The Effect of Short-Term Rentals on Residential Investment

Ron Bekkerman* Maxime Cohen† Edward Kung‡
John Maiden§ Davide Proserpio¶

August 23, 2021‖

Abstract

We provide new evidence that short-term rental (STR) platforms like Airbnb incentivize residential real estate investment. We exploit two complementary identification strategies. First, we use variation in the timing of STR regulations to estimate the effect of regulation on both Airbnb listings and residential permits. We find that over the first 12 months following the start of the regulation, STR regulations reduce Airbnb listings by 11.3% and residential permits by 11.2%. We estimate an elasticity of permits with respect to Airbnb listings of 0.769. Second, we show that residential permits decline discontinuously across jurisdictional boundaries in which one side of the boundary has a STR regulation and the other side does not. The effect is especially pronounced for accessory dwelling units, which decline by 16.5% across regulatory boundaries. Our results imply that STRs incentivize residential investment, and especially so for housing units that are well suited for short-term renting.

JEL Codes: R31, R38, L86

*Cherre, Inc.; ron@cherre.com

†McGill University; maxime.cohen@mcgill.ca

‡California State University, Northridge; edward.kung@csun.edu

§johnmaiden@fb.com

¶University of Southern California; proserpi@usc.edu

‖Authors listed in alphabetical order. We would like to sincerely thank David Wachsmuth from McGill University for sharing AirDNA data with us, and Kaihang Zhao and Elizabeth Slavin for research assistance.

1 Introduction

Home-sharing and short-term rentals (STRs), facilitated by web-based platforms such as Airbnb, have attracted extensive public debate and academic scrutiny in terms of their effect on the housing market. Critics argue that home-sharing platforms cause some landlords to switch from supplying the market for long-term rentals—which serves local residents—to supplying the STR market which serves non-residents. Since the total housing supply is inelastic in the short-run, this drives long-term rental rates up. A growing body of research supports this thesis.[1] Motivated by this research and by the growing concern for housing affordability in many metropolitan areas, policymakers around the world have implemented more stringent regulations of home-sharing platforms.

Yet, there is a reason to be more sanguine about the long-term effects of home-sharing on housing costs. As noted in Barron et al. (2021), the main effect of home-sharing is to increase the option value of spare capacity in residential housing. For example, Airbnb lets home owners rent extra rooms (capacity that would otherwise go unused) to consumers in the STR market.[2] The ability to better utilize spare capacity increases the economic value of and the demand for residential housing. Over the short-run when housing supplies are inelastic, this raises house prices and rental rates. Over the long-run, however, home-sharing may lead to a greater investment in residential housing. In particular, we might see an increase in supply of housing units that can be flexibly allocated to both long- and short-term rental markets.

In this paper, we investigate whether Airbnb has had a causal effect on the

[1] See Horn and Merante (2017); Garcia-López et al. (2020); Barron et al. (2021); Valentin (2021); Koster et al. (2021). Using a variety of research designs and settings, all these studies show a positive causal effect of Airbnb market penetration on house prices and rental rates.

[2] One could ask what was innovative about Airbnb that allowed this to happen, and why was it more difficult before. Einav et al. (2016) discuss the innovations that gave rise to peer-to-peer markets, centering on reductions in transactional and information frictions associated with trust.

quantity and type of residential real estate investment. We leverage a comprehensive dataset of public records—residential permits, tax records, and residential sales transactions—issued in 15 U.S. metropolitan areas from 2008 to the end of 2019. We supplement this data with Airbnb listings data over the same period. Altogether, our data consists of 2.9 million residential permits, $750,000$ Airbnb listings, and 4 million residential sales transactions from 15 metropolitan areas covering a period of 12 years.

We test for the causal effect of Airbnb on residential investment using two approaches. Our first approach exploits differences in the timing of STR regulations across cities. Regulations have proven effective in limiting the growth and usage of STRs.[3] We thus test whether the growth in Airbnb listings and residential permits both declined after the implementation of STR regulations. Our method follows the staggered difference-in-differences framework introduced in Callaway and Sant'Anna (2020), in which average treatment-on-treated (ATT) effects are estimated separately by group and by time period. We estimate that the average of these group-time ATTs across all groups and treated time periods is -21.4% on Airbnb listings and -10.4% on residential permits. The results support our hypothesis that the ability to home-share has a positive causal effect on residential investment.

To estimate the elasticity of growth in residential permits with respect to growth in Airbnb listings, we regress residential permits on Airbnb listings, using the presence of regulation as an instrumental variable. We estimate that a 1% increase in Airbnb listings is associated with a 0.769% increase in residential permits.

To translate our findings on residential investment into an approximate dollar value, we analyze more than 480k properties from our 15 metropolitan areas that were bought and sold more than once (i.e., repeat sales) within our sample period.

[3]See Basuroy et al. (2020); Koster et al. (2021); Valentin (2021).

We find that the presence of an approved permit application between sales accounted for a 38% increase in the latest sales price relative to properties without an approved permit between sales. For cities with STR regulations, we estimate that the decrease in permit applications due to the regulations have collectively reduced property values by $2.8 billion and tax revenue by $40 million per year for the 15 cities in our sample.

Our second approach to estimating the effect of STR on residential investment exploits differences in STR regulations across jurisdictional boundaries. For this analysis, we focus on Los Angeles County, which provides a unique setting for our study because it represents one interconnected housing and labor market but contains multiple jurisdictions that set their own land-use policies. We then test whether the number of residential permits varies discontinuously across jurisdictional boundaries when home-sharing is regulated on one side of the boundary but not on the other side. This approach is analogous to Koster et al. (2021), who used a similar method to test the effect of STR regulations on house prices and rental rates. In our paper, we leverage this method to test the effect of STR regulations on residential investment. We find that the total number of residential permits is on average 18% lower on the regulated side of the boundary relative to the unregulated side. The effect is more striking on permits for Accessory Dwelling Units (ADUs) than for other permit types. ADUs are housing units located on the same lot as a main property but are intended for use as an independent living space, such as a converted garage or an additional unit in the backyard. ADUs are especially well suited for renting to either the long or short-term rental markets, so they should be more affected by STR regulations. We find that the difference in ADU permits between regulated and unregulated boundaries is 17%. In comparison, the difference in non-ADU permits between regulated and unregulated boundaries is only 9%. Besides being informative about the heterogeneous impact of Airbnb, the fact that we see a stronger effect for permits of units well suited for STR provide further support to our ability to identify the causal effect of Airbnb on residential investments. Overall, our border discontinuity results suggest that STR regulations reduced residential investment, and especially so for housing units that are well suited for short-term rentals.

