would deliver upward biased estimates of the impacts of the regulation on listings.

Figure 3: Airbnb Entire Home Listings in Santa Monica and Venice.

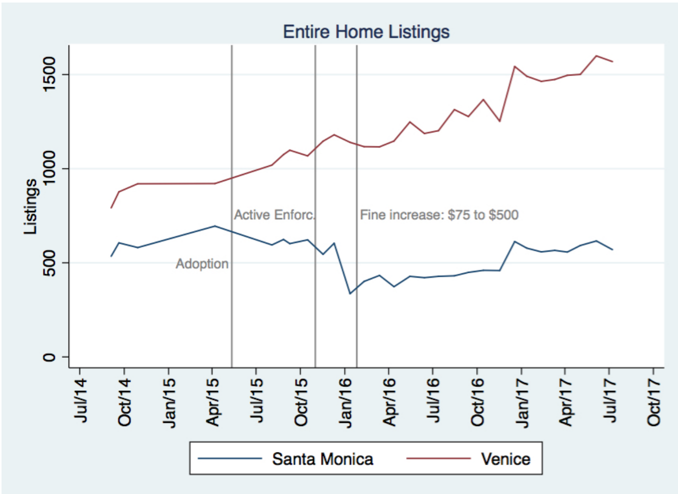

4.2 Synthetic Control Method (SCM)

Given the downsides of the methods outlined in the last section, a more adequate empirical strategy to estimate the causal impact of the ordinance on the number of Airbnb listings in the city of Santa Monica is to use a Synthetic Control Method approach. This method was laid out in a series of papers by Abadie and coauthors and it is particularly adequate to estimate impacts of policy interventions taking place at the aggregate level (city, state, country). What follows below is based on the exposition of the Synthetic Control Method contained in Abadie, Diamond, and Hainmueller (2010).

The idea behind this method is that the combination of many potential control units does a better job reproducing the treated group than any control unit is able to do in isolation. In the context of this study, this is to say that a combination of many cities potentially similar to Santa Monica is more efficient in approximating its number of Airbnb listings in the absence of the regulation than any single control city.

[11] Technically, Venice is a neighborhood of the city of Los Angeles rather than a city. However, for clarity of exposition and to avoid having to repeatedly name different terms for geographical areas (city, neighborhood, district, incorporated town and so on), I hereafter loosely refer to either one of these administrative units with the term "city" whenever they have an Airbnb market which is large enough for them to be considered as separate spatial units of analysis.

More formally, let $j = [1, 2, ...J + 1]$ represent each city in the Los Angeles County and $t = [1, 2, ...T_0, ...T]$ each wave of the panel dataset on Airbnb listings. Denoting Santa Monica as city $j = 1$, the SCM estimated effect of the Home-Sharing Ordinance at each time period t after its adoption is given by the difference between the observed number of listings in Santa Monica and an weighted average of the number of listings observed in control cities

$$\alpha_{1t} = Y_{1t} - \sum_{j=2}^{J+1} w_j Y_{jt} \tag{1}$$

where the weights assigned to each control city are chosen to minimize the following expression

$$\sum_{m=1}^{k} v_m (X_{1m} - X_{0m} W)^2 \tag{2}$$

in which $W = [w_2, ..., w_{J+1}]'$ is a vector containing the weights of each one of the J control cities, X_{1m} is the observed value for a given predictor m of the number of listings in Santa Monica for periods before the regulation, $X_{0m} = [X_{2m}, ..., X_{J+1m}]$ is a vector containing the values of the same predictor m for all control units, and v_m is the weight assigned to predictor m, to reflect the idea that each predictor has a different importance in explaining the number of listings in a given city. Intuitively, what the method does is to choose the weighted average of other cities that best approximates Santa Monica before the Home-Sharing Ordinance in terms of k predictors of the outcome variable, number of Airbnb entire home listings. Crucially, these predictors may include any linear combination of the outcome variable itself, for instance, average pre-regulation entire home listings.

4.3 SCM to estimate effect of the Home-Sharing Ordinance

This section describes in detail the implementation of a SCM approach to estimate the extent to which the Home-Sharing Ordinance has reduced the number of entire homes listed on Airbnb in the city of Santa Monica. Table 2 summarizes the study.

Table 2: Summary of the Study

Treated Unit	Santa Monica
Potential control units	Cities (geographical units) within the Los Angeles County
Outcome variable	Airbnb Entire Home Listings
Predictor Variables	Income, population, share of young adults, establishments in food and accommodation
Date of Regulation	May 12th, 2015
Pre-regulation period	Sep/2014 to Apr/2015 (in 4 waves)
Post-regulation period	Aug/2015 to Jul/2017 (in 25 waves)

One potential downside of using the Synthetic Control Method here is that I don't have a very long pre-regulation period. Ideally, I would have a longer series of observed outcomes before the Home-Sharing Ordinance so that I could be more confident that the synthetic Santa Monica is able to reproduce the series of the actual Santa Monica for an extended period of time. Nonetheless, given the already mentioned drawbacks of the alternative estimation methods, I still judge the

SCM to be the best that can be done to accurately estimate the impacts of the ordinance on the number of entire home units active of Airbnb in Santa Monica.

4.3.1 The pool of controls cities

The analysis starts with the definition of the cities that will be considered as potential control units. Or, as it is commonly said in the literature, the definition of the relevant pool of donors. In the context of the present study, the goal is to only consider cities that arguably have some degree of similarity to Santa Monica in terms of a short-term rental activity. Table 3 presents the summary of the main steps taken in the process of getting to the final pool of relevant control cities.

Table 3: Steps in the process of defining the pool of relevant control cities

Step	Result
Initial pool of potential controls units	113 cities (geographical units) for which I have Airbnb data
Eliminate cities for which I don't have listings in all waves	48 cities remain
Eliminate cities which also implemented STR regulations	40 cities remain
Eliminate cities that may be indirectly affected by the HSO (spillover effects)	38 cities remain
Rank cities in terms of pre-regulation mean number of listings	Keep cities above the median pre-regulation listings
Final pool of donors	19 control units

I start by considering as potential control units all the 113 cities in the Los Angeles county for which I have Airbnb listings.[12] The initial restriction I make to the pool of donors is to exclude all cities for which I have missing information on the total number of entire home listings in any of the 29 waves of data collection. The reason for that restriction is that the SCM, although being flexible enough to deal with irregularly spaced waves of data collection, it does require the final panel, including treated unit and the pool of donors to be balanced (Abadie, Diamond, and Hainmueller 2010). That is, it requires Santa Monica and every control city to have valid information on the outcome variable (listings) for all waves of data collection.

The second restriction I make to the pool donors is to eliminate cities in which strict regulations on short-term rental activity was also undertaken. This is necessary because the SCM methodology uses post-regulation listings in control units to build the counterfactual to which Santa Monica should be compared to. Therefore, if control cities undergo any severe shock in listings in the post-regulation period, the estimated synthetic Santa Monica would reflect these shocks and the estimated effect of the regulation would be biased. To illustrate this idea suppose I don't exclude from the pool of donors Redondo Beach, which in March 2016 banned short-term rentals of residential units.[13] The number of Airbnb listings in this city will likely reduce after the ban, and if Redondo Beach receives a non-negative weight in the SCM, the synthetic Santa Monica will reflect

[12] Again, I stress that cities here is being used very loosely. In fact, some geographical units may be formally considered districts of the city of L.A., neighborhoods, incorporated towns, etc. I use the word cities to avoid long and confusing sentences every time I mention the control units.

[13] Source: http://www.scpr.org/news/2016/03/03/58174/redondo-beach-cracks-down-short-term-rentals/.

this negative shock on Airbnb listings in Redondo Beach, producing an estimated effect of the Home-Sharing Ordinance lower than that the true effect really is.

