The Effects of Occupancy Taxes on the Short-Term Rental Market: Evidence from Boston

Undergraduate Honors Thesis

Alan Liang

May 2020

Advisor: Conrad Miller

Department of Economics

University of California, Berkeley

Abstract

This paper analyzes market price and quantity changes caused by a newly introduced occupancy tax in the short term rental market. We conduct a comparative case study on Boston and a synthetic control through the synthetic controls method originally proposed by Abadie et al. (2001). The paper determines a causal effect in the reduction of nights stayed to be around 25% due to a 14.95% occupancy tax introduced in Boston in 2019. From these effects, we find demand in the short term homestay market to be significantly more price elastic than previous analyses of occupancy taxes on hotels.

1 Introduction

1.1 Airbnb and Occupancy Taxes

Airbnb is an American travel-tech firm that facilitates a peer-to-peer online marketplace for short term lodging experiences around the world. Compared to similar firms offering vacation rental services such as VRBO or HomeAway, Airbnb is the largest and most prominent, with more than 7 million listings worldwide and 2 million people staying in one of its listings per night in 2018 (Airbnb 2020). Since its founding in 2008, hosts on the platform have served more than 750 million guests, and the firm has grown at an exponential rate globally.

As a popular alternative to hotels, Airbnb has become subject to occupancy taxes in cities around the world. These taxes are applied to short term lodgings and are typically enforced by local or state governments in the United States. Proponents of the tax argue that existing rules that apply to hotels should also be enforced for Airbnb. Opponents instead argue that these laws do not apply to services created out of excess capacity, and the tax would stifle innovation and growth in the industry (Wilking 2016). Although occupancy taxes have largely been in place in most American cities for decades, Airbnb's original stance was that the inherent small-scale and informal nature of the transactions it facilitated were not subject to the tax.

However, after years of fighting with local legislators, Airbnb retracted this view in late 2013 and announced they "believed that it makes sense for our community to pay occupancy taxes" (Airbnb blog 2013). Hosts were originally expected to remit occupancy taxes directly to local governments, but between 2014 and 2015 Airbnb began collaborating with major cities including New York City, Portland, and San Francisco to help facilitate the collection and remittance of the tax from its hosts (Kaplan & Nadler 2017). This change forwarded the statutory incidence of the tax from producers to consumers: guests became responsible for paying the up-front cost of the tax when a reservation was booked on Airbnb.

1.2 Boston

Despite beginning to remit taxes in cities that had passed occupancy tax legislation applying to short term rentals, Airbnb actively attempted to lobby against similar legislation in other cities. Such a bill would only come into effect in 2019 for Massachusetts, which is the setting of our natural experiment. In the week following Christmas in 2018, Governor Charlie Baker signed into law *An Act Regulating and Insuring Short-Term Rentals*, which formalized existing short-term rental activity, mandating hosts to register with the state and to provide insurance for their listings. The law further extended the state's existing room occupancy excise tax to short term rental bookings, imposing a 5.7% statewide excise on all reservations beginning July 1st 2019, in addition to a local option excise ranging from 0 to 6.5% and a convention center finance fee of 2.75% that only applied in select Massachusetts cities to finance the construction of convention centers statewide. Notably, the bill came into effect almost immediately on January 1st, allowing two different natural experiments to be conducted. The procedure will be discussed in more detail in section 3.

In Boston, the total occupancy tax amount for Airbnb listings set forth by the act was 14.95%. This included a 5.7% statewide excise tax, a 6.5% local option excise, and a 2.75% convention center finance fee. Immediately after the introduction of the tax, Airbnb worked with the state and local governments to help remit the tax on behalf of hosts. Figure 1 details the cost breakdown guests see when booking a listing on Airbnb's website, including the tax. Before the COVID-19 pandemic, Airbnb was expected to collect $27.5 million in fiscal year 2020 due to the tax on behalf of the state (Lannan 2019).

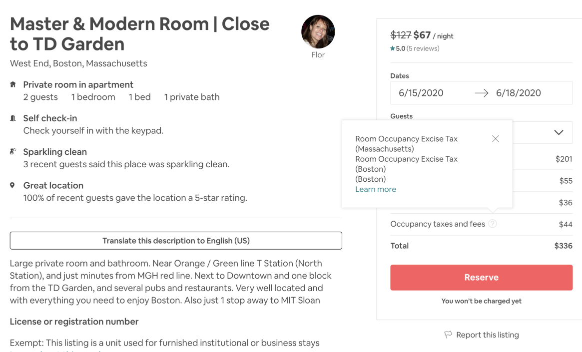

Figure 1: Airbnb reservation interface in April 2020 for a listing in Boston

This paper will analyze the effects of levying an occupancy tax on the short term rental market booked through Airbnb in the city of Boston, a sizable market with approximately 3000 listings. It will evaluate the causal impact of the tax in terms of changes in the quantity transacted (nights reserved) and transaction price. From these results, the paper will determine the price elasticity of demand due to the introduction of the occupancy tax, which can potentially provide valuable insight to policymakers on the consequences and efficiencies of such taxes for short term rentals in other cities around the world.

2 Literature

Clay and Collins (2016) note that "although hotel taxes are now a common revenue source for state and local governments, they have been relatively little studied." There exist a handful of papers that examine of occupancy taxes on hotel demand based on policies enacted in different states and time frames. Despite the heterogeneity of these 'natural experiments', the studies generally conclude that hotel demand is price inelastic and taxes are largely passed on to hotel guests.