We conclude the paper by testing our prediction that home-sharing increases the value of and demand for residential housing. To test whether STR regulations have a causal impact on residential demand, we again exploit differences in the timing of STR regulations across cities, while using home sales volume and house prices as outcomes. We do not find a statistically significant effect of STR regulations on sales volumes, but we estimate an average group-time ATT of -3.3% for house prices. Consistent with our prediction, these results imply that STR regulations reduced the economic value and the demand for residential real estate. In addition, these results are consistent with the literature that shows a positive effect of STRs on house prices and rental rates.[4]

Taken all together, our results show that STR platforms like Airbnb incentivize residential real estate investment. Increased residential investment can benefit cities in a number of ways, such as providing additional tax revenues, adding to the quality and quantity of existing housing stock, and improving neighborhood attractiveness and character. We observe an increase in the number of independent housing units, as shown by our analysis of the growth in ADUs in Los Angeles. This means that not all STRs come from the reallocation of the existing housing stock: some will also come from investment in increasing housing capacity. The increase in economic value of housing and the expansion of supply should ultimately translate into higher tax revenues for the city. These revenues can then be spent on social programs to develop

[4]See Horn and Merante (2017); Garcia-López et al. (2020); Barron et al. (2021); Valentin (2021); Koster et al. (2021).

affordable housing that can offset the increased housing cost driven by STRs. Our results suggest that the time may be ripe to revisit stringent STR regulations that ultimately can be more detrimental than beneficial for the cities that enact them. If the migration trend from cities to suburban areas spurred by work-from-home and COVID-19 continues, STRs may play a crucial role in revitalizing city centers when travel demand rebounds.

Literature Review

There is a large growing body of research studying home-sharing and its effect on various socioeconomic indicators. Papers closely related to our work study the impact of Airbnb on the housing market and, more generally, on urban development. Using different approaches and focusing on different municipalities, Barron et al. (2021); Koster et al. (2021); Horn and Merante (2017); Garcia-López et al. (2020); Valentin (2021) all show that Airbnb has a positive causal effect on the prices of both the rental and housing markets. Calder-Wang (2019) focuses on the welfare effects associated with the changes in the housing market due to Airbnb in New York City. The author shows that renters have suffered a loss of $178mm; moreover, these losses fall disproportionately on high-income, educated, and white renters.

Turning to the effect of home-sharing on cities, Almagro and Domınguez-Iino (2019) study how changes in tourism patterns driven by Airbnb affect the provision of amenities, and how these changes impact city residents with different preferences. Basuroy et al. (2020) measure the impact of Airbnb on the local economy of the neighborhoods it penetrates and show that restaurant revenues grow faster in areas with a higher Airbnb demand. Additionally, Jain et al. (2021) show that Airbnb data can be used as a leading predictor of gentrification, suggesting that Airbnb contributes to the gentrification of the neighborhoods it enters.

We contribute to this stream of research by studying the effect of home-sharing on residential real estate investments. We show that Airbnb induces an increase in residential investments measured by the number of residential permits in a given zipcode. We do so by implementing two identification strategies. The first approach is similar to Bibler et al. (2018) and exploits the variation in Airbnb policy timing across cities in our dataset, while the second approach is similar to Koster et al. (2021) and exploits the variation in home-sharing ordinances across the urban part of the Los Angeles County.

The rest of the paper is organized as follows. Section 2 discusses our data and institutional setting. Section 3 presents our results exploiting the variation in the timing of STR regulations. Section 4 presents our results using the border discontinuity strategy. Finally, Section 6 concludes.

2 Data and Setting

The goal of this paper is to study the effect of short-term rentals on residential investments. For this purpose, we collected data on 15 U.S. metropolitan areas from different sources including Airbnb, Cherre (a real estate data analytics company), the California Department of Housing and Community Development (CDH), Zillow, and the American Community Survey (ACS).[5] We next discuss each of our data sources.

[5]Considering the largest U.S. cities by population while capturing geographic and demographic diversity, we focus on the following 15 metropolitan areas that implemented short-term rental regulations in the last ten years: Austin (TX), Boston (MA), Chicago (IL), Columbus (OH), Denver (CO), Las Vegas (NV), Los Angeles (CA), Nashville (TN), New Orleans (LA), New York City (NY), Portland (OR), San Diego (CA), San Francisco (CA), San Jose (CA), and Seattle (WA).

2.1 Airbnb Data

Launched in 2008, Airbnb is a peer-to-peer platform for short-term rental accommodations. Hosts can list their properties for rent on the platform, and guests can book these properties for a short stay (i.e., a few days or weeks). Starting in 2008, Airbnb has experienced exponential growth, going from a few hundred listings in San Francisco to six million active listings in 220 countries by the end of 2020. Recent statistics suggest that more than 800 million guests have stayed with Airbnb hosts as of September 30, 2020.[6]

We obtained information on all available listings (including their approximate address and the longitude-latitude coordinates) and their reviews for the above set of metropolitan areas from three different sources. First, one of the authors of this paper collected consumer-facing information about the complete set of Airbnb properties located in the U.S. and about the hosts who offer them. The data collection process spanned a period of approximately six years, from mid-2012 to mid-2018. Second, we downloaded data from `insideairbnb.com`, a website that regularly scrapes Airbnb data. This data contains all the listings and reviews from the available listings on Airbnb during the period 2015–2020. Third, we obtained Airbnb listings and reviews from `airdna.co` covering the period 2015–2020.[7] Our raw dataset for the 15 target metropolitan areas includes around 30 million reviews and listing information for approximately 1 million properties listed on the Airbnb website during the period

[6]See `https://news.airbnb.com/about-us/`.

[7]The additional data from `insideairbnb.com` and `airdna.co` allow us to construct a panel of Airbnb listings that ends at the end of 2020 (even though, as explained later, we exclude the year 2020 to avoid potential COVID-19 effects).

2008–2020.[8]

Measuring Airbnb supply. Once we have collected the Airbnb data, the next step is to define a measure of Airbnb supply. This task entails selecting the geographic granularity of our measure and defining the entry and exit dates of each listing in the Airbnb platform. Regarding the geographic aggregation, we conduct our main analysis at the zipcode level. Zipcodes are the lowest level of geography for which we can reliably assign listings without error. Moreover, neighborhoods (which we approximate by zipcodes) are a natural unit of analysis for housing markets because there is significant heterogeneity in housing markets across neighborhoods within cities, but comparatively less heterogeneity within neighborhoods.

Regarding the entry and exit dates of each listing, this choice comes less naturally. Measuring *active* supply is challenging due to the presence of "stale vacancies," that is, properties that are listed on the Airbnb platform but for which the host has no intention of renting.[9] Thus, to construct the number of active Airbnb listings for each zipcode, we employ a method similar to the one used by Zervas et al. (2017) and by Barron et al. (2021). Specifically, we use the listing reviews to create a Time-to-Live (TTL) for each property based on the reviews it receives. A listing enters the market when the host registers with Airbnb and remains active for m months (the TTL). Each time a listing receives a review, the TTL is extended by m months from the review date. If a listing exceeds the TTL without receiving any reviews, it is then considered inactive. A listing can then become active again if it receives a new

[8]Since one of the authors started scraping data in mid-2012, it is possible that some listings created between 2008 and mid-2012 were deleted from the platform and therefore not included in our dataset. To make sure that these potentially missing listings are not driving our results, we tested the robustness of our finding by considering several time windows that exclude Airbnb's early years and we obtained consistent results.

[9]Fradkin (2017) estimates that about 15% of guest requests are rejected by U.S. hosts because of stale vacancies.

review. In all the analysis reported in the paper, we use a TTL of six months.[10]

2.2 Permits, Sales, and Prices Data

Residential permits. Our permits data is sourced directly from all of the 15 cities, leveraging their open data initiatives to obtain 2.9 million residential permits that span from 2008 until the end of 2020 (see Table 1 for summary statistics). In the U.S., public records such as permits data are collected at the county level, which means that each city dataset contains permit entries for the focal city as well as for the surrounding municipalities.