Then, I also discard cities which have direct borders with Santa Monica. This exclusion is important because the SCM assumes that the policy intervention under study has no relevant impacts on any of the control units. And to the extent that the Home-Sharing Ordinance may have generated spillover effects on cites very close to Santa Monica, including them in the pool of control units would be problematic. For instance, suppose I keep Venice in the pool of donors and imagine that the outright prohibition to short-term rentals in Santa Monica causes the number of Airbnb listings in Venice to increase. In this case, if post-regulation Venice listings are used to build the synthetic Santa Monica, the SCM would overstate the true effect of the ordinance. After removing the neighboring cities of Pacific Palisades and Venice, I am left with 38 potential control units.[14]

Lastly, given that Santa Monica is a huge Airbnb market, featuring before the ordinance a total number of entire homes listed in the online platform only lower the city of L.A. itself and Venice (already excluded because of potential spillover effects), I rank cities in terms of their pre-regulation mean number of listings and only keep cities above the median rank, which leaves me with the 19 largest Airbnb markets before the ordinance. Here, the question of whether this restriction to the pool makes sense may arise. The reason for taking this step is to try to have in the final pool of control units, only cities that are not extremely different than Santa Monica with respect to the main outcome variable. There is a trade-off between the capacity of the synthetic Santa Monica to fit the trend in number of listings before the Home-Sharing Ordinance and the risk of over-fitting the trend by using cities that have little to do with Santa Monica in terms of the determinants of short-term rental activity. If I keep the all the 38 cities in the pool of control units, it is likely that the pre-regulation trend in entire home listings in Santa Monica will be better approximated by the synthetic version of the city than if I only have 19 control cities. On the other hand, if the very small cities are kept in the pool and they have little to do with Santa Monica with respect to short-term rental activity, there is the risk of over-fitting Santa Monica's trend and having a synthetic version of the city that, after the regulation, will not display the behavior that actual Santa Monica would have displayed in the absence of the ordinance. In order to have a better understanding of how each one of control units is contributing to the synthetic Santa Monica and to avoid simply fitting on the listings' trend, I choose to stick to the smaller pool of control cities.

4.3.2 Constructing the synthetic Santa Monica

Given the final pool of 19 control cities, the next step is to construct the synthetic Santa Monica. To make more meaningful comparisons and avoid obtaining a statistically significant effect of the regulation exclusively due to the large size of Santa Monica's Airbnb market,[15] I normalize the number of entire homes listings by the number housing units in each city. Thus, the outcome variable of interest is expressed as entire home units listed in Airbnb per 1000 housing in a given

[14] I only eliminate Venice and Pacific Palisades because the other three geographical areas neighboring Santa Monica, namely Brentwood, West Los Angeles, and Mar Vista are classified in my dataset as part of the city of L.A., which for its very large size relative to other cities, I assume the spillover effects of the Santa Monica Home-Sharing to be marginal.

[15] To evaluate the statistical significance of the results, the SCM methodology suggests running placebo analysis as if each one of the control cities was the treated unit and then comparing the estimated placebo effects with the ones estimated to Santa Monica. If the effect for Santa Monica is markedly higher than the ones estimated for other cities (placebos), one may conclude that the effect of the regulation is statistically significant. Given that Santa Monica has a large number of Airbnb listings compared to most cities in the pool of control units, a change in absolute number of Airbnb listings in Santa Monica may look very different than changes in other cities more due its market size rather than a true discrepancy in a relevant measure of the magnitude of the effects.

city[16].

Figure 4 plots the evolution of normalized Airbnb entire home listings for Santa Monica as well as for the average of all other cities in the pool of control units. An initial indication that the impact of regulation in reducing the total number of entire home listings in Santa Monica was substantial is that after all the regulation actions (adoption, active enforcement, and fine increase) were in place, listings in the city were roughly equal to the average over all control units, in spite of being approximately three times higher in April 2015, before the Home-Sharing Ordinance was enacted.

Figure 4: Airbnb entire home listings per 1000 housing units in Santa Monica and in the rest of cities.

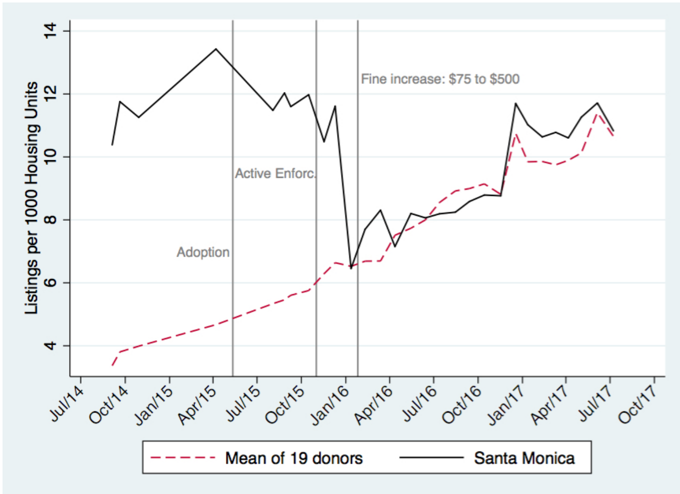

The main question, and one that cannot be answered by looking at Figure 4, is how listings would have evolved in Santa Monica in case there had been no law restricting short-term rentals. To get an estimate of this hypothetical situation I construct a synthetic Santa Monica according the ideas laid out in Section 4.2.

First, the set of predictors that could explain the number of Airbnb listings per housing unit in a city has to be determined. Entire home units listed on Airbnb may be predicted by income, population, the percentage of the population constituted of young adults (between 20 and 34 years old), the percentage of housing units occupied by owners versus renters, the number of establishments in food and accommodation, and the number of establishments in arts and entertainment. Additionally, as it is usually done in most applications of the Synthetic Control Method, I include the mean of the outcome variable itself as a predictor. That is, one of the predictors is the mean number of normalized entire home listings before the period when the Home-Sharing Ordinance was implemented (mean over the four waves pre-regulation). Figure 5 shows the actual and estimated synthetic Santa Monica using different sets of predictor variables.

In the upper-left graph, the only predictor used is the mean normalized listings across pre-regulation periods. As we move right and down, each graph plots the synthetic Santa Monica obtained using, in addition to pre-ordinance mean listings, some extra predictor variables. The top-right graph depicts the result of the SCM when adding income and population (in logs) as predictors. In the bottom-left picture, I add the fraction of the population between 20 and 34 years old and the share of occupied houses inhabited by owners rather than renters. And lastly,

[16]Hereafter, when referring to the main outcome of interest in this analysis, to avoid repetitively using the overly long "entire home listings per 1000 housing units", I may occasionally just write "listings", which should be understood as the normalized version of the variable.

the bottom-right graph plots the result of the SCM when I further add the number of establishments in food/accommodation as well as the number of establishments in arts and entertainment. The main message is that, although there is some degree of variation in the quality of fit before the regulation, the long-run trend of the synthetic version of Santa Monica does not change much by including more or less predictor variables.

Figure 5: Actual and synthetic Santa Monica using different predictor variables.

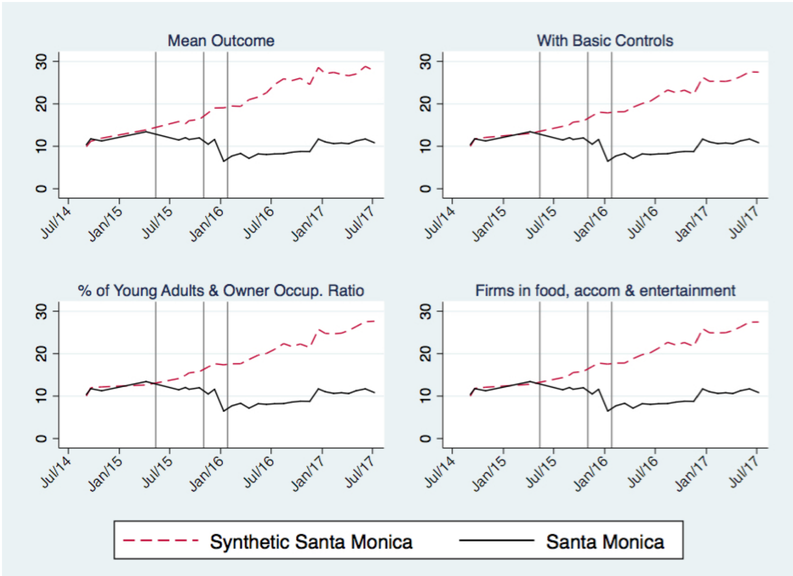

One can directly see that adding extra predictor variables does not change the synthetic Santa Monica in any significant way by looking at Figure 6, which plots in the same graph these different versions of the synthetic city. The specification that uses as predictor only the pre-regulation mean listings is always a little bit above the other synthetic versions of the city. The other three synthetic versions of the city produce very similar counterfactual scenarios for analysis. One thing that is interesting to note is that when I use mean listings as the only predictor variable, the synthetic Santa Monica is a weighted average of all 19 controls cities. That is, all control cities receive a positive weight in building the synthetic versions of the regulated city. This is because the SCM is simply using all the cities' outcome to fit as best as it can the trend in normalized listings in Santa Monica before the Home-Sharing Ordinance. However, when I include the other predictors that correlate with Airbnb listings (the already mentioned predictor variables), many control cities receive a weight of zero in the estimation procedure and the synthetic Santa Monica is, in practice, a weighted average of less cities. This is expected, since by using this extra information, the SCM will tend to choose cities that approximate Santa Monica before the ordinance not only in listings but also in terms of these other predictors. Following what is suggested in Abadie, Diamond, and Hainmueller (2010), I try not to exclusively fit the outcome variable without considering other relevant characteristics and, for that reason, I choose to keep all the predictor variables in my preferred specification.