One of the earliest ex-post studies on the effects of occupancy taxes was Bonham et al.'s analysis (1992) of Act 340 in Hawaii, which imposed a 5% transient accommodation tax beginning in 1987. Through examining supply side revenues, the authors find that "the effects of the 1987 Hawaii hotel room tax did not have a significant negative impact on hotel rental receipts," but make a caveat that their findings may not be generalizable to other locations. Nonetheless, their findings imply that most of the occupancy tax's economic incidence is passed on to guests, indicating that demand — largely from tourists — is relatively price inelastic.

Through a national survey of the members of the American Hotel and Motel Association (AH&MA), Hiemstra and Ismail (1992) also conclude that price elasticities towards hotel occupancy taxes are indeed inelastic. Their analysis breaks up the heterogeneity of hotels into different sizes and price points, and indicate that the price elasticity of demand for smaller hotels (< 150 rooms) is larger in magnitude than that for larger hotels. In addition, across hotel sizes, hotels with higher prices also have price elasticities that are greater in magnitude. These findings suggest that the price elasticity of demand may be greater for Airbnb units given their much smaller scale.

Clay and Collins (2016) analyze the effects of a \$5 per night hotel tax imposed in Georgia in 2015 through a panel data setting. The model they construct shows that the tax approximately decreased total bookings by 2.7%, implying a consumer elasticity of about -0.7. In addition, Clay and Collins observe that although the introduction of the tax did have a \$1.5 decrease in daily prices, the result was not statistically significant.

Wilking (2016) studies the shift in statutory incidence of an occupancy tax from producers onto consumers on Airbnb, when the platform began to remit occupancy taxes from guests directly on behalf of the hosts. She concludes that "shifting the legal obligation to remit hotel taxes from small, independent hosts to Airbnb increases after-tax prices paid by consumers," but the magnitude of the effect is heterogeneous across hosts and listings. Notably, Wilking's results contradict the traditional theory that the incidence of a consumption tax is exclusively determined by market-wide demand and supply elasticities, and that factors such as the assignment of the remittance obligation may affect the incidence in practice.

This paper will examine the changes in price and quantity from introducing a transient occupancy tax on short term homestays, and thus determine the corresponding price elasticities of consumers in response to the occupancy tax. The study outlined differs from previous literature focusing on hotels as homestays typically operate at much smaller scales, and serve more casual travelers. These factors combined may cause demand responses to be more price elastic in response to an occupancy tax.

3 Data

3.1 *Inside Airbnb* and Data Schema

The data for this paper originate from *Inside Airbnb*, an independent investigatory project that collects and hosts substantial Airbnb data on more than 100 cities around the world. The data collected by *Inside Airbnb* are web-scraped from the Airbnb website on a roughly monthly basis. Although *Inside Airbnb* was originally started by its creators to investigate the effects of Airbnb on affordable housing and exacerbating gentrification, the data are made public for free and open for use.

Since the data on *Inside Airbnb* are collected through web scraping the Airbnb website, they only contain information that a visitor to Airbnb's site can see. Approximately every month, data are scraped for all listings in a particular city. For each listing, its availabilities for the next 365 days and quoted price per night before tax, service fees, and cleaning fees over the next year are collected. In addition, information particular to the listing such as its location, host, review count, and room type are also collected.

3.2 Imputing transactions

One key limitation of utilizing the web-scraped data from *Inside Airbnb* is that they do not reflect actual bookings and transactions made on the platform. For example, it is impossible to discern whether some dates for a unit may be 'blocked out' (unavailable to reserve) due to the unit being occupied, or it being unavailable on the market. Thus, I resort to imputing the booking from the changes in listing availabilities across scraping dates. To impute transaction quantity and price, I rely on 2 key assumptions:

1. A unit that becomes unavailable across 2 consecutive scraping periods is due to the fact that it was booked. This assumption allows me to impute a unit's dates that it was booked and the approximate time period, on the month level, when a booking was made. For example, if a scraping occurring on April 1st reported the listing as available on June 1st and the next scraping occurring on May 1st reported the listing was no longer available on June 1st, we assume that the listing is reserved for June 1st and booked during the April-May scraping period. However, this assumption will likely somewhat overestimate the number of actual transactions. Fradkin (2017) finds that stale vacancies, the phenomena in which hosts don't promptly block specific dates on a listing's calendar even though it is not available, accounts for 15.3% of first-contacts by guests on Airbnb.

2. The price per night of a booking is the same price as that recorded when the listing was most recently available. For example, an imputed transaction with reservation date June 1st that was booked between April and May will have its transaction price imputed to be the price as seen in April for June 1st. This may not be completely true: it is possible that hosts may decrease their prices very near the reservation date in order to maintain high occupancy, and these changes would not be detected by our periodic scraping. However, this limitation should not bias our analysis of changes in the transaction price, since we can reasonably assume that the same hosts adjust their

prices dynamically in the same manner (in terms of magnitude and recency) before and after an occupancy tax is introduced.

As a 'sanity check', I validate the number of imputed nights of a listing based on its review count by constructing a linear fixed effects model. The results are generally consistent and can be found in the appendix. In addition, there is one major flaw in this model of imputation: reservations made between 2 scraping periods that also have the check-in date between the same period will go undetected. For example, if a scraping occurred on June 15th and the next scraping occurred on July 15th, a reservation that was made on June 17th for the night of June 18th will not be detected. As a result, the analysis describes an average treatment effect on generally non last-minute bookings. May (2016) estimates that 30% of reservations on Airbnb are made more than 30 days in advance, which sets a lower bound for our estimates, but this number doubles to 60% by 15 days before and 80% by 10 days before the check-in date. Thus, we can roughly expect the transaction imputation method employed to capture around two-thirds of all reservation bookings on Airbnb. Since it is likely for guests who book far in advance on Airbnb to have greater magnitudes of price elasticity, the elasticity estimates determined in this paper subsequently may be larger than the overall price elasticity.