Permits data (as with other public data such as tax records and deed transactions) are reported at the county level and regulated at the state level, making it hard to standardize the records across multiple jurisdictions. To handle the variation in the type and quality of data provided by each municipality, we analyzed only approved permits, utilizing the following elements for analysis: the city-provided permit identifier, the permit issue date, and the location of the property (which could be either an address or the longitude-latitude coordinates).

To separate residential permits from commercial permits (some cities segregate these data, whereas others combine them), we leverage Cherre's data platform to standardize location data and map permits to residential tax lots across the U.S. The residential tax lot data itself is sourced from one of the major data providers that regularly consolidates publicly available records from all 3,000+ U.S. county-level *tax assessor* and *recorder of deeds* offices into a single harmonized proprietary dataset, indexed at the tax-lot level. This ultimately allows us to join permit applications, tax records, and residential sales transactions to their respective properties via a tax lot identifier. We provide summary statistics of our residential permits by metropolitan

[10]We have also checked that our results are robust to using a TTL of three months.

area in Table 1.

California Department of Housing and Community Development (HCD). From the California HCD, we obtained the full list of residential permits submitted between 2018 and 2019 for the 88 incorporated cities located in the Los Angles County. This data includes the exact address of the property and the type of property for which the permit was requested (e.g., whether the property was a single-family, multi-family, or an attached dwelling unit). There are 32,571 permits in this dataset.

Residential sales. Our transaction-level residential property sales data is sourced and standardized by Cherre's data provider, who collects publicly available deed transactions from all 3,000+ county-level recorder of deeds offices. Additional processing is then performed by Cherre to standardize use codes and addresses. From this nationwide dataset, we pulled the four million residential property transactions for our target cities covering the period 2008–2020, aggregating the transactions at the zipcode level to track the sales transaction volume on a monthly basis. We excluded transactions with a price per square foot less than $10,[11] or greater than $2,000, treating them as outliers.

House prices. We obtained a measure of house prices at the zipcode level from Zillow.com, an online real estate company that provides estimates of house and rental prices for over 110 million homes across the U.S.[12] In addition to reporting house value estimates, Zillow provides a set of indices that track and predict home values and rental prices on a monthly level and at different geographical granularity levels. In our analysis, we use the zipcode-level Zillow Home Value Index (ZHVI) that estimates

[11]Lower-end transactions tend to be transfers of deeds between family members and do not represent the market-based price for the property.

[12]`https://www.zillow.com/research/data/`.

the median transaction price for the actual stock of homes as a measure of house prices.

[Table 1 here.]

Summary statistics. Table 1 provides a summary of the Airbnb, permits, sales, and price data that we collected for each of our 15 metropolitan areas. We restrict our time period to January 2008–December 2019 to limit any effects of COVID-19 on our analyses and results. The resulting analysis sample contains 540,000 Airbnb listings, 1.8 million residential permits, and 2 million home sales over 608 zipcodes in 15 metropolitan areas from 2008 to 2019.

[Table 2 here]

2.3 Home-Sharing Policy Data

Home-sharing regulation and tax policies. Each of the cities in our analysis besides San Diego has implemented some type of regulation related to home sharing in the last decade. We manually collected information on these regulation and policies from the Airbnb and city websites. The list of cities and the policy implementation dates are reported in Table 2.

Los Angeles home-sharing ordinances. In the period 2014–2020, 18 out of the 88 cities in Los Angeles County implemented Home Sharing Ordinances (HSO). The list of the 18 cities in Los Angeles County and the start dates (year-month) of the HSO enforcement were manually collected and are reported in Table 3.

[Table 3 here]

2.4 Additional Data Sources

Zipcode-level time-varying characteristics. To control for zipcode-level economic conditions, we leveraged the American Community Survey (ACS) to obtain zipcode-level annual estimates of median household income, population size, percentage of 25–60 years old with a bachelors' degree, and employment rate. These estimates are available for zipcode tabulation areas (ZCTA) for the years 2011–2019. We map ZCTAs to zipcodes using a crosswalk provided by the U.S. Department of Housing and Urban Development. The data were downloaded from `data.census.gov`.

3 Analysis Using Regulatory Variation

Regulation has proven to be effective in limiting the growth and usage of short-term rentals (Basuroy et al., 2020; Koster et al., 2021; Valentin, 2021). The introduction of new short-term rental regulations can thus be leveraged as a variation that shifts Airbnb usage in a neighborhood, but at the same time is exogenous to other factors affecting residential investment. The time period between 2008 and 2019 is a fruitful period to study this variation because many cities have passed short-term rental regulations at various times.[13] Table 2 shows the cities we use in our analysis and the dates at which short-term rental policies went into effect.

Our first empirical strategy exploits the variation in the time at which short-term rental regulations went into effect across different U.S metropolitan areas. In doing so, we implement a difference-in-differences (DD) strategy in which we compare the change in residential real estate investment for areas where a STR regulation has been implemented to the change in residential real estate investment for areas where no

[13] Although our raw data includes the observations from 2020, we only use data up to the end of 2019 to avoid effects from the COVID-19 pandemic.

regulation has yet been implemented over the same period.

Using variation in the regulation timing offers two main advantages. First, since the regulation timing is different in each city, the variation in regulation is not confounded with national time trends in residential investment. Second, because we are comparing changes over time, our variation is also not confounded with permanent unobserved heterogeneity across neighborhoods. To identify the causal effect, we require the classic parallel trends assumption—namely, that in the absence of treatment, the outcome for treated and control units would have evolved in a similar fashion. We discuss this assumption in more detail in Section 3.1.

We measure residential investment as the count of approved residential permits, as described in Section 2. Table 1 summarizes the data we use in this section.

3.1 Treatment Effects Estimation

We estimate the treatment effects using the staggered DD framework proposed by Callaway and Sant'Anna (2020). This group-time treatment effect framework fits well our setting because we have multiple units of observation (zipcodes) that are treated (regulated) at different time periods. As discussed, we use zipcodes as our unit of analysis to account for the highly localized nature of housing markets, especially in regards to the demand for short-term rentals.

We let i denote a zipcode and t a time period (month). We let Y_{it} be the observed outcome of interest, which is either the log number of Airbnb listings or the log number of residential permits.[14] For each i, we let G_{ig} be a binary variable equal to one if i was *first* regulated at time g (so g is the date at which the regulation was implemented in i). Finally, we let $Y_{it}(0)$ denote i's potential outcome at time t if untreated at time t, and $Y_{it}(g)$ denote i's potential outcome at time t if i was *first*

[14] When taking the log of a count variable, we add a factor of one to avoid taking logs of zero.

treated at time g. This framework thus allows for heterogeneous treatment effects with respect to the treatment date.