Furthermore, given the importance attributed to pre-regulation mean listings (very high weight relative to other predictors), it is crucial to define a systematic way to choose the final specification in terms in which linear combination of pre-ordinance listings to use. Since the SCM allows for the use of any kind of linear combination of the observed pre-regulation outcome variable itself, one could use many different specifications, such as the mean outcome, the even periods outcome, the outcome observed in the last period before the event of interest, etc. Due to this huge flexibility

Figure 6: Comparison of synthetic versions of Santa Monica using different predictor variables.

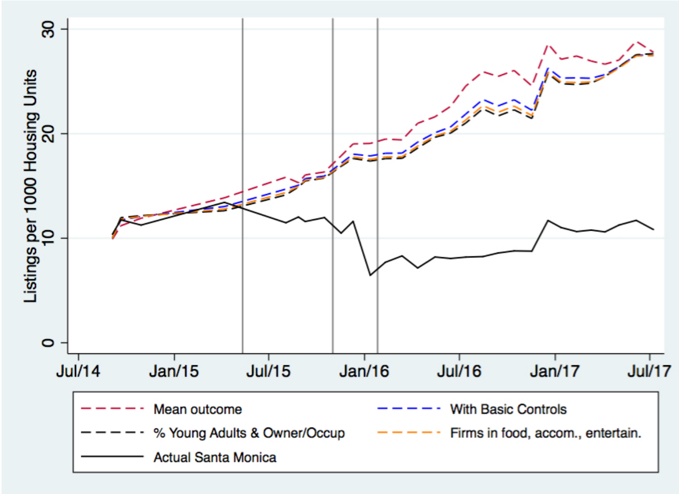

in the hands of the researcher combined with the incentive to find significant results, it arises the worry that one may not be able to commit and ends up engaging in specification search and cherry picking for the specification that delivers the most "desirable" results for the researcher (Ferman, Pinto, and Possebom 2018).

In order to avoid this issue, I run the SCM using different types of linear combinations of the pre-regulation normalized listings as one of the predictor variables and before carrying out any analysis of the statistical significance of the estimated effect of the regulation, I follow Ferman, Pinto, and Possebom (2018) and choose the specification that minimizes the post-regulation root mean squared prediction error (RMSPE) for the control units. In order to do that, for every specification, I run the SCM using each control city as if it were actually the one going through the Home-Sharing Ordinance and I get a synthetic version of each city. Then, I choose the specification that reduces the out of sample prediction error for the control units. The intuition behind this strategy is that, assuming that the control cities did not actually go through any major event throughout the period of analysis, the synthetic control method should be able to efficiently approximate the actual observed trend for the control cities after the regulation. Figure 7 illustrates the product of the process just described for three cities as well as for Santa Monica. It plots the results for five different linear combinations of the pre-regulation normalized listings as one of the predictor variables. While each specification varies in terms of the particular way in which it uses the main outcome variable itself as a predictor of normalized listings, all of them have the same supplemental predictors, namely log income and population, fraction of young adults and owner occupiers, and log number of establishments in food, accommodation, and entertainment.

In terms of the pre-ordinance normalized listings, I test the mean across all pre-regulation periods, the mean over the last two waves before the regulation, the observed outcome in the last pre-regulation period, the observed values for the last two waves before the ordinance, and lastly, the observed outcome in the even periods only. In Figure 7, we may note that in spite of the low number of pre-regulation waves available (four only), all the specifications produce synthetic versions that track post-ordinance trends reasonably well for Marina Del Rey, North Hollywood, and Los Angeles,. Clearly, for Santa Monica, the picture looks completely different, with all the synthetic versions of this city evolving in a very different way than the observed listings in the actual city. Regarding the choice of specification, one cannot really decide anything only by looking at the picture. Again, using the rule of favoring the specification that minimizes the post-regulation prediction error for control units, the chosen one was the specification that uses the mean of the normalized listings over the last two waves before the ordinance was adopted (besides the already mentioned supplemental predictors).

Figure 7: Specifications with different linear combinations of pre-ordinance listings.

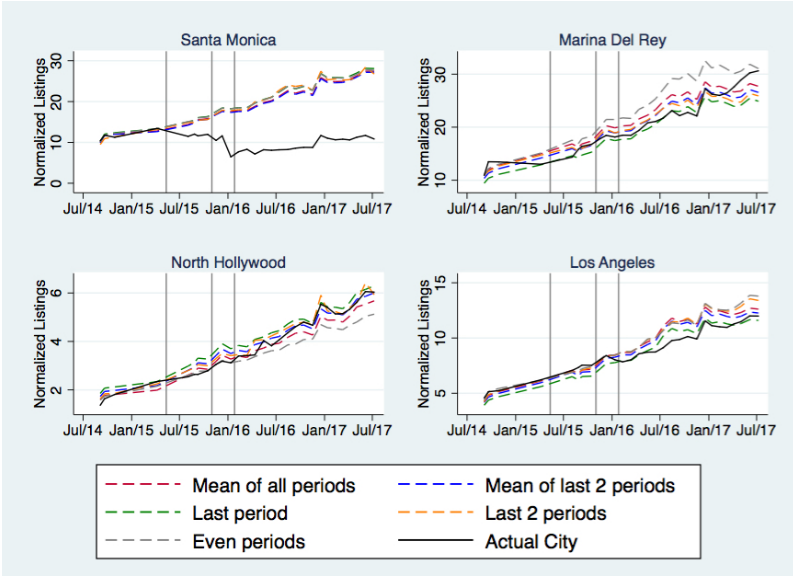

4.3.3 Estimated impact on normalized listings: Results

Following the method explained above, I construct a synthetic Santa Monica as the weighted average of cities within the pool of control units that best approximates the evolution of listings in the city before the Home-Sharing Ordinance was adopted. Table 4 summarizes the results and compares the pre-ordinance characteristics of actual Santa Monica, its synthetic version, and the average of all 19 control cities.

As expected, the synthetic version of Santa Monica is much closer to the actual city than the mean of all control cities is. The difference is most striking when it comes to the mean of normalized Airbnb listings over the last two waves before the ordinance. While the average across all 19 control cities displays approximately one third of the observed value in Santa Monica, the convex combination produced by the SCM was able to virtually match the 12.34 normalized entire homes units in the city under study. The pattern is similar, although less pronounced, for the other predictors. However, for some variables the fit between Santa Monica and its synthetic counterpart is not as good. That is related to the weight (v_m in Equation 2) each predictor receives in the estimation method. The higher the weight a predictor receives, the more important it is in predicting entire home listings, thus the more the SCM will prioritize approximating Santa Monica and the synthetic control in terms of that predictor (Abadie, Diamond, and Hainmueller 2010). For instance, establishments in arts and entertainment received the lowest weight, therefore the difference between Santa Monica and its synthetic counterpart is "larger" when it comes to these predictors. The weight assigned to each predictor is determined by a data driven process where a cross section regression of normalized entire home listings on all the suggested predictors is carried out for each wave before the regulation. Then, the final weight each predictor receives is related to how much of the variation in listings across cities can be explained by each predictor. Intuitively, the higher the regression coefficient of a predictor in these cross-section regressions, the higher the weight the SCM will assign to that specific predictor (Galiani, Quistorff, et al. 2016).