3.3 Restrictions in Analysis

To better satisfy the first assumption listed above, I attempt to restrict my analysis to units that are dedicated to Airbnb rentals. My hypothesis is that units dedicated to Airbnb rentals will typically only become unavailable between 2 scraping periods due to the listing being booked. However, this assumption may be less dependable for units that are also the primary residence of the host, since hosts who list their primary residences on Airbnb are more likely to change the dates of availability due to changes from their personal schedules. Hence, only listings with more than 100 nights available in the next year, at any stage of scraping, are considered. The cutoff of 100 nights was chosen to be a good tradeoff point between popular listings and primary residences. It is reasonable to assume that primary residences are available for at most 3 months in a year; while for very popular listings, an educated guess for the max annual occupancy rate is around 70%, so that approximately 30%, or 100 nights a year, are available. Furthermore, hosts in the treatment group are exempted from the tax if they declare their intention to rent out the unit for no more than

14 days in a calendar year beforehand, which would likely remove units in the treatment group that are generally unavailable.

I further restrict my analysis to listings that were both on the market in the month immediately before and after the treatment effect. This removes units that become deleted before the treatment was administered and units that enter after the treatment was administered, accounting for attrition and entry.

3.4 *Reservation Date* vs. *Date Booked*

Throughout this paper, there is an important yet subtle distinction between the terms *reservation date* and *date booked*. Specifically, I define the *date booked* as when the reservation is made, while the *reservation date* is the check-in date of the reservation. For example, if a reservation for an Airbnb unit was made on April 1st for July 1st to the 7th, its reservation dates are July 1st to July 7th, but its date booked is April 1st. Although in the long run the effects of taxation affect both equally, in the short run when a tax is introduced it is more clear to separate these and investigate them both.

This distinction becomes particularly important in investigating the treatment effect in our natural experiment. The Massachusetts state legislature enacted the legislation for all reservation dates after July 1st, 2019, but this came into effect beginning on January 1st, 2019. Hence, only a reservation made on or after January 1st 2019 for a check-in date on or after July 1st was subject to a tax; if a reservation for after July 1st was booked before January 1st, or if a reservation was made after January 1st but for before July 1st, it was not subject to the tax. Thus, I will conduct 2 separate sets of analyses on the effects of the tax:

1. The effect of the occupancy tax for all reservations that were booked after January 1st. In this study, the treatment was 'administered' effectively on July 1st, where the panel describes reservation dates.

2. The effect of the occupancy tax on all bookings with reservation dates on or after July 1st. In this study, the treatment was 'administered' on January 1st, where the panel describes reservations' date booked. Since we are only able to narrow down the date booked to be some time between 2 scraping periods in an approximate period of a month, there are only 5 observations in each of the pre and post treatment periods.

> Notably, the proximity between when the tax was signed into law (December 28th, 2018) and its enforcement date (January 1st, 2019) largely rules out potential tax avoidance by guests looking to book their reservations before the tax came into effect.

As a shorthand, the first set of analyses will be denoted as the *reservation date* analysis, while the second set will be denoted as the *date booked* analysis.

4 Empirical Framework

4.1 The Synthetic Control Method

To determine the causal effect of the introduction of an occupancy tax, I utilize the synthetic control method (SCM). Proposed initially by Abadie and Gardeazabal (2001) to study the effects of terrorism on GDP in the Basque country, the synthetic control method is a causal inference method used in comparative case studies in which one unit receives some treatment while others do not. In order to determine the causal effect of treatment, the method creates a non-negative linear combination of control units that represents the counterfactual potential control outcome of the treated unit known as the 'synthetic control'. SCM best fits the synthetic control by minimizing the Euclidean norm of the difference between the synthetic control and the treatment unit in the pre-treatment period subject to certain constraints. The method can be expressed as the following constrained optimization problem:

$$\begin{aligned} &\min_{\gamma} ||X_0'\gamma - X_1||_2^2 \\ \text{subject to } &\sum_i \gamma_i = 1 \\ &\gamma_i \geq 0 \quad \forall i \end{aligned}$$

Where X_1 is a vector denoting outcome values across time of the treated unit in the pre-intervention period, and X_0 is a matrix denoting outcome values of the control units in the pre-intervention period.

The package I use to calculate Synthetic Controls is `AugSynth`, an R package that implements the *augmented synthetic control method* (ASCM), developed by Ben-Michael et al at UC Berkeley. Although this paper does not utilize ASCM, the package implements both plain synthetic controls as well as synthetic controls with ridge regression, which is described below.

4.2 Ridge Regression

Ridge regression is a common regularization procedure to reduce overfitting, trading model variance for added bias. Ridge regression augments the minimization of sum of squared residuals with an added regularization parameter:

$$||X_0'\gamma - X_1||_2^2 + \lambda||\gamma||_2^2$$

Due to the relatively small amount of time periods in the pre-treatment, I utilize ridge regression to prevent over-fitting, keeping weights more stabilized. Ridge regression is implemented in the synthetic controls method by penalizing absolute deviations in weight from the uniform weight value. For example, if there are 10 controls, the penalization value for a weight of 0.15 would correspond to $0.15 - 0.10 = 0.05$. Shen (2018) proposes that introducing regularization in the form of ridge regression can decrease model variance, thus reducing the post-treatment error bound despite increasing pre-treatment error.

Notably, introducing ridge regression deviates from the original synthetic controls method by allowing for negative weights of control units in the synthetic control. Ben-Michael et al (2018) show through simulation that this relaxation of constraints reduces the bias to ensure better balance if the potential control outcome $Y_1(0)$ is indeed a linear combination in the control outcomes.