We assume that there is no anticipation of treatment, so that $Y_{it}(0) = Y_{it}(g)$ for all $t < g$. We also assume that once a unit is treated, it remains treated, and that the intensity of treatment is the same for all units and time periods.[15] The relationship between observed and potential outcomes is as follows:

$$Y_{it} = Y_{it}(0) + \sum_{g=1}^{T} \left[Y_{it}(g) - Y_{it}(0)\right] G_{ig}. \tag{1}$$

We then define the group-time treatment effect of interest as

$$ATT(g,t) = \mathbb{E}\left[Y_{it}(g) - Y_{it}(0) | G_{ig} = 1\right], \tag{2}$$

where $ATT(g,t)$ measures the average treatment effect at time t for the group of observations that were first treated at time g. The group-time treatment effect framework is flexible enough to accommodate for the estimation of many different types of aggregated treatment effects. For example, if one is interested in estimating the average treatment effect by length of exposure to treatment, one could estimate a weighted average of $ATT(g,t)$ conditional on $t - g = e$, where e is the desired length of exposure.

Callaway and Sant'Anna (2020) show that the group-time treatment effects are identified from data on Y_{it} and G_{ig} as long as the parallel trends condition holds, in addition to standard independence and support conditions. Formally, the parallel trends assumption for our setting can be stated as follows:

[15]We thus estimate an average effect over all types of policies implemented in our sample.

Assumption 1 *For each g, h, and t such that $g \leq t \leq h$,*

$$\mathbb{E}\left[Y_{it}(0) - Y_{i,t-1}(0) | G_{ig} = 1\right] = \mathbb{E}\left[Y_{it}(0) - Y_{i,t-1}(0) | G_{ih} = 1\right]. \tag{3}$$

Assumption 1 requires that if treated groups had instead not been treated, then their outcome would evolve in the same way as groups that have not yet been treated. Assumption 1 cannot be directly tested because the left-hand side of Equation (3) is not observed (i.e, we do not observe the counterfactual outcome for treated units). However, we are able to test whether groups of observations have a different trend with respect to the amount of time left until the regulation in the pre-treatment period. We will return to this point when we discuss the estimation results.

Each group-time treatment effect, $ATT(g,t)$, is estimated by computing a weighted DD estimate where the reference time period is $g-1$. The treatment group includes the observations with $G_{ig} = 1$, whereas the control group contains the observations with $G_{ig} = 0$ *and* have not yet been treated by time t. We refer to Callaway and Sant'Anna (2020) for the technical details on estimation and inference. We implement the estimator in `R` using the package `DID`, written by Brantly Callaway and Pedro Sant'Anna.

Overall treatment effect. We report the results in two ways. First, we report the overall treatment effect, θ_o, which is a weighted average of $ATT(g,t)$ for $t \geq g$ over all groups and time periods, that is,

$$\theta_o = \frac{1}{\kappa_o} \sum_{g} \sum_{t>g} \omega_g ATT(g,t), \tag{4}$$

where κ_o is the number of combinations of g, t with $t \geq g$, and the weights ω_g are proportional to the number of observations in each group. We report the estimates in Table 4. We find that the introduction of STR regulations reduces Airbnb listings by 21.4% on average in the post-treatment period, and the number of residential permits per month by 10.4% in the post-treatment period. Therefore, our estimates suggest that short-term rental regulations substantially reduce both Airbnb listings and residential investment as measured by the number of issued permits.

[Table 4 here]

Treatment effects by length of exposure. Second, we compute the treatment effect by length of exposure. We define $\theta(e)$ as the weighted average of $ATT(g, t)$ for all t and g such that $t - g = e$, that is,

$$\theta(e) = \frac{1}{\kappa_e} \sum_g \omega_g ATT(g, g + e), \tag{5}$$

where κ_e is now the number of groups and the weights ω_g are once again proportional to the group size. The parameter $\theta(e)$ can be interpreted as an event study. More precisely, for $e < 0$, $\theta(e)$ captures the trend in outcomes for groups that are e periods away from the regulation relative to other groups that are not yet regulated. Conversely, for $e > 0$, $\theta(e)$ captures the trend in outcomes for groups that are e periods into the regulation relative to groups that are not yet regulated. Figure 1 reports our estimates (and 95% confidence intervals adjusted for multiple hypothesis testing) of $\theta(e)$ for both Airbnb listings and residential permits, for e running between -36 and 36 months (i.e., three years in either direction).

[Figure 1 here]

No differential pre-trends. Figure 1 bears several important implications. First, there are no differential trends in the pre-treatment period between groups that are

e periods away from being regulated and other groups that are not yet regulated. This provides strong evidence in support of the parallel trends assumption discussed in Assumption 1. Treated units do not trend differently from untreated units in the periods leading up to the treatment. Thus, we should not expect them to have trended differently in the absence of treatment. The fact that treated and untreated units only diverge once the regulation occurs provides strong evidence that regulation is causing the change.

Increasing effect by length of exposure. Second, Figure 1 shows that the treatment effect is increasing with the length of exposure. The estimates imply that over a 12-month period, STR regulations reduce Airbnb listings by 11.3% and monthly residential permits by 11.2%. Over a 36-month period, STR regulations reduce Airbnb listings by 29.2% and monthly residential permits by about 16.9%. It should be noted that the number of listings is a stock (the number of listings counted as active in a given month), whereas the number of permits is a flow (the number of new permits issued that month), which explains why the treatment effect seems to be increasing more over time for listings than for permits.

3.2 Elasticity of Residential Investment w.r.t. Airbnb

In addition to the effect of regulation, policymakers may be interested in the elasticity of residential investment with respect to short-term rentals. For example, knowing this elasticity could help policymakers decide to what extent they should allow short-term rentals so as to achieve a target residential investment. To compute this elasticity, we estimate several regression specifications of the following form:

$$\ln(\text{permits}_{it}) = \alpha \ln(\text{airbnb}_{it}) + \mathbf{X}_{it}\beta + \delta_i + \gamma_t + \epsilon_{it}, \tag{6}$$

where $\mathbf{X}_{it}$ is a vector of time-varying zipcode characteristics from the ACS that serve as controls, δ_i captures zipcode fixed effects, and γ_t represents time fixed effects.

To account for the potential endogeneity between permits and Airbnb listings, we use regulation as an instrumental variable (IV) for the number of Airbnb listings. Regulation is a valid IV if it satisfies the relevance and exclusion restrictions. Relevance is testable and well established both in the literature (Koster et al., 2021; Basuroy et al., 2020) and from the previous section of this paper.

The exclusion restriction says that regulation should have no direct effect on permits except via its effect on Airbnb. Moreover, regulation should not be correlated with other unobserved factors that affect the number of permits. The exclusion restriction is not directly testable, but we believe that it is likely satisfied. First, we did not observe differential pre-trends in permits, which suggests that the regulation timing is not correlated with unobserved factors that affect the number of permits. Second, STR regulations are highly targeted policies that focus only on specific uses of a home. Third, we are not aware of any other land use or housing policies over this period that coincide with the timing of STR regulations. We thus doubt that these very specific regulations would affect residential investment except via their impact on short-term rentals.

The specification of our first-stage regression is given by:

$$\ln(\text{airbnb}_{it}) = \zeta \text{regulation}_{it} + \mathbf{X}_{it}\beta + \delta_i + \gamma_t + \varepsilon_{it}, \tag{7}$$

where regulation_{it} is a binary indicator equal to one if zipcode i is regulated by time t and zero otherwise.