Table 4: Airbnb Listings Predictor Means

Variables	Santa Monica		Average of 19 controls cities
	Real	Synthetic	
Mean Listings (Oct/2014 and Apr/2015)	12.34	12.31	4.33
Ln(Median Income)	11.28	11.27	11.18
Ln(Population)	11.45	11.41	11.02
Share of young adults	25.48	25.46	22.18
Owner-occupancy rate	0.28	0.35	0.47
Ln(Establishments Food & Accommod.)	6.17	5.76	4.98
Ln(Establishments Arts & Entertain.)	6.60	5.86	5.06

Note: All predictor variables except normalized entire home listings are averaged over the four pre-regulation waves of data collection. Normalized entire home listings are averaged over the last two pre-regulation waves.

As mentioned before, one advantage of the SCM is that it makes it explicit the contribution of each control unit to the final synthetic control used as counterfactual. Table 5 shows all the cities and the weights assigned to each potential control unit. Although we have 19 cities in the pool of donors, Santa Monica's trend in Airbnb entire home listings per 1000 housing units is best represented by a convex combination of just three of them, Marina Del Rey, Los Angeles, and Topanga. All other cities received weight zero. This is common in applications of SCM, which only assigns non-negative weights to control units that are reasonably similar to the treated unit in the most relevant predictor variables. For instance, Abadie, Diamond, and Hainmueller (2010), compute a synthetic control in order to mirror California's trends in per capita cigarette consumption using the states of the U.S. as control units and the method only assigns non-negative weights to 5 out of 38 potential control states (some of the 50 U.S. states were not even included in the donor pool for being exposed to external shocks over the period of analysis).

Table 5: City weights in the Synthetic Santa Monica

City	Weight	City	Weight
Altadena	0	Playa Del Rey	0
Burbank	0	San Pedro	0
Culver City	0	Sherman Oaks	0
Encino	0	South Pasadena	0
Glendale	0	Studio City	0
Long Beach	0	Topanga	0.129
Los Angeles	0.351	Valley Village	0
Malibu	0	Van Nuys	0
Marina Del Rey	0.520	Woodland Hills	0
North Hollywood	0	Santa Monica	-

Figure 8 displays Airbnb entire home listings per 1000 housing units in Santa Monica and its synthetic counterpart from September 2014 to July 2017. Differently than what we observe in Figure 4, where the simple average of all control cities had a pre-regulation trend in listings very far from the observed in the city of Santa Monica, the synthetic version of the city resembles much more the observed trend. It is true that even the synthetic Santa Monica does not reflect perfectly the pre-ordinance evolution of normalized listings, displaying a lower level of normalized listings right before the regulation adoption. Nonetheless, the method provides a hypothetical city that is much closer to Santa Monica's trend than any other city in isolation. Moreover, the fact that synthetic Santa Monica has lower (rather than higher) levels of listings just before the regulation compared to the actual city, suggests that I obtain conservative estimates of the Home-Sharing Ordinance causal effect on the number of listings operating in the city, specially in the time periods shortly after the regulation adoption.

Figure 8: Trends in normalized entire home listings in Santa Monica and its synthetic version

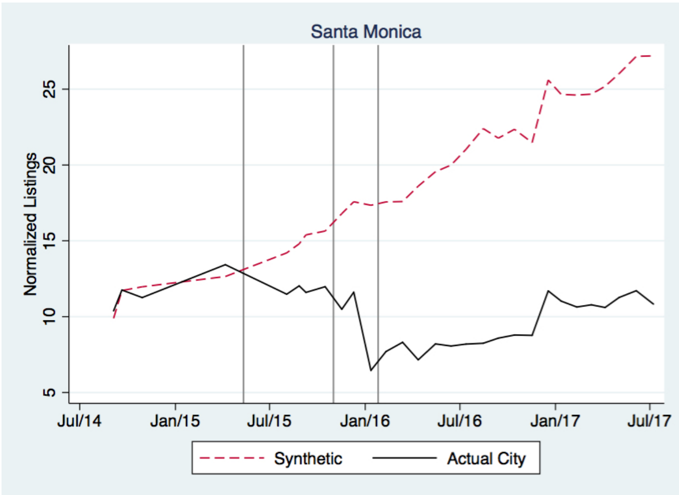

My measure of the regulation's impact on Airbnb entire home listings is simply the difference between the black solid line reflecting the observed listings in Santa Monica and the red dashed line representing the synthetic version of the city, which is my estimate of how Santa Monica's listings would have evolved in the absence of the regulation. Promptly after the first vertical line represented in the picture, which represents the regulation's adoption, the two lines start to diverge sharply. The SCM is suggesting that, if the Home-Sharing Ordinance had not been adopted, entire home listings in Santa Monica would have kept its increasing pre-ordinance trend. That comes from the fact that, differently than what happened in Santa Monica, in the cities that constitute the synthetic counterpart of the regulated city, listings kept on increasing after May 2015.

As mentioned before, the SCM delivers one point estimate of the effect of the regulation for each time period after the regulation adoption. Table 6 brings the results for three of these post-regulation time periods. It includes the estimate of the impact for October 2015, three months after the regulation had been adopted but neither active enforcement or the fine increase had taken place; for January 2016, around 2 months and a half after the proactive enforcement had been implemented but before the fine had been increased; and for July 2017, the last period of my sample and after the structure of the regulation was stable for one year and half. Besides displaying the effect of the ordinance in terms of entire home listings normalized by housing units, I also include the results in percentage terms and I use the total number of housing units in Santa Monica to transform back the outcome variable to absolute number of listings.

The first point worth mentioning is that the regulation had a substantial impact of the number of entire home listings operating in Santa Monica. A little over two years after the ordinance's adoption, the number of listings per 1000 housing in Santa Monica was around 16 units lower than it would have been. A more intuitive interpretation can be derived by looking at absolute number of Airbnb entire home listings in the city, which by the same time period, were approximately 861

Table 6: Estimated Effect of the Ordinance on Entire Home Listings

Variable	October 2015	January 2016	July 2017
Listings / 1000 housing units	-3.67	-10.89	-16.36
Number of Listings	-190.66	-566.68	-861.20
In percentage terms (%)	-23.46	-62.78	-60.17

units lower comparing to a hypothetical situation in which Santa Monica had not been regulated. In plain words, 26 months after the ordinance was adopted, it had reduced the number of entire homes listed on Airbnb in Santa Monica by around 861 units.

Moreover, the reduction in the number of entire home listings in Santa Monica has been large even just a few months after the adoption of the regulation. In October 2015, only three months after the adoption, and before any proactive enforcement had been taken, the total number of entire home listings had already been reduced by almost 191 units, a non-trivial 23.46% decline in listings. This suggests that just by disseminating an educational campaign about the new law and sending warning letters to potential violators, the city was already able to reduce the number of listings by almost one fourth of what they would have been.

When we look at the effect by January 2016, figures become even more striking. It shows that eight months after the passage of the Home-Sharing Ordinance and two months and a half following the introduction of the proactive enforcement strategy, the city of Santa Monica had 567 entire home units on Airbnb less than what it would have had in the absence of the ordinance, an impressive 62.78% reduction in listings in the city. This suggests that the proactive enforcement was very effective in reducing listings operating in the city, since although there are only three months separating October 2015 and January 2016, the estimated effect of the ordinance on total entire home listings almost tripled. One could also argue that this strong reduction in listings by January 2016 is not only a result of the proactive enforcement strategy, but is also mixed with an anticipation effect in which hosts already expected that the fine would suffer a substantial increase by the end of that month and this may have intensified the incentives to stop operating an illegal vacation rental. I argue that this was not the case, because the fine increase was decided during a City Council meeting on January 26th, 2016 (after January's wave of data collection, which was on January 12th, 2016) and there is no indication that people knew the fine would raise by that amount. And even if they expected the fine to increase, there is no obvious reason why rational hosts would stop illegal operation of vacation rentals as a causal response to a higher fine before the fine is actually increased. Therefore, I interpret the massive reduction in entire home listings by January 2016 mostly as a causal effect of the proactive enforcement. I hypothesize that through an increase in the number of effective enforcement actions taken against violators (pro-actively issuing fines), the Short-Term Rental Team was able to force fined entire home units out of the market as well as to induce other non-fined hosts to also stop operating illegal vacation rental units after perceiving a real threat of being fined. At the moment, this is just a hypothesis and with the analysis shown here I cannot provide sufficient evidence to prove that this was the mechanism driving the sharp decrease in listings in January 2016. Further research investigating the specific mechanisms through which the regulation has reduced the number of vacation rental units operating in Santa Monica is needed.