To tune the coefficient penalization hyperparameter λ, I utilize leave-one-out cross validation by leaving a pre-treatment time period out in each fold. The λ value with the lowest average error across all folds is chosen for fitting. Since the hyperparameter determined through cross validation is not exact or necessarily optimal, I tune one set of hyperparameters for each outcome (nights, log nights, and prices) for the analysis, which is then consistently applied across different models for better stability and replicability.

4.3 Controls

Ten control cities in North America were used to construct the synthetic control. Unlike conventional control units, these cities were already all ‘treated’ before the experiment began, i.e. occupancy tax laws for Airbnb had already been implemented in place in all of the cities (Boston was relatively late to implement such a tax). Regardless, since these cities did not see a change in the treatment effect throughout the experiment time frame, they are suitable to construct a synthetic control if we assume that their post-treatment outcomes (in terms of quantity and price), some time after treatment had been applied, vary similarly with Boston's pre-treatment outcomes. The ten cities were selected based on proximity to Boston and data availability on *Inside Airbnb*, which had data for around 30 cities/regions in North America. Due to data availability issues, Rhode Island and Chicago were dropped from the date booked study, while Quebec City was dropped from the reservation date study. Table 1 lists basic information about each control city's Airbnb market and their Airbnb occupancy tax rates.

Table 1: Control Cities

	# units	Avg price/night[1]	Est. nights/year[2]	Tax %
Asheville†‡	2170	168	162	13
Boston†‡	**3507**	**179**	**153**	**14.95**
Chicago†	8533	181	136	10.5
Washington DC†‡	9369	206	97	14.8
Denver†‡	5222	155	147	10.95
Montreal†‡	19495	109*	88	3.5
Nashville†‡	5921	213	147	8
Quebec City‡	3240	111*	124	3.5
Rhode Island†	2758	277	107	13
San Francisco†‡	7072	213	153	14
Twin Cities†‡	6680	185	62	14

[1] Data based on scraped *listing* prices.
[2] Data imputed by *Inside Airbnb*.
* Price in Canadian Dollars (CAD).
† In reservation date study.
‡ In date booked study.

We observe that Airbnb markets differ widely in terms of size: Boston tends towards the smaller end, as neighboring cities such as Cambridge are not included. Average prices across cities are somewhat similar, with Canadian cities being slightly cheaper and also quoted in Canadian dollars. The average price per night does not simply reflect rental-market prices, but also the type of units available: for example, a significant portion of Nashville and Rhode Island's listings are entire homes or apartments (as opposed to private rooms or shared rooms), hence increasing the average price. In particular, Boston's occupancy tax is the highest out of the cities examined, partly due to the application of both state and local taxes, in addition to convention center fees.

5 Results

5.1 By Reservation Date

We begin our examination of the occupancy tax's effects on listing reservation dates, with the treatment effect beginning on July 1, 2019. Only reservations with dates of booking in 2019 or after are of interest, since the tax had only come into effect on January 1st, 2019. In practice, the analysis begins on February 18th, 2019, which was the latest February scraping date out of all cities. This is because reservation dates are imputed from differences in-between scraping periods, so that the earliest imputable date after January 1st was immediately after the January-February scraping across all cities. On the other end, the analysis ends on December 31st, 2019. We examine the two outcomes of the (imputed) number of nights and price separately by constructing multiple synthetic control models. These analyses are conducted on log outcome scales to determine proportional changes of the tax, and are examined at two granularities: on a monthly and semi-monthly basis. For ridge SCM results, the plots for cross validation errors across different λ values are in the appendix. In addition, the treatment effect for each model is the average predicted outcome across post-treatment periods. Table 2 reflects the results of the analysis; results that are statistically significant ($p = 0.05$) are marked with an asterisk.

We observe that the introduction of the occupancy tax decreased the number of nights transacted by around 25%, and was statistically significant both in the ridge and non-ridge cases at both frequencies. For log prices, the analysis indicates the effect of the tax to be around a 5% decrease in pre-tax prices before potential service fees and cleaning fees. It is important to note that this change in prices is not necessarily due to producers responding to the tax by decreasing listing prices, but perhaps from consumers simply opting for cheaper listings. In addition, the models do not find the results to be statistically significant. Across log differences semi-monthly and monthly, the treatment effects are generally similar for both prices and nights, providing some robustness to the results.

We can calculate price elasticities of demand from the changes in quantity transacted and price. For example, in the monthly non-ridge analysis, a price hike of 10.15% (from

Table 2: Treatment Effect on Reservation Dates

	SCM	Ridge SCM
Number of Nights		
Monthly	-0.269^{*}	-0.238^{*}
	(0.118)	(0.077)
Semi-Monthly	-0.245^{*}	-0.236^{*}
	(0.105)	(0.084)
Price[1]		
Monthly	-0.048	-0.032
	(0.067)	(0.033)
Semi-Monthly	-0.050	-0.024
	(0.046)	(0.069)
Elasticity		
Monthly	-2.65	-2.11
Semi-Monthly	-2.46	-1.88

[1] Based on pre-tax per-night price of listing, before service and cleaning fees.

14.95% - 4.8%) is additionally paid by consumers. With a quantity decrease in the number of nights by 26.9%, this in turn indicates that the consumer elasticity is around 2.65. Note that since the analysis is conducted on an aggregate market level, the price elasticities of demand capture that of the aggregate market, including effects from both quantity and quality margins. Remarkably, the findings are not consistent with the conclusions of inelastic demand from previous ex-post studies of occupancy taxes on hotels. The direction of this difference in elasticities however matches our intuition, since Airbnb consumers are likely to be more tourist-dominated and thus more elastic to changes in price.