Table 5 reports the results of our regressions. Specifically, we report both the OLS and the IV regression results. We report the results using both a log-log and a levels-

levels specification.. Our preferred specification is Column 2, which is the IV estimate using the log-log specification. The coefficient on ln(listings) in column 2 tells us the elasticity of residential permits with respect to Airbnb listings, which we estimate is 0.769. This implies that a 1% increase in Airbnb listings is associated with a 0.769% increase in residential permits. To sharpen the interpretation of our estimates' magnitudes, we run the same regressions without applying log transformations. The results are reported in Columns 4 and 5 of Table 5. In this case, we estimate that each additional Airbnb listing is associated with 0.092 additional residential permits.

[Table 5 here]

3.3 Implications for Property Values

In this section, we perform an approximate calculation to estimate the dollar value of the reduced residential investment caused by STR regulations. To do so, we estimate repeat-sales regressions (i.e., multiple transactions for the same property that occur within our analysis timeframe), while controlling for a binary indicator for whether a permit was approved for the focal property in between sales. Let i denote a property in city j that was sold at price p_i^s at time s and at price p_i^t at time t, with $s > t$. Let permit_i be a binary indicator for whether a permit was approved for property i in between times t and s. We then estimate the following regression specification:

$$\ln(p_i^s) - \ln(p_i^t) = \alpha\text{permit}_i + \delta_{js} - \delta_{jt} + \varepsilon_i, \tag{8}$$

where δ_{jt} are city-time dummies that capture arbitrary city level time trends. The coefficient α captures the average change in property value that results from having some renovation work done to the property (measured by an approved permit application).

To estimate this model, we use property-level residential sales transactions from Cherre for our 15 cities. Out of all the residential sales transactions in this data, 36.4%, or 486,886 properties, were sold at least twice between 2008 and 2020. Thus, the data we use consists of repeat sales transactions for 486,886 properties with sales pairs (purchase and sale dates) occurring between January 2008 and December 2020. The average time between sales is 1,451 days or about 4 years. We observe that 15.1% of the repeat sales transaction pairs have an approved permit in between the purchase and sale dates.

Using this data, we estimate that $\hat{\alpha} = 0.387$ with a standard error of 0.002. The estimate implies that properties in our sample which had a permit approved between sales have increased their sale price by 38.7% relative to similar properties that did not have a permit.[16]

What do these results imply in terms of the economic value of the residential investment foregone due to STR regulations? The average number of residential permits issued per year in our sample is 154,253 (total number across all zipcodes used in our analysis). The average property value in our sample, as measured by the ZHVI, is $\$452,438$. If a permit increases the property value by 38% on average, then a 10.4% reduction in permits due to STR regulations represents a loss in property value of \$2.8 billion per year across our 15 cities. If property tax rates are around 1.5%, then this represents about \$40 million in potential lost revenue per year for the cities in our sample. Admittedly, this is a rough calculation. There are a number of effects we did not take into account, such as potential heterogeneous treatment effects of permits across different property types, different tax rates across cities, and potential

[16] Our estimate of the effect of a permit is similar in magnitude to the effects of renovation reported in Bogin and Doerner (2017). They used a repeat-sales methodology that is similar to ours, using data from the FHFA property transactions database. They measured renovations by changes in property characteristics, whereas we use the presence of an approved permit application.

sample selection bias in our estimates due to using only properties that were sold twice within our timeframe. We also did not include the effects of reduced property values which STR regulations are shown to induce—the effect we calculate here is entirely due to reduced investment. Nevertheless, this back-of-the-envelope estimate suggests that cities with strict STR regulations may miss out on a significant increase in property value along with the resulting tax revenues.

4 Border Discontinuity Analysis

In this section, we present additional evidence to further support the results from Section 3.1. Specifically, to further study the effect of STRs on residential investment, we exploit the rich micro-dataset on residential building permits made available to us by the California Department of Housing and Community Development and the information about HSO implementations across the Los Angeles County. As discussed in Section 2, this data consists of all residential building permits issued in the Los Angeles County between 2018 and 2019 and the date in which HSO have been implemented in 18 out of the 88 Los Angeles County cities (see Table 3 for the list of cities with HSO). Since this data allows us to know the exact address of each permit's project site, we can observe how the number of permits varies across city borders when the city on one side of the border has a restrictive home-sharing ordinance and the city on the other side does not. If short-term renting indeed incentivizes residential investment, then the number of permits should be higher on the side of the border without a HSO. Moreover, the difference in permits should be higher for units that are especially suited for short-term renting, such as accessory dwelling units (ADUs).[17] The border discontinuity method was used in Koster et al. (2021) to estimate the effect of HSOs on house prices and rents and in Basuroy et al. (2020) to estimate the effect of HSOs on restaurant reviews. In this paper, we use the same strategy to study the effect of HSOs on residential investment. The identification assumption behind this analysis is that, when focusing on a small geographical area around the border, the only factor affecting the probability of observing a new residential permit is whether the permit is on the left or the right side of the border, and not as a function of changes in external factors (e.g., economic conditions) that favor the construction of residential properties on one side of the border.

[Figure 2 here]

We focus on the Los Angeles County because it is a large metropolitan area that represents one interconnected housing and labor market, yet includes multiple cities each with its own housing and land use policy. Out of the 88 cities in the Los Angeles County for which we have permits data, 18 of them had implemented HSOs by the time of our analysis period. We plot the geographic distribution of permits in Figure 2. Altogether, we have data on 32,368 permits issued in 2018–2019, 35% of which are single-family, 13% are multi-unit, and 52% are ADUs.

4.1 Graphical Analysis

We first present our analysis by plotting the count of permits by distance to a boundary between a HSO and a non-HSO city, while modelling the distance as positive

[17] An ADU is a smaller, independent residential dwelling unit located on the same lot as a stand-alone (i.e., detached) single-family home. ADUs go by many different names throughout the U.S., including accessory apartments, secondary suites, and granny flats. ADUs can be converted portions of existing homes (i.e., internal ADUs), additions to new or existing homes (i.e., attached ADUs), or new stand-alone accessory structures or converted portions of existing stand-alone accessory structures (i.e., detached ADUs).

if the permit is in the city with a HSO and negative otherwise. The top panel of Figure 3 shows the count of ADU permits using a 0.05km distance bin size, whereas the bottom panel of Figure 3 focuses on permits excluding ADUs. On the top panel Figure 3, we can see a striking discrete drop in the number of ADU permits on the HSO side of the border. This suggests that restrictive STR regulations dis-incentivize residential investment, at least in the form of ADUs.

[Figure 3 here]

An alternative explanation could be that the city planning departments in HSO cities are more restrictive in all manners, and not just for HSOs. If this were the case, then we would expect to see a similar drop in permits across the border for non-ADU permits. However, the bottom panel of Figure 3 shows that this is not the case. Although the number of non-ADU permits is lower on the HSO side of the border, the effect is not as substantial as for ADUs. Moreover, ADU regulation in California is standardized due to a number of ADU-related laws passed in 2016 and 2017 (AB494, SB1069, and SB229). These laws limit the ability of city planning departments to apply discretionary rules to the ADU permit approval process. Taken together, the above facts suggest that HSOs are responsible for the observed discontinuity in ADU permits across HSO and non-HSO borders, implying a causal effect of STR regulations on residential investment.