Lastly, it is interesting to note that although the estimated effect on the total number of entire home units listed on Airbnb in Santa Monica kept increasing (in magnitude) from January 2016 to July 2017, the effect in percentage terms did not change much, and in fact, it has slightly dropped.

This can be visually verified in Figure 8, where we see that after the sharp drop in listings in between the start of proactive enforcement and the introduction of the higher fine, Santa Monica's entire home listings increase until the end of the period of analysis. When compared to the increase in listings in its synthetic counterpart, the increase in Santa Monica was markedly lower in the number of normalized listings, but it actually was slightly higher in percentage terms. More specifically, normalized listings in Santa Monica from January 2016 to July 2017 have increased from 6.46 to 10.83, a 67.65% increase, whereas in the synthetic version of the city listings went from 17.34 to 27.19, a 56.79% increase.

4.3.4 Inference about the effects of the Home-Sharing Ordinance

The Synthetic Control Method allows for statistical inference about the estimated effects of the policy intervention of interest through running a series of placebo analysis. In the context of the present paper, following Abadie, Diamond, and Hainmueller (2010), I hypothetically assume that the Home-Sharing Ordinance took place in each one of the cities in the control group and estimate placebo effects via the SCM. Then, I compare the estimated effects for Santa Monica versus all the other placebo effects. The idea is that if the effects of the Home-Sharing Ordinance are statistically significant, then they will be markedly higher than the estimated placebo effects.

The first inference exercise I run consists in plotting the estimated effects for Santa Monica versus the placebo effects for the cities in the control group. Figure 9 includes four graphs depicting the gap between the actual observed normalized listings in a city and the normalized listings estimated for its synthetic counterpart. The top-left graph includes all the 20 cities, Santa Monica

Figure 9: Gap in listings per 1000 housing units: actual versus synthetic

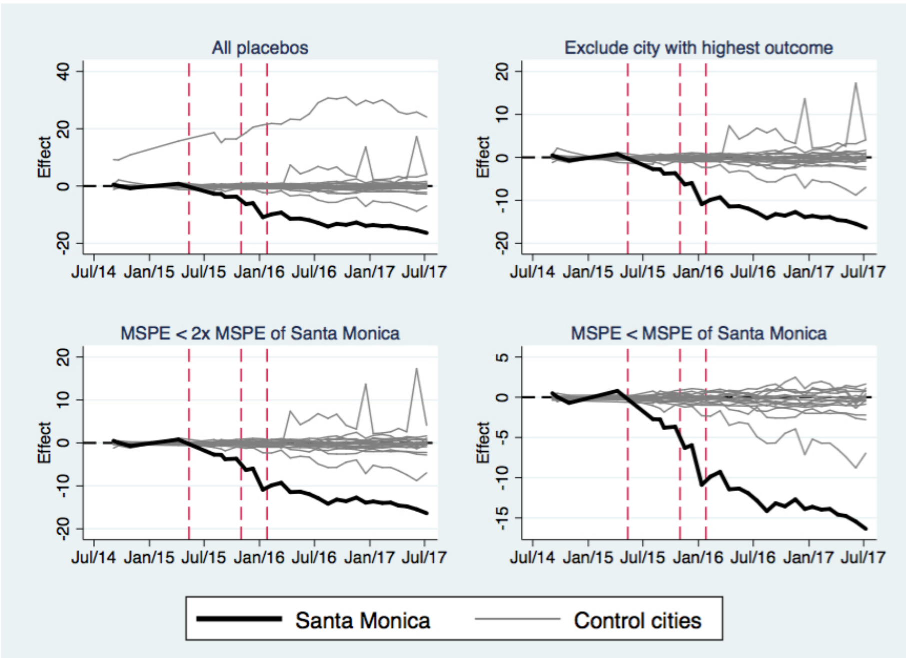

and the 19 control units. We immediately notice that there is one city's trend which cannot be represented by a convex combination of the other ones. This happens by construction, because there will never be a convex combination of cities that will mirror the trend of the city that has the highest level of the outcome variable among all cities in the pool. Given that out goal here is to compare the gaps between actual and synthetic city post-regulation across different cities, it makes little sense to include in the comparison cities for which the pre-regulation fit was not good either, that is, cities that even before the regulation had a large gap between its actual trend and its synthetic counterpart.

Thus, each plot in Figure 9 excludes cities for which the pre-ordinance fit was worse than some specified threshold. The top-right graph excludes Topanga, the one city that had its pre-regulation gap far from the zero horizontal line in the first plot. Already in that picture, we can notice that the magnitude of the estimated gap for Santa Monica is much higher than for any other city. It is interesting to note how the gap (or the effect) for all cities is close to zero before the first vertical line, which represents the adoption of the Home-Sharing Ordinance. This says that, excluding Topanga, the pre-regulation trend in listings in all cities could be well approximated by a convex combination of the other ones. Then, after the regulation's adoption, the lines evolve in a very different way. Santa Monica's gap immediately becomes negative, while the line representing other cities stays roughly close to zero for some months. Although the gap departs from zero for some cities around January 2016, in magnitude they are still smaller than the gap for Santa Monica. The fact that the gap is very different from zero for a few cities reflects the importance of carrying out statistical significance analysis for the estimated effect of the ordinance because even though other cities did not go through any major event that would change their number of Airbnb listings, there is a possibility that the SCM, by chance, estimates an effect where in fact there is none.

In order to compare Santa Monica's effect only with cities that had a good fit in the pre-regulation trend, in the bottom-left plot I further exclude Marina Del Rey, which had a pre-regulation mean squared prediction error (MSPE) higher than the double of the one obtained for Santa Monica. Lastly, in the bottom-right picture I only consider cities that had a pre-ordinance prediction error equal or lower to the one obtained in Santa Monica, which leads me to further exclude Malibu. In that plot, it becomes evident that while pre-regulation gaps were similar and close to zero for all cities, post-ordinance Santa Monica's gap stands out as the one that most strikingly departs from zero. Given that this plot includes 17 cities, if one were to assign the Home-Sharing Ordinance at random in the data, the probability of estimating an effect as large as the one for Santa Monica is $1/17 = 0.0588$.

Yet another way to evaluate the estimated effect for Santa Monica relative to the placebo effects estimated for control units is to compare post and pre-regulation measures of fit across different cities. Figure 10 plots the ratios of post/pre-ordinance Root Mean Square Prediction Errors (RMSPE) for all the 20 cities considered in the study. The hypothesis is that if the estimated effect for Santa Monica is significantly higher than the placebo effects estimated for control cities, then the ratio of post/pre-regulation predictor errors should be the highest for Santa Monica. Intuitively, the idea behind this test is that for Santa Monica, where I argue there is a true effect to be estimated, the divergence of the post-ordinance synthetic control and the observed listings should be higher than for the other cities. Dividing these post-regulation prediction errors by the pre-regulation ones corrects for the issue of comparing Santa Monica to cities that were poorly approximated by a convex combination of others and has the advantage of not requiring the exclusion of any city for meaningful comparisons (Abadie, Diamond, and Hainmueller 2010). Figure 10 shows that the highest ratio is the one for Santa Monica. And given that this picture does not exclude any of the 20 cities considered, the probability of estimating a ratio as large as the one for Santa Monica under a random permutation of the regulation in my data is $1/20 = 0.05$, or 5%.