Figure 2 visualizes the analysis by plotting the difference between the treated unit and the synthetic control across time. In both outcomes, we observe a decrease in log nights and prices beginning after the intervention since the synthetic control represents the potential control outcome of the treated unit. In addition, we could visually check for robustness of the synthetic control by examining whether it fits the treated unit well in the pre-treatment periods; the differences in outcomes should be constant and near zero before intervention.

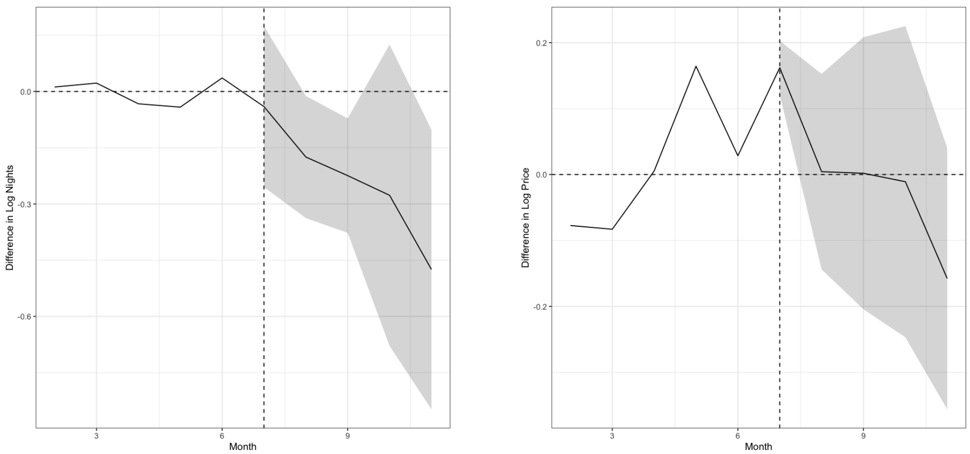

(a) Change in Log Nights

(b) Change in Log Price

Figure 2: Synthetic Control Method Plots Across Time (Monthly, No Ridge)

5.2 Weights of the Synthetic Control

We can examine the weights that composed the synthetic control as a check to support or refute our intuition. For the sake of brevity, we will only focus on the synthetic control analyses on the month granularity.

Table 3: Synthetic Control Weights (Month Granularity)

	Nights (SCM)	Price (SCM)	Nights (Ridge)	Price (Ridge)
Asheville	0.00	0.00	0.015	−0.186
Chicago	0.00	0.129	0.072	0.240
Washington DC	0.00	0.00	0.069	0.075
Denver	0.127	0.00	0.152	−0.087
Montreal	0.00	0.00	0.085	0.104
Nashville	0.449	0.641	0.259	0.434
Rhode Island	0.091	0.00	0.056	0.133
San Francisco	0.270	0.00	0.258	−0.021
Twin Cities	0.064	0.230	0.034	0.309

Table 3 shows that the regular synthetic control method tends to shrink weights of ‘irrelevant’ control cities to 0, while using ridge regression does not. The weights confirm ridge regression’s behavior that it outputs control weights with less deviation from the uniform weight scenario (in this case, 1/9 would be the uniform weight value). However, ridge regression outputs negative weights for control cities, which Abadie et al. (2010) argue are not as interpretable and come at the cost of extrapolation.

Cities that contribute significantly to the non-ridge synthetic control when examining the number of nights are Nashville, San Francisco, and Denver. These weights are similar to those with ridge regression, but the penalty in weights cause cities like Chicago and Montreal, which had not contributed to the non-ridge synthetic control, to now exert a minor influence. Nashville, San Francisco, and Denver have a somewhat similar number of Airbnb listings as Boston, and also have quite similar average estimated occupancies per listing per year. Furthermore, they also have very similar population sizes of around 700,000 residents.

Except for Nashville, a somewhat different set of cities create the synthetic control when examining changes in prices. This discrepancy between weights can be partially explained by the fact that the quantity transacted in a market may be due to different factors that those for changes in price. Notably, only 3 cities were selected in the non-ridge analysis: Chicago, Nashville, and Twin Cities. These 3 cities had somewhat similar average listing prices per night.

5.3 Robustness of Results and Potential Issues

5.3.1 L2 Imbalance

The L2 imbalance describes the extent to which the synthetic control matches the treated unit's pre-treatment outcomes. The L2 imbalance measures the "root sum of squared errors" in the pre-intervention period:

$$\text{L2 imbalance} = \sqrt{||X_0'\gamma - X_1||_2^2}$$

To get a better relative sense, we can compare the L2 imbalance of the fitted synthetic control with the L2 imbalance in the case if uniform weights were used to see the method's improvement.

The larger L2 imbalance in price outcomes confirms the prediction plots in figure 2, which display a relatively 'jumpy' and poor fit of log price in the pre-intervention period. This hence indicates that the synthetic control may not be a robust interpolation of the potential control outcome of Boston. Notably, the L2 imbalance may be reduced from

Table 4: L2 Imbalance in Reservation Date Analysis (Month Granularity)

	SCM	Ridge SCM
Number of Nights		
L2 imbalance	0.069	0.096
% improvement over uniform weights	83.9%	77.7%
Price		
L2 imbalance	0.202	0.177
% improvement over uniform weights	67.2%	71.3%

conducting ridge regression, as seen in the case of log price, since the constraint on non-negative weights is relaxed.