4.2 Regression Analysis

In this section, we formalize the above analysis using a regression. To do so, we collapse the data into permit counts by border segment, HSO, year, and month. Namely, for each border segment between a HSO and a non-HSO city, we count all the permits issued within a 1-km bandwidth of that segment by year-month, and by the side of the border (i.e., HSO or non-HSO). We define border segments using a K-means clustering algorithm while constraining each cluster to contain points whose pairwise maximum distance is 3 km.[18] For each permit, we find the nearest point along a border between a HSO and a non-HSO city. We then cluster these border points, resulting in approximately 100 clusters. Each cluster represents a segment of a border between a HSO and a non-HSO city. For each border segment, we count how many permits are on the HSO and the non-HSO sides of the border for each year and month. The unit of analysis is thus the side of a HSO/non-HSO border segment in a year-month.

[Table 6 here]

We let it designate the unit of analysis, where i is the side of the border segment and t is the year-month. We the consider regressions of the following form:

$$\ln(\text{permits}_{it}) = \alpha HSO_i + \gamma_i + \delta_t + \epsilon_{it}, \tag{9}$$

where $\ln(\text{permits}_{it})$ is the log of the number of permits, $HSO_i = 1$ if i is on the HSO side of the border and 0 otherwise, γ_i is a dummy variable for the border segment corresponding to i, and δ_t is a dummy variable for the year and month. The border segment fixed effects, γ_i, capture unobserved spatial heterogeneity at the border segment level. The time dummies, δ_t, capture any unobserved temporal heterogeneity at the year-month level. We run the regression separately for the total number of residential permits, the number of ADU permits, and the number of non-ADU permits. We cluster the standard errors by border segment. Table 6 presents the estimation results. The results show that, after controlling for border segment

[18] The results are consistent when using a smaller or a larger value for the maximum distance between each point belonging to the same cluster.

and time effects, residential permits are 18% lower on the HSO side of the border relative to the non-HSO side (Column 1). A decrease in ADU permits on the HSO side seems to be the driver of this effect, being 17% lower on the HSO side of the border relative to the non-HSO side (Column 2). Finally, non-ADU permits are 9% lower on the HSO side relative to the non-HSO side (Column 3). These estimates are consistent in magnitude with our findings from Section 3, where we found that STR regulations decrease the total number of residential permits by 23%.

[Table 7 here]

Sensitivity analysis. We next show that our results are not sensitive to the length of the bandwidth. In Table 7, we report the results using distance bandwidths of 2 km, 1.5 km, 1 km (our baseline), and 0.5 km.[19] Table 7 shows that our regression results are not sensitive to the choice of the bandwidth length. For conciseness, we only report the results for ADU permits but we draw the same conclusion for non-ADU permits and for total permits.

[Table 8 here]

Placebo test. We present a placebo test aimed at reinforcing the causal interpretation of our results. We create fictitious borders by shifting the real borders by ± 2 km inside or outside the HSO city. We then recalculate the count of permits within a 1-km distance of the fictitious border segments. Table 8 reports the results for four levels of border shifts (-2 km, -1 km, +1 km, +2 km). When relying on the fictitious borders, we no longer estimate any statistically significant decline in ADU permits across borders. As before, for conciseness, we only report the results for ADU permits but the conclusion is the same for non-ADU permits and for total permits.

[19] Counts are normalized by the length of the bandwidth to account for the change in land area. Since the buffer zones are rectangular, the land area scales linearly with the bandwidth distance.

5 Effect on Home Sales and Prices

As we argued in the introduction, the reduction in residential investments associated with the introduction of STR regulations should be due to a reduction in the economic value of and the demand for residential housing. This means that STR regulations should lead to fewer home sales and lower home prices. Here, we test this prediction by estimating the effect of STR regulations on home sales and home prices using the same identification strategy used in Section 3.[20]

As discussed in Section 2, we obtained data for transaction volumes from the Cherre data, whereas for house prices, we use the Zillow ZHVI, which is publicly available at the zipcode-month level. The overall treatment effects using these variables as outcomes are reported in Table 9. We find that the introduction of STR regulations is not significantly associated with transaction volumes (Column 1), but we estimate a treatment effect of 3.3% on house prices (Column 2).

[Table 9 here]

In Figure 4, we present the effects by length of exposure. We do not see much of an effect on transaction volumes (either pre-trend or post-treatment), but we observe parallel pre-trends for house prices followed by a divergence post-treatment. Specifically, over the first year following the policy, house prices are reduced by about 1.2% and by 5.7% over 36 months.

These results show that STR regulations reduced the economic value and demand for residential real estate, which translated into reduced house prices. These results are also aligned with the large body of work demonstrating a causal effect of STRs

[20]We lag house prices by one month to account for the time difference between agreement date and closing date.

on house prices.[21]

[Figure 4 here]

6 Conclusion

In this paper, we show that home-sharing platforms like Airbnb can increase residential investments in the cities they enter. We do so using two empirical strategies: a difference-in-differences strategy that exploits the differences in the timing of adoption of home-sharing regulations across cites, and a border strategy that exploits differences in home-sharing ordinances across municipalities in the Los Angeles County. A back-of-the-envelope calculation suggests that the reduction in residential permits due to home-sharing regulations is associated with an annual loss in property value of $3 billion across the 15 cities we analyze. Moreover, we demonstrate that home-sharing increases residential investments by boosting the demand and the value of residential housing. Specifically, we show that home-sharing regulations lead to reduced house prices, consistent with the prior literature.

Overall, our analysis shows that municipalities can enjoy a clear and steady revenue stream through the increased residential investments generated by home-sharing. This additional revenue may be then directed toward social programs, such as programs related to affordable long-term rentals for residents. This calls into question the need to strictly regulate home-sharing, which can instead be leveraged by local governments as a tool to encourage investments and urban renewal.

Today, this seems more important than ever. To survive the recent economic downturn caused by COVID-19 and prosper, cities need to attract both visitors and

[21] See Horn and Merante (2017); Garcia-López et al. (2020); Barron et al. (2021); Valentin (2021); Koster et al. (2021).

residents by (i) looking appealing; (ii) providing a large volume of attractive accommodations; and (iii) be attractive destinations for businesses and large-scale events. Home-sharing can play a critical role in all three of these goals. In this paper, we mainly focus on (ii) and to some extent on (i). Specifically, we empirically demonstrate that homeowners are willing to invest more in making their real estate assets attractive to short-term renters. Then, the short-term rental market elasticity enabled by home-sharing rentals makes (iii) also possible. Thus, home-sharing is already working for the city's needs—it is now time for policy makers to recognize the benefits of home-sharing and to lighten its regulations for every city to flourish.

References

Almagro, Milena and Tomás Domınguez-Iino, "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," 2019.

Àngel Garcia-López, Miguel, Jordi Jofre-Monseny, Rodrigo Martínez-Mazza, and Mariona Segú, "Do Short-Term Rental Platforms Affect Housing Markets? Evidence from Airbnb in Barcelona," *Journal of Urban Economics*, 2020, *119.*

Barron, Kyle, Edward Kung, and Davide Proserpio, "The Effect of Home-Sharing on House Prices and Rents: Evidence from Airbnb," *Marketing Science*, 2021, *40* (1), 1–191.