Figure 10: Ratio of post-regulation RMSPE and pre-regulation RMSPE

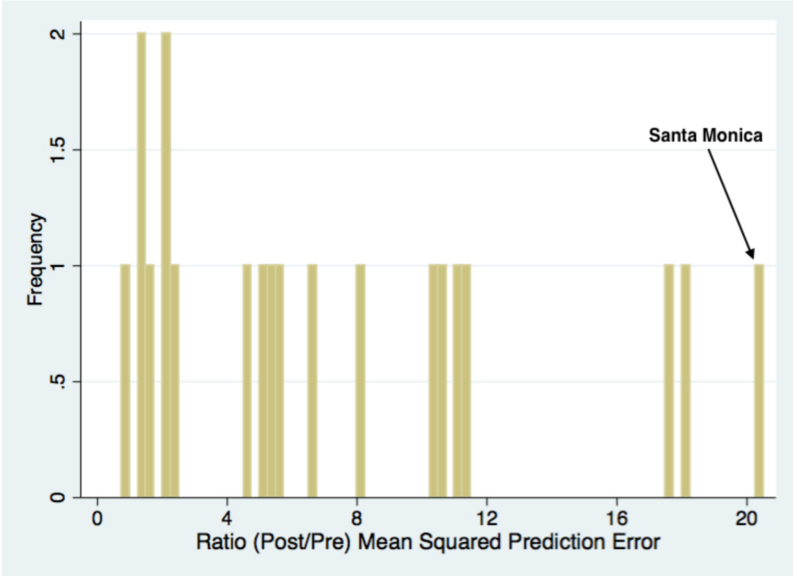

5 The Effect of the Ordinance on Long-Term Rents

In this section I try to evaluate whether the Home-Sharing Ordinance significantly impacted the long-term rents faced by Santa Monica's residents. Given that in the last section I find significant effects of the regulation in reducing the number of homes being used for vacation rental through the platform Airbnb, the natural hypothesis to be tested is whether this decrease in Airbnb listings has also translated into a reduction in the long-term rents.

Moreover, Santa Monica's City Council, when adopting the ordinance in May 2015, explicitly stated that one of the main reasons for restricting vacation rentals was related to increasing rents for the city's residents. This is clear in the following passage extracted from the an official report by the Short-Term Rental Team:[17] "The ordinance was passed to ensure that residential rental housing remains available to long-term tenants, and because short-term rentals have undesirable impacts that threaten the stability and character of the City's neighborhoods and result in increased rents".

To investigate whether the regulation had any significant impacts on long-term rents I use a database of 121 cities in the L.A. county for which I have rental market data from Zillow. From this source I obtain each city's monthly time series of ZRI (Zillow Rent Index) from November 2010 to April 2018. The ZRI is a smoothed, seasonally adjusted measure that reflects the median market rent for the entire stock of homes in a given region at a given point in time and is a good measure of the rents faced by residents (Barron, Kung, and Proserpio 2018).

5.1 Fixed Effects Estimation

My first attempt to assess the impact of the regulation on rents is based upon a simple fixed-effects framework. Let Y_{it} be the rent index in city i at time t and assume the following relationship

$$Ln(Y_{it}) = \alpha + \delta_i + \gamma_t + \beta(Reg_{it}) + \epsilon_{it} \tag{3}$$

where δ_i captures city-specific time-invariant characteristics, γ_t captures time-specific shocks that affect all cities in the same way, Reg_{it} is a dummy variable indicating whether city i is under

[17]The report is available at https://www.smgov.net.

short-term rental regulation at period t, and ϵ_{it} are city-specific time-varying unobservable shocks that affect local rents. Effectively, Reg_{it} will only have value one when the observation is of Santa Monica in some month-year after May 2015. Table 7 shows the results for the coefficient of interest β. Given that Zillow provides city level rent index disaggregated by the number of bedrooms a housing unit has, I include the results for seven different regressions, one for the rent index of all homes, and a separate one for each category provided by Zillow, from Studio apartments to homes of 5 bedrooms or more.

Table 7: Effect of the regulation on long-term rents (in logs)

	All Homes	Studio	1 Bed	2 Bed	3 Bed	4 Bed	5+ Bed
Regulation	0.1128***	0.0541***	-0.0136*	0.0868***	0.110***	0.0814***	0.0345***
	(0.0069)	(0.0105)	(0.00815)	(0.00659)	(0.00674)	(0.00725)	(0.00870)
City FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs	10,350	2,340	5,670	9,900	10,350	9,990	8,280
R-squared	0.996	0.980	0.980	0.989	0.993	0.995	0.995

Standard errors in parentheses
***p<0.01, ** p<0.05, * p<0.1

Most results of Table 7 may initially look puzzling, given that they would suggest that the regulation has actually increased long-term rents in Santa Monica for most types of houses, except for units of one bedroom. However, one cannot interpret the results of this fixed-effects regression as causal. It is very likely that ϵ_{it} is correlated to Reg_{it}, which implies that the estimated β coefficient is probably biased. In other words, most likely the pre-regulation trend in long-term rent prices was very different in Santa Monica from the rest of cities, causing the estimated effect of the regulation to be biased.

For example, suppose the Home-Sharing Ordinance was enacted in a period when long-term rents in Santa Monica were increasing significantly faster than elsewhere. In fact, Figure 11 suggests that this was likely the case. Under this scenario, the fixed-effects regression would capture this faster increase around the regulation period as an effect of the ordinance, delivering positive estimates of β. This is a common problem with the use of fixed-effect models to estimate public interventions that usually respond to trends in the main outcome of interest, which produces a spurious correlation between the policy and the outcome. In Santa Monica's context, this means that it is not unreasonable to think that the ordinance was passed precisely as a response to observed strong increases in rents, which explains why the regression produces positive estimates for the regulation's impact on long-term rents for most types of housing units.

In the same spirit, given the reasons mentioned above for being very suspicious about interpreting the fixed-effects model β coefficient as a causal effect of the Home-Sharing Ordinance, one should not use the results of the regression that uses one bedroom units as the dependent variable as convincing evidence of a negative impact on rents. The estimated impact is indeed negative, which would be in line with the basic argument that the reduction in vacation rentals increases the supply of housing units to long-term tenants, reducing the housing costs for residents. However, for its very nature the regulation is endogenous and the β coefficient in Equation 3 cannot be interpreted as the causal impact of the regulation.

Figure 11: Evolution of long-term rent index

5.2 SCM to estimate the impact on long-term rents

Given the issues related to the estimation of the impacts of the Home-Sharing Ordinance on residential rents using a fixed-effects model, I evaluate the regulation's effect on long-term rents using a Synthetic Control Method approach. For briefness, I do not detail here each step of the analysis as it was done with the application of the method to estimate the impact on the number of Airbnb listings. The step by step details of the application of the SCM to estimate the effects on house rents (as well as on house prices) is available upon request and is also part of a separate paper exclusively focused on the task of estimating the effects of the ordinance on housing costs.

Here, I focus on the estimation of the ordinance's impact on long-term rents of one bedroom homes. I choose one bedroom units because most Airbnb listings in Santa Monica are one bedroom apartments and due to the results of Table 7, which suggest that the ordinance's impact on rents was probably stronger for this type of housing unit. The analysis starts from the 121 cities for which I have long-term rent data. However, not all of these cities have the rent index measure specific to one bedroom homes. After dropping the cities for which this measure is not available, I am left with 63 cities in the L.A. county.

Next, to avoid including in the pool of donors cities with very different pre-regulation trends in long-term rents compared to Santa Monica, I rank the 63 cities in terms of their rate of increase in long-term rents of one bedroom units in the 12 months immediately before the adoption of the Home-Sharing Ordinance (June 2014 to May 2015) and I keep in the final pool of donors only the 24 control cities with pre-ordinance rates of increase in rents most similar to Santa Monica's rate (that amounts to keeping in the pool of donors the cities which have experienced high rent increases). Thus, I have as the final pool of control cities, the 24 cities in the L.A. county that best resemble Santa Monica's pre-regulation rate of increase in long-term rents of one bedroom homes.

I construct a synthetic counterpart of Santa Monica that resembles its pre-regulation trend in long-term rents of one bedroom units. As predictors of rents I use different linear combinations of the main outcome variable itself: the yearly average rent in 2011, 2012, 2013, 2014, as well as the mean over the first 4 months in 2015 and the observed rent in May 2015, the last month before the regulation became effective. Using the observed value of the outcome variable in the last period before "treatment" is common in the literature and it is a way to match the synthetic and actual trend right before the intervention under study takes places and see how they depart from each other following the intervention. I also use as predictors other variables that usually correlate with rent levels, namely population, median income, employment rate, and share of residents with college degree or higher.

Looking at the trend for Santa Monica itself (Figure 11 top-right graph), one is tempted to think that the slow down in rent levels after the ordinance was caused by this regulation. However, in Figure 12, which plots this trend contrasted with the evolution in rents for the synthetic version of the city, there is no clear indication that the regulation had a significant impact on reducing residential rents for one bedroom units. Differently than what we observe in Figure 8, the synthetic trend does not seem to systematically depart from the actual trend observed in Santa Monica after the passage on the ordinance restricting vacation rentals. To provide a better graphical illustration of how the method generated synthetic controls that were able to approximate well the actual trend in each of the cities in the pool, Figure 13 plots the comparison of actual and synthetic trends not only for Santa Monica, but also for other three cities, Alhambra, Burbank, and Los Angeles.