5.3.2 Refitting Weights

We attempt to see if a similar treatment effect can be observed if we utilized the weights from the analysis for log nights into the analysis for log price, and vice versa. Fitting weights from the SCM analysis on log nights (Table 3, Column 1) onto log price yielded a decrease of 0.020, less than half the 0.048 decrease reported in Table 1. This result had a statistically insignificant t-value of 0.28. On the other hand, fitting weights from the SCM analysis on log price (Table 3, Column 2) onto log nights yielded a decrease of -0.580, more than double the reported effect of -0.269 in Table 1. In both cases, as depicted in Figures 3 and 4, the synthetic control was relatively poorly fit in the pre-treatment period, casting less confidence on results from the robustness check. Nonetheless, this check shows the signs of the treatment effect are consistent and confirm our intuition for the occupancy tax; the downwards trend can be seen in both figures below.

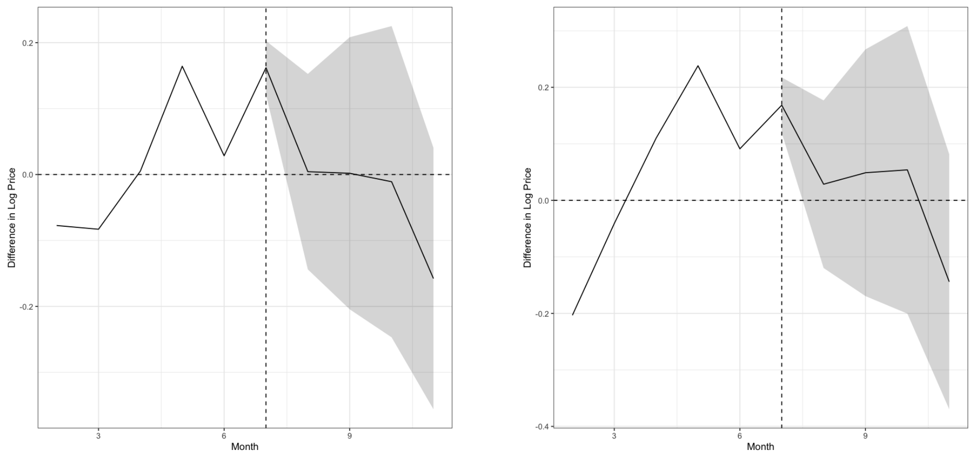

(a) Log Price (Fitted on Log Price's Weights) (b) Log Price (Fitted on Log Nights' Weights)

Figure 3: Comparison of Log Price Outcomes on Different Fitted Weights

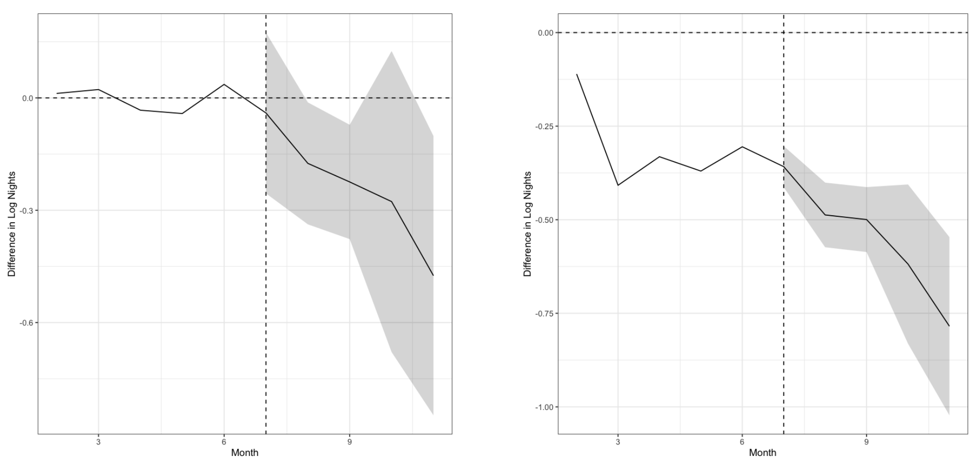

(a) Log Nights (Fitted on Log Nights' Weights) (b) Log Nights (Fitted on Log Price's Weights)

Figure 4: Comparison of Log Nights Outcomes on Different Fitted Weights

5.3.3 Potential Issues

One potential issue with the analysis is that plain SCM may overfit the synthetic control in the pre-treatment period, which only has 5 periods of data. This can be seen in the somewhat different weights for the synthetic control across different models, and could lead to results that may not be as robust or replicable.

I attempt to mitigate this by introducing ridge regression, which reduces overfitting in the pre-treatment periods. However, ridge regression is justified if the synthetic control is distributed around the case of uniform weights, which may not entirely hold. Another method utilized to decrease overfitting was to conduct the analysis on a semi-month level, thus increasing the number of pre- and post-intervention periods. The model in this analysis exhibited similar treatment effects as those on the month level, further adding some extent of validity to the findings. However, as seen in Figure 5, the fluctuations in changes throughout semi-monthly periods are significantly noisier, making the effects of intervention less evident. It would be ideal to extend the analysis to include more pre- and post-treatment periods, although factors including the lack of data and treatment effects being introduced in control cities in previous periods inhibited me from adding more periods.

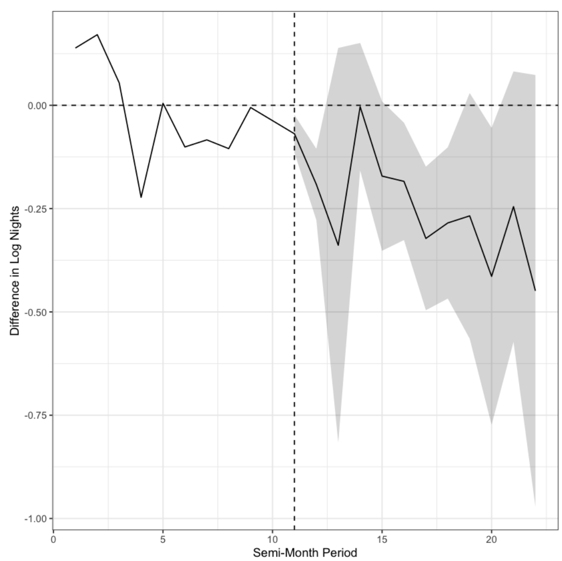

Figure 5: Synthetic Control Prediction (Semi-Month Granularity)

A somewhat opposite issue arises in the analysis of changes in log price. The model produces a synthetic control that does not well capture the price movements of Boston in the pre-intervention period. This implies that the synthetic control may not be a reliable construction of the potential control outcome of Boston, and thus may not generalize well into the post treatment period.