Basuroy, Suman, Yongseok Kim, and Davide Proserpio, "Estimating the impact of Airbnb on the local economy: Evidence from the restaurant industry," *Available at SSRN 3516983*, 2020.

Bibler, Andrew J, Keith F Teltser, and Mark J Tremblay, "Inferring tax compliance from pass-through: Evidence from Airbnb tax enforcement agreements," *Review of Economics and Statistics*, 2018, pp. 1–45.

Bogin, Alexander N. and William M. Doerner, "Property Renovations and Their Impact on House Price Index Construction," Working Paper 17-02, FHFA 2017.

Calder-Wang, Sophie, "The Distributional Impact of the Sharing Economy on the Housing Market," *Job Market Paper*, 2019.

Callaway, Brantly and Pedro Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2020.

Einav, Liran, Chiara Farronato, and Jonathan Levin, "Peer-to-peer markets," *Annual Review of Economics*, 2016, *8*, 615–635.

Fradkin, Andrey, "Search, matching, and the role of digital marketplace design in enabling trade: Evidence from airbnb," 2017.

Horn, Keren and Mark Merante, "Is Home Sharing Driving Up Rents? Evidence from Airbnb in Boston," *Journal of Housing Economics*, 2017.

Jain, Shomik, Davide Proserpio, Giovanni Quattrone, and Daniele Quercia, "Nowcasting gentrification using Airbnb data," *Proceedings of the ACM on Human-Computer Interaction*, 2021, *5* (CSCW1), 1–21.

Koster, Hans R.A., Jos van Ommeren, and Nicolas Volkhausen, "Short-Term Rentals and the Housing Market: Quasi-Experimental Evidence from Airbnb in Los Angeles," *Journal of Urban Economics (forthcoming)*, 2021.

Valentin, Maxence, "Regulating Short-Term Rental Housing: Evidence from New Orleans," *Real Estate Economics*, 2021, *49*, 152–186.

Zervas, Georgios, Davide Proserpio, and John W Byers, "The rise of the sharing economy: Estimating the impact of Airbnb on the hotel industry," *Journal of marketing research*, 2017, *54* (5), 687–705.

Table 1: Summary statistics.

City	#Zipcodes	#Listings	#Permits	#Transactions	Avg. ZHVI
Austin	48	45,645	445,950	495†	299,259
Boston	17	12,858	25,971	27,646	619,877
Chicago	59	42,018	145,106	174,473	272,232
Columbus	31	5,238	190,687	156,006	150,234
Denver	32	19,765	176†	215,646	349,396
Las Vegas	27	18,935	38,486	293,585	208,679
Los Angeles	70	89,777	285,766	181,645	725,113
Nashville	28	20,440	6,622	116,763	275,610
New Orleans	26	22,676	145,119	58,294	253,137
New York	100	129,578	63,649	71,887	1,064,439
Portland	32	16,427	5,031	176,348	374,226
San Diego	42	36,644	244,238	186,153	544,619
San Francisco	31	44,253	198,829	77,998	1,125,682
San Jose	29	11,179	6,854	127,902	754,239
Seattle	36	26,629	48,550	152,855	531,461
TOTAL	608	542,062	1,851,034	2,017,696	452,438

† *Note:* Because of a lack in data quality, Denver is excluded from any analysis using permits and Austin is excluded from any analysis using transaction counts.

Table 2: Short-term rental regulations.

City	Effective date (yyyy/mm)
Austin	2012/08
Boston	2019/01
Chicago	2016/06
Columbus	2018/07
Denver	2017/01
Las Vegas	2018/12
Los Angeles	2019/07
Nashville	2015/07
New Orleans	2017/01
New York City	2016/10
Portland	2014/07
San Diego	—
San Francisco	2015/02
San Jose	2014/12
Seattle	2019/09

Table 3: Los Angeles County cities that implemented home-sharing ordinances.

City	Effective date (yyyy/mm)
Arcadia	2017/07
Beverly Hills	2014/09
Burbank	2014/06
Calabasas	2018/01
Cerritos	2016/08
Hermosa Beach	2016/06
Lawndale	2017/07
Manhattan Beach	2015/06
Maywood	2018/04
Palos Verdes Estates	2016/09
Pasadena	2017/10
Rancho Palos Verdes	2016/07
Redondo Beach	2016/06
Rolling Hills	2016/12
Rolling Hills Estates	2016/12
Santa Monica	2015/06
Torrance	2016/04
West Hollywood	2015/09

Table 4: Overall treatment effect of regulation on Airbnb listings and permits.

	Dependent variable:	
	ln(listings)	ln(permits)
	(1)	(2)
Overall Treatment Effect (θ_o)	-0.2135^{***}	-0.1040^{***}
	(0.0303)	(0.0321)

Note: *p<0.1; **p<0.05; ***p<0.01

Table 5: OLS and IV regressions of residential permits on Airbnb listings.

	Dependent variable:			
	ln(permits)		permits	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
ln(listings)	0.071^{***}	0.769^{***}		
	(0.009)	(0.195)		
listings			0.004	0.092^{***}
			(0.003)	(0.028)
ln(population)	-0.264^{***}	-0.006	-8.009	0.195
	(0.076)	(0.207)	(5.209)	(6.649)
pct bachelors	0.086	-1.363^{*}	16.055	-52.051^{*}
	(0.291)	(0.741)	(10.155)	(28.072)
unemployment	-0.753^{***}	1.987^{*}	-28.668^{**}	-22.701
	(0.276)	(1.102)	(12.740)	(16.603)
ln(median income)	0.105	0.684^{***}	1.573	-5.695
	(0.068)	(0.232)	(3.447)	(4.947)
Zipcode FE	✓	✓	✓	✓
Year-Month FE	✓	✓	✓	✓
Observations	51,017	51,017	51,017	51,017

Note: *p<0.1; **p<0.05; ***p<0.01

Table 6: Border discontinuity regression results.

	Dependent variable:		
	ln(permits)	ln(ADU permits)	ln(non-ADU permits)
	(1)	(2)	(3)
HSO	-0.184^{***}	-0.165^{***}	-0.090^{*}
	(0.031)	(0.030)	(0.046)
Border segment FE	✓	✓	✓
Year-month FE	✓	✓	✓
Observations	1,946	1,946	1,946
R^2	0.291	0.437	0.359

Note: *p<0.1; **p<0.05; ***p<0.01

Table 7: Sensitivity of border discontinuity to bandwidth size.

	Dependent variable:			
	ln(ADU permits)			
BW:	<2km	<1.5km	<1km	<0.5km
	(1)	(2)	(3)	(4)
HSO	-0.183^{***}	-0.178^{***}	-0.165^{***}	-0.186^{***}
	(0.025)	(0.027)	(0.030)	(0.050)
Border segment FE	✓	✓	✓	✓
Year-month FE	✓	✓	✓	✓
Observations	2,498	2,301	1,946	1,291
R^2	0.542	0.476	0.437	0.412

Note: *p<0.1; **p<0.05; ***p<0.01

Table 8: Border discontinuity placebo tests.