Figure 12: Trends in rents for one bedroom units: actual versus synthetic Santa Monica

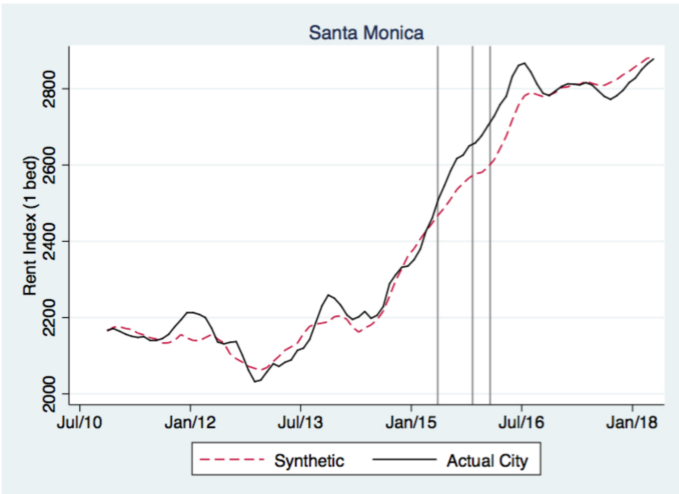

Figure 13: Trends in rents for one bedroom units: actual versus synthetic city

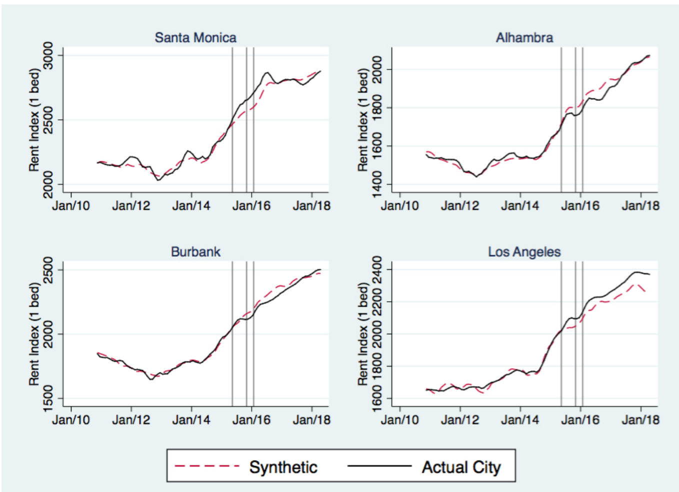

The comparison between the top-left plot (Santa Monica) and the other three plots provides anecdotal evidence that the Home-Sharing Ordinance did not have a specially strong effect in reducing rents for residents looking for one bedroom units. In the same spirit of the analysis on Airbnb listings, to be able to make statistically meaningful claims about the regulation's effect on rents in Santa Monica, we need to carry out inference exercises based on placebo runs of SCM. Thus, I run placebo analysis assuming each one of the cities in the pool to be the treated one and I obtain Figure 14, which plots the gaps between observed rents and the ones estimated via the SCM for each city. It is true that, as usual, some cities are not well approximated by a convex combination of the others (the gap is large even before the adoption of the regulation), but there is no need to successively exclude them from the graph and conduct inference analysis as it was done with the analysis on Airbnb listings, because here it is very clear that post-regulation Santa Monica does not stand out as the city in which the gap (effect) is the largest. In fact, there are many cities for which the gaps look larger than for Santa Monica. Therefore, the SCM does not find any significant effect of the regulation on long-term rents of one bedroom home units in Santa Monica. I also conduct analogous SCM analysis for the rent index of all homes as well as for different number of bedrooms (studios, two bedrooms, three bedrooms, four bedrooms, and five or more bedrooms) and the results point in the same direction, there is no evidence of a causal effect of the ordinance in reducing long-term rents.

Figure 14: Gaps between observed and synthetic trends in rents (one bedroom unit)

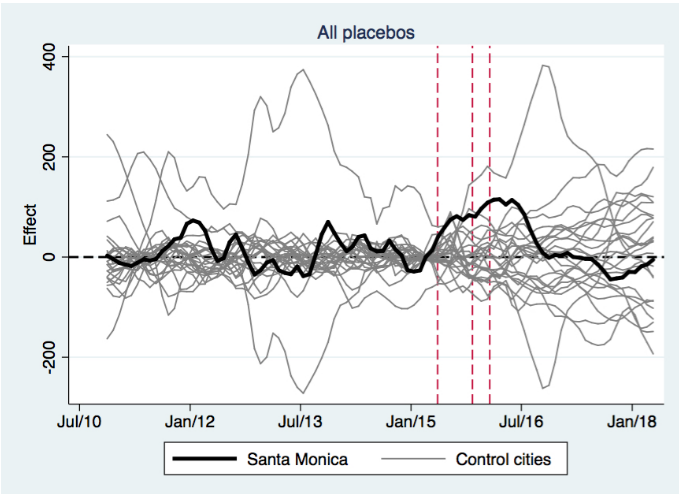

6 The Effect of the Ordinance on Housing Reallocation

As previously showed, I do not find significant effects of the Home-Sharing Ordinance on local rents faced by Santa Monica residents. It seems very intuitive that by reducing the number of homes that are operating as vacation rental apartments on Airbnb, the ordinance would have caused residential rents to decrease. In this section I provide some suggestive evidence of why the regulation was not effective in reducing local rents. I stress that the evidence I bring here is only intended to be suggestive and to provide some initial insights into the question of why the regulation did not affect residential rents.

The argument that a law restricting the number of vacation rentals will lead to a reduction in local rents rests upon the assumption that housing units prevented from operating as short-term rentals by the regulation will be offered for long-term tenants looking for houses to live in a permanent basis. Here I provide suggestive evidence that the reason behind the ineffectiveness of the ordinance to reduce residential rents may be related to the fact that, in reality, landlords of would-be vacation apartments are not willing to immediately supply their units for long-term tenants. That is, from the perspective of a landlord, these two alternative allocations for housing units (long-term and short-term renting) do not seem to be as close substitutes as one may initially think. To explore this idea, I use a panel dataset built using information available in the American Community Survey on 139 cities in the state of California from 2007 to 2016.

First, hypothesize that if landlords prevented from using their apartments as vacation rentals indeed decide to allocate their housing units to long-term tenants, we should see an increase in the share of occupied houses that are inhabited by renters versus owners. To illustrate this idea, suppose a long-term rental contract finishes in a city where there is no regulation restricting vacation rentals. In this scenario, if the owner looks around and observes that many people are renting to short-term visitors and obtaining a higher profitability with this activity than with the usual long-term tenants, it is very likely that she will not look for a next permanent renter and choose to allocate this apartment to short-term guests. On the other hand, if a similar situation happens in Santa Monica, after taking into account the high fines issued for illegal short-term rentals, the owner may decide to list her unit for long-term rent and wait for another permanent renter (indeed this was exactly the hope of the City Council when enacting the Home-Sharing Ordinance). And given that short-term rental units (vacation rentals) are considered vacant houses, whereas units rented for long-term tenants are considered renter-occupied houses, under this hypothetical scenario, Santa Monica post-regulation share of houses occupied by renters would increase relative to non-regulated cities. Table 8 test this hypothesis by regressing the share of occupied houses on a city fixed effect, a time fixed effect, and an indicator variable for Santa Monica post-regulation. The specification also includes basic city specific time-varying controls, namely income, population, employment rate, and the share of population with college degree or higher.

The results do not support the idea that, in Santa Monica, would-be vacation rental units are being supplied in the long-term market. Since the dependent variable is expressed in percentage terms, the regulation coefficient is actually indicating that the share of occupied houses inhabited by renters in Santa Monica reduced by 0.21 percentage points after the regulation's adoption. In any case, this result was highly insignificant.