One interesting scenario to consider is if a host conducts "multi-channeling" by posting their listing on other short term homestay sites such as VRBO or Homeaway, which also began remitting taxes on behalf of hosts around when the tax was introduced. Hence, a booking on another site could cause the host to mark the corresponding dates as unavailable on Airbnb as well. This means that the method employed in this paper to impute transactions will also pick up on this change and assume it to be an Airbnb transaction. As a result, it is possible that the treatment effect outlined in the analysis describes the effect of the tax on not just Airbnb listings but of the entire short-term homestay market.

5.4 By Date Booked

A parallel set of analyses is conducted for the effects of the occupancy tax on the booking dates of reservations occurring after July 1, 2019, with the intervention beginning on January 1, 2019. Since Airbnb only shows listing availabilities up to 1 year advance, the analysis begins in July 2018 and extends until June 2019. However, due to the limitation of the method, the date booked can only be narrowed down to be between 2 scraping periods, which is approximately a one-month time frame. As a result, only a month level analysis can be conducted.

Table 5: Treatment Effect on Date of Booking

	SCM	Ridge SCM
Number of Nights	-0.200	-0.327^{*}
	(0.346)	(0.148)
Price[1]	-0.127	-0.064
	(0.251)	(0.147)
Elasticity	-8.89	-3.82

[1] Based on pre-tax per-night price of listing, before service and cleaning fees.

The analyses indicate effects in the same direction as those outlined in table 2. Furthermore, the magnitude of the effects, when examining log outcomes, is in the ballpark of that in the reservation date study. It is also likely that elasticities are indeed higher in this case, which could be explained by guests being able to adjust their plans months in advance.

The nature of the experiment leads to large variance in the predictions, since the number of nights booked with check-in dates after July 1st grows exponentially as the period gets closer to July 1st. Thus, the analysis does not possess sufficient statistical power to determine any results of statistical significance. The exception here is the result in the number of nights in the ridge SCM case, but this may simply be due to having a "fortunate" penalization hyperparameter. Since most results are not of statistical significance, not much additional insight can be concluded from the analysis and will not be examined for robustness.

6 Conclusion

We conducted a comparative case study on the changes in price and quantity of a newly introduced occupancy tax on short term homestay units in Boston. Through conducting multiple synthetic control models, the paper determined a causal effect in the reduction of reservation nights of around 25% due to a 14.95% occupancy tax. A causal effect in the reduction of transaction prices due to the tax was estimated to be around 5%, but this was result was not statistically significant. From these effects, the paper found the price elasticity of demand in the short term homestay market to be quite elastic, which is much larger in magnitude than results from previous literature examining the effects of occupancy taxes on hotels.

The findings presented in this paper could be valuable to policymakers. As the demand price elasticity is high, introducing or increasing occupancy taxes would likely result in relatively large deadweight losses and hence be less efficient compared to instituting the same tax in the hotel market. One potential direction of interest to policymakers is the effects of occupancy tax on overall tourism, and consequently the local economy.

More research and analysis is required to improve the external validity of the results. Conducting more comparative case studies, or utilizing multi-outcome models using different cities around the world in different time periods could potentially help validate or disprove the treatment effects uncovered in this study. One additional direction for further research of interest to policymakers is whether hosts respond to the occupancy tax similarly. It is conceivable that professional Airbnb hosts may differ in their response to the tax from smaller hosts.

References

Abadie, A., Diamond, A., & Hainmueller, J. (2007). Synthetic control methods for comparative case studies: Estimating the effect of californias tobacco control program. https://doi.org/10.3386/t0335

Abadie, A., & Gardeazabal, J. (2001). The economic costs of conflict: A case-control study for the basque country. https://doi.org/10.3386/w8478

Airbnb blog. (2013). https://blog.atairbnb.com/who-we-are/

An act regulating and insuring short-term rentals. (2018). https://malegislature.gov/Laws/SessionLaws/Acts/2018/Chapter337

Ben-Michael, E., Feller, A., & Rothstein, J. (2018). The augmented synthetic control method.

Bonham, C., Fujii, E., Im, E., & Mak, J. (1992). The impact of the hotel room tax: An interrupted time series approach. *National Tax Journal*, *45*(4), 433–441. https://ntanet.org/NTJ/45/4/ntj-v45n04p433-41-impact-hotel-room-tax.pdf?v=%CE%B1&r=5943970666266978

Collins, C. G., & Stephenson, E. F. (2018). Taxing the travelers: A note on hotel tax incidence. *Journal of Regional Analysis Policy*, *48*(7). https://jrap.scholasticahq.com/article/3701.pdf

Cox, M. (n.d.). Inside airbnb.