	Dependent variable:			
	ln(ADU permits)			
	-2.0km	-1.0km	+1.0km	+2.0km
	(1)	(2)	(3)	(4)
HSO (Placebo)	−0.071	−0.037	−0.038	−0.053
	(0.048)	(0.024)	(0.037)	(0.037)
Border segment FE	✓	✓	✓	✓
Year-month FE	✓	✓	✓	✓
Observations	282	953	2,121	1,749
R^2	0.674	0.550	0.447	0.514

Note: *p<0.1; **p<0.05; ***p<0.01

Table 9: Overall treatment effect of regulation on home sales and prices.

	Dependent variable:	
	ln(tx vol)	ln(price)
	(1)	(2)
Overall Treatment Effect (θ_o)	-0.0474	-0.0326^{***}
	(0.0267)	(0.0053)

Note: *p<0.1; **p<0.05; ***p<0.01

Figure 1: Treatment effects by length of exposure.

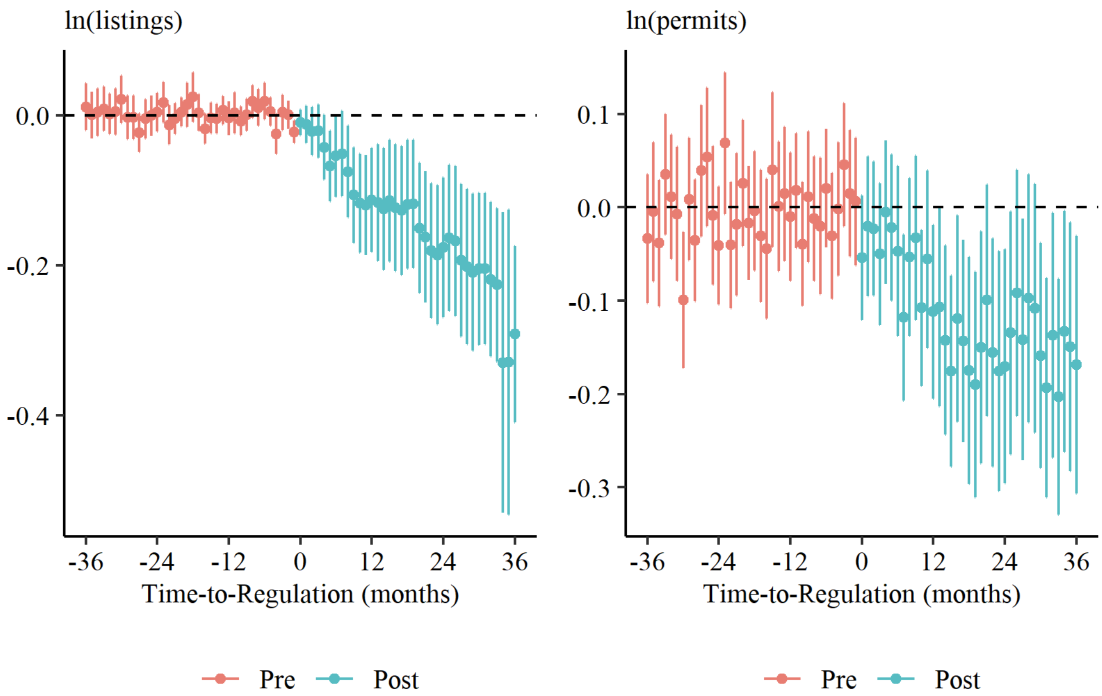

Figure 2: City borders and residential permits in the LA County (2018–2019).

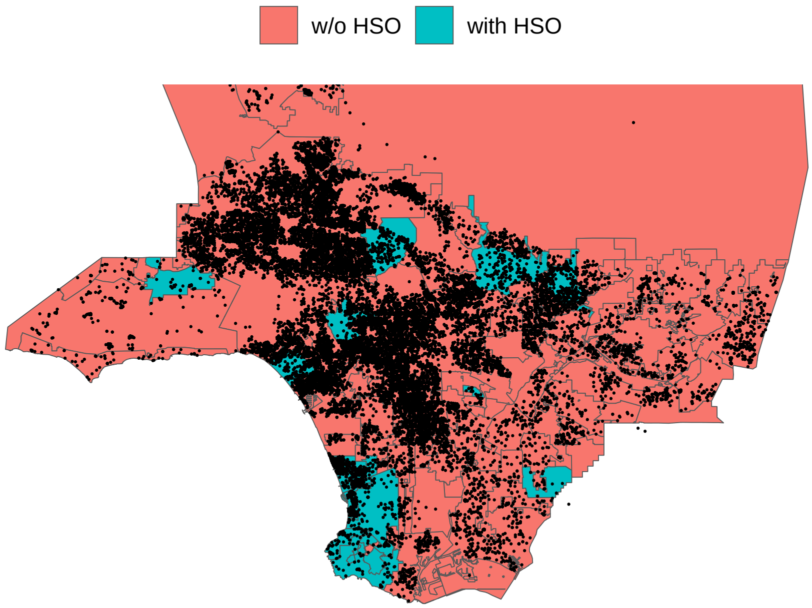

Figure 3: Residential permits across non-HSO/HSO borders.

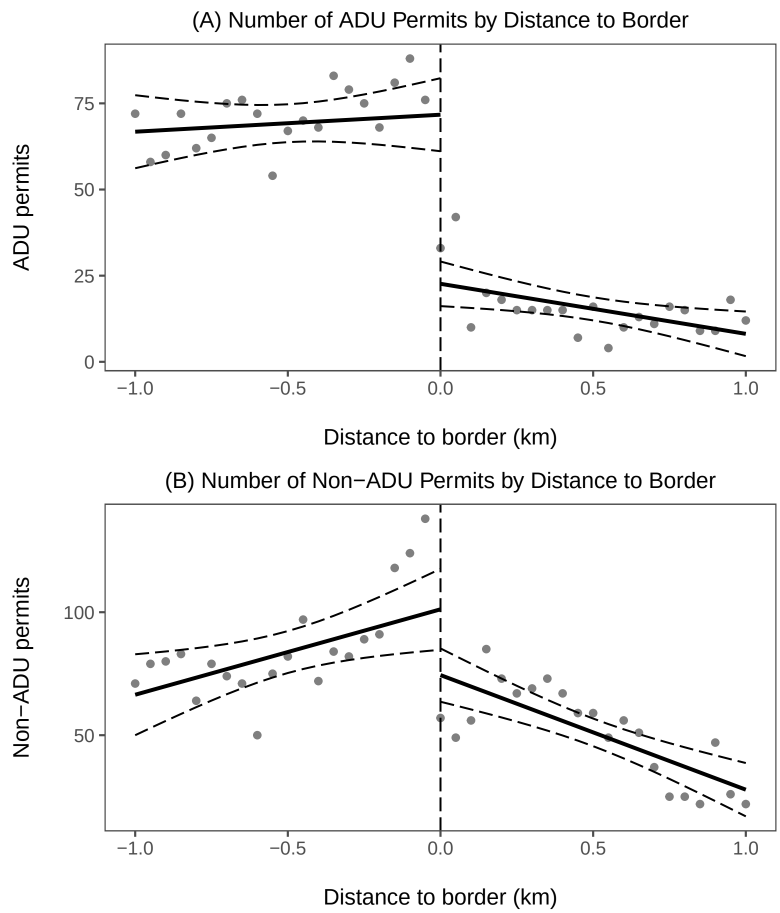

Figure 4: Treatment effects by length of exposure (transaction volume and prices).