Table 8: Effect of the regulation on renter occupancy rate

	Renter Occupancy Rate
Regulation	-0.210
	(2.761)
City FE	Yes
Time FE	Yes
Controls	Yes
Obs.	1,543

Standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

I also explore the same dataset to study the effects of the regulation on vacancy rates. Under the assumption that homes prevented from operating as vacation rental units will be supplied to long-term tenants as residential houses, we should expect vacancy rates in Santa Monica to evolve differently than in other cities where there was no regulation. More specifically, in the ACS survey, vacant housing units are classified as either for-rent, for-sale, for vacation and seasonal purposes, or for other/unknown reasons. If we believe that housing units deterred from operating on Airbnb as short-term vacation rentals are going to be supplied for residents in the long-term rental market, we should observe that in Santa Monica, relative to other cities, the rate of vacant-for-rent units increased after the Home-Sharing Ordinance. The idea is that, if the regulation is achieving its goal, at any given point in time, there is a flow of houses that in non-regulated cities are being converted to vacation apartments, whereas in Santa Monica they are actually being offered as residential rentals instead. Table 9 shows the results of regressions similar to ones performed to study renter-occupancy rates. Once more, the relevant dependent variable is regressed on a city fixed effect, a time fixed effect, and an indicator variable that assumes value one for Santa Monica after May 2015. All regression include a set of four basic controls, median income, population, employment rate, and the share of the population with college education of higher. Each dependent variable is constructed as the number vacant houses divided by the total number of housing units and multiplied by 100, representing the percentage of total housing units that is vacant in a city at a given year. Each column restricts the vacant houses by a specific reason why that unit is vacant. While in the first column the dependent variable includes all types of vacant homes, each one of the other columns restricts vacant units respectively to vacant for seasonal and recreational use, vacant-for-rent homes, vacant-for-sale houses, and units which are vacant for unknown reasons.

Table 9: Effect of the regulation on vacancy rates

	All Vacant	Seasonal Units	Vacant for Rent	Vacant for Sale	Vacant (unknown)
Regulation	2.790***	0.022	-0.073	0.566	2.359***
	(0.904)	(0.329)	(0.523)	(0.405)	(0.458)
City FE	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Obs.	1,239	1,239	1,239	1,239	1,239

Standard errors in parentheses
***p<0.01, ** p<0.05, * p<0.1

If one looks the results of the first column, one may be led to think that the Home-Sharing Ordinance had a significant impact in increasing the number of houses being offered for long-term tenants, since the vacancy rate has increased relative to other cities. However, when we look at the analysis for vacant houses disaggregated by the reason why the unit is vacant, the results do not support the hypothesis that the regulation has increased the number of houses offered for residential rent. In fact, the estimated effect of the ordinance on the rate of houses vacant for rent was slightly negative and highly non-significant. Similarly, the effect of the regulation on vacant houses for seasonal purposes or for sale was also non-significant. In reality, virtually all the effect that was captured in the regression displyed in the first column comes from houses vacant for unknown reasons. Although I cannot make any robust conclusion from these results, there is no indication that the supply of apartments for long-term rent available for residents has increased as a result of the regulation.

In the context of the same example already discussed, when a long-term rental contract ends and the owner has the option between looking for a next long-term tenant or allocating his unit to short-term guests, it does not seem that the regulation was capable of inducing Santa Monica owners to choose the first option more often when compared to owners in other cities. One possible interpretation of the positive and significant effect of the regulation on the rate of housing units vacant for unknown reasons may be that these owners are choosing not to allocate their housing units to the long-term rental market and instead are waiting to find out whether the regulation will remain in place for a long time, therefore the rate of vacant houses for unknown reasons in Santa Monica is increasing more than in other locations in California. Another possibility is that not all owners who use their units as vacation rentals on Airbnb are on the margin between allocating their empty house to the short versus long-term rental market. Suppose an owner wants the flexibility to have their Santa Monica apartment available for them to spend some of the summer months or even a holiday weekends. In that case, they may be wiling to offer this house as short-term rental on Airbnb, but not as a residential unit to long-term tenants, since that would prevent them from also being able to use the place occasionally. One potential drawback of the analysis presented above is that the ACS data only goes until 2016, which may be too short of a post-ordinance period for housing reallocation to actually go through significant changes.

7 Conclusion

This paper has presented a detailed study estimating the effect of the Home-Sharing Ordinance on the number of Airbnb entire home listings active in the city of Santa Monica. Through the application of a Synthetic Control Method approach, I show that, in roughly two years after its adoption, the ordinance caused the number of Airbnb entire home listings to decrease by approximately 861 units, a 61% reduction comparing to the number of listings that would have existed in the city in the absence of the regulation.

The fact that the regulation has reduced the total number of listings advertised on Airbnb is hard to argue against. All cities in the region have seen a persistent growth in Airbnb listings over the period of analysis, except Santa Monica, which displayed sharp drops followed by a more recent come back in the number of listings. It is still a question for future research to analyze the reasons why the regulation, although reducing listings, was not able to completely eliminate illegal vacation rentals listed Airbnb. Additionally, it would be interesting to conduct future studies to understand the determinants of who complied and who continued to illegally offer vacation rental units in the city.

Furthermore, I investigate the impact of the ordinance on residential rents and I don't find any evidence that the regulation has achieved one of its ultimate goals, local long-term rents were not significantly reduced in Santa Monica. I show suggestive evidence arguing that the reason why the regulation did not achieve its central goal may be related to the fact that many housing units prevented from operating as vacation rentals are not actually being supplied in the long-term rental market. This may be because owners are waiting before deciding to completely give up on allocating their units to short-term guests, because some owners were not actually on the margin between short-term and long-term renting, or because the short-term rental market was not big enough relative to the long-term one to have a substantial impact on residential rents.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2010). "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program". In: *Journal of the American statistical Association* 105.490, pp. 493–505.

Abadie, Alberto and Javier Gardeazabal (2003). "The economic costs of conflict: A case study of the Basque Country". In: *American economic review* 93.1, pp. 113–132.

Barron, Kyle, Edward Kung, and Davide Proserpio (2018). "The Sharing Economy and Housing Affordability: Evidence from Airbnb"". In: *Unpublished working paper.*

Cox, Murray and Tom Slee (2016). "How Airbnb's data hid the facts in New York City". In: *Insideairbnb. com.*

Coyle, Diane and Timothy Yeung (2016). "Understanding Airbnb in fourteen European cities". In: *The Jean-Jacques Laffont Digital Chair Working Papers.*

Edelman, Benjamin G and Damien Geradin (2015). "Efficiencies and regulatory shortcuts: How should we regulate companies like Airbnb and Uber". In: *Stan. Tech. L. Rev.* 19, p. 293.

Ferman, Bruno, Cristine Pinto, and Vitor Possebom (2018). "Cherry picking with synthetic controls". In: *Working Paper, FGV - São Paulo School of Economics.*

Galiani, Sebastian, Brian Quistorff, et al. (2016). "The synth_runner package: Utilities to automate synthetic control estimation using synth". In: *Unpublished paper, University of Maryland.*

Horn, Keren and Mark Merante (2017). "Is home sharing driving up rents? Evidence from Airbnb in Boston". In: *Journal of Housing Economics* 38, pp. 14–24.

Kaplan, Roberta A and Michael L Nadler (2015). "Airbnb: A case study in occupancy regulation and taxation". In: *U. Chi. L. Rev. Dialogue* 82, p. 103.

Kaul, Ashok et al. (2016). "Synthetic control methods: Never use all pre-intervention outcomes as economic predictors". In: *University of Hohenheim, Department of Economics.*

Lee, Dayne (2016). "How Airbnb short-term rentals exacerbate Los Angeles's affordable housing crisis: Analysis and policy recommendations". In: *Harv. L. & Pol'y Rev.* 10, p. 229.

Li, Jun, Antonio Moreno, and Dennis J Zhang (2015). "Agent behavior in the sharing economy: Evidence from Airbnb". In: *Ross School of Business Working Paper Series* 1298.

Quattrone, Giovanni et al. (2016). "Who benefits from the sharing economy of Airbnb?" In: *Proceedings of the 25th international conference on world wide web.* International World Wide Web Conferences Steering Committee, pp. 1385–1394.

Samaan, Roy (2015). "Airbnb, rising rent, and the housing crisis in Los Angeles". In: *Los Angeles Alliance for the New Economy (LAANE) Report.*

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

Zervas, Georgios, Davide Proserpio, and John W Byers (2017). "The rise of the sharing economy: Estimating the impact of Airbnb on the hotel industry". In: *Journal of Marketing Research* 54.5, pp. 687–705.