Fast facts. (n.d.). https://news.airbnb.com/fast-facts/

Fradkin, A. (2017). Search, matching, and the role of digital marketplace design in enabling trade: Evidence from airbnb. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2939084

Hiemstra, S. J., & Ismail, J. A. (1992). Occupancy taxes: No free lunch. *Cornell Hotel and Restaurant Administration Quarterly*, *33*(5), 84–89. https://doi.org/10.1177/001088049203300529

Kaplan, R. A., & Nadler, M. L. (2015). Airbnb: A case study in occupancy regulation and taxation. *University of Chicago Law Review Online*, *82*(1). https://chicagounbound.uchicago.edu/cgi/viewcontent.cgi?referer=&httpsredir=1&article=1038&context=uclrev_online

Lannan, K. (2019). New rental tax expected to generate $27.5 million in revenue. The Patriot Ledger, Quincy, MA. https://www.patriotledger.com/news/20191204/new-rental-tax-expected-to-generate-275-million-in-revenue

May, K. (2016). Number-crunching reveals best time to book on airbnb. PhocusWire. https://www.phocuswire.com/Number-crunching-reveals-best-time-to-book-on-Airbnb

Shen, D. (2018). *Robust synthetic control* (Doctoral dissertation).

Wilking, E. (2016). Hotel tax incidence with heterogeneous firm evasion : Evidence from airbnb remittance agreements.

Acknowledgement

I would like to thank Professor Miller for his time and expertise in advising this thesis. Throughout the research process, he provided valuable insight and impactful suggestions that helped guide my research. I would also like to thank Professor Tang for early guidance during *ECON H195A* that allowed me to discover and select this topic. In addition, my gratitude goes to Eli Ben-Michael for his mentorship in the synthetic control method and much beyond. Lastly, I would like to thank my sister Adeena for helping me collect the massive amounts of Airbnb data used in this paper, and Michael Young, Zhenyu Lai, Umar Maniku, and Andrei Caprau for editing and proofreading this paper. All errors are my own.

Appendix I. Panel Data Results for Estimating Imputed Nights from Reviews

A fixed effects least-squares regression model that regresses the number of imputed nights of a listing based on the number of its reviews is constructed to validate the imputation of nights:

$$\text{Imputed Nights}_{it} = \beta_0 + \beta_1 \text{Reviews}_{it} + \delta_t + \alpha_i + \epsilon_{it}$$

Where i refers to a listing unit, and t refers to a month-level time period. Thus, both entity effects and time effects are incorporated in the model.

Dep. Variable:	nights	**R-squared:**	0.0155
Estimator:	PanelOLS	**R-squared (Between):**	0.1589
No. Observations:	11925	**R-squared (Within):**	0.0155
Date:	Thu, May 07 2020	**R-squared (Overall):**	0.1382
Time:	00:37:02	**Log-likelihood**	-3.926e+04
Cov. Estimator:	Clustered		
		F-statistic:	163.83
Entities:	1531	**P-value**	0.0000
Avg Obs:	7.7890	**Distribution:**	F(1,10393)
Min Obs:	1.0000		
Max Obs:	11.000	**F-statistic (robust):**	101.49
		P-value	0.0000
Time periods:	11	**Distribution:**	F(1,10393)
Avg Obs:	1084.1		
Min Obs:	782.00		
Max Obs:	1342.0		

	Parameter	**Std. Err.**	**T-stat**	**P-value**	**Lower CI**	**Upper CI**
reviews	0.5135	0.0510	10.074	0.0000	0.4136	0.6134

The model implies that each review corresponds to approximately 2 imputed nights per listing, which is roughly consistent with estimations from *Inside Airbnb* that the average duration of a reservation is 3-4 nights, and the review rate is around two-thirds for each reservation.

Appendix II. Cross Validation Plots for Ridge Regression

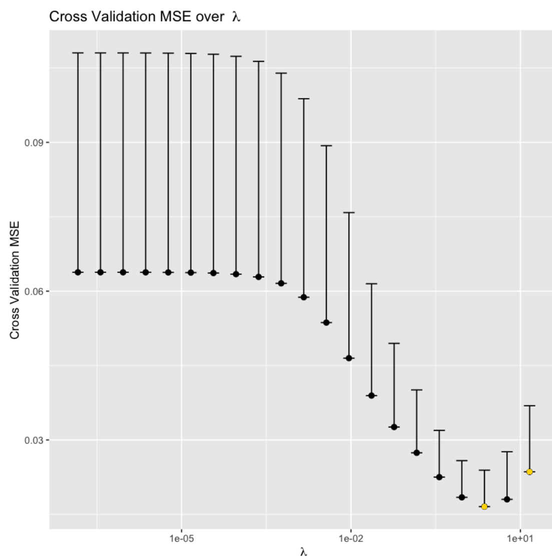

Figure 6: Mean squared error over different λ values for log nights prediction

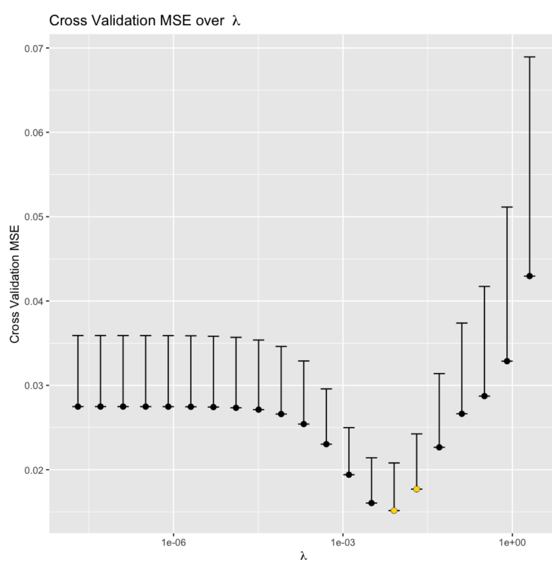

Figure 7: Mean squared error over different λ values for log price prediction