Do Foreign Buyer Taxes Affect House Prices?*

Jonathan S. Hartley, Li Ma, Susan Wachter and Albert Alex Zevelev†

Stanford, Baruch, Wharton, Baruch
October 14, 2022

Abstract

This paper studies the impact of foreign buyer taxes on house prices using recent law changes in Canada, Australia, and New Zealand. Counterfactual house prices are estimated for each treated location combining prediction techniques from machine learning with inference methods from the Synthetic Control Method literature. In general, the foreign buyer taxes we studied had negative, large, and persistent effects on house price growth. We find bigger effects in locations with bigger taxes and with higher immigrant shares. Alternative outcome variables, including population growth, GDP growth, and unemployment rates were either unaffected or slightly affected in ways that do not confound our results.

*We thank Sonia Gilbukh, Arpit Gupta, Joseph Gyourko, Cameron LaPoint, Matteo Maggiori, Waldo Ojeda, Andrey Pavlov, Monika Piazzesi, Stijn Van Nieuwerburgh, Chenzi Xu, Yildiray Yildirim, as well as audience members at AsRES-GCREC-AREUEA, Baruch College, and Stanford for helpful comments. All errors are our own. Data and code are available from the authors.

†Email: hartleyj@stanford.edu, li.ma@baruch.cuny.edu, wachter@wharton.upenn.edu, albert.zevelev@baruch.cuny.edu

1 Introduction

Housing affordability is a growing problem in many cities around the world. House prices have grown faster than income in many metropolitan areas.[1] In response, governments have experimented with several policy tools including zoning changes, rent control, housing vouchers, tax credits, and Tobin flip taxes.[2] Out of town and foreign home buyers are often accused of playing a key role.[3] Several cities and countries have recently passed foreign buyer taxes (FBTs) to be paid once at purchase as well as vacancy taxes (VTs) to be paid annually (with property taxes) to try to make housing more affordable.[4] This paper seeks to identify if and to what extent these policies worked in reducing house price growth and increasing affordability.

This paper studies the causal effect of the new taxes on house price growth using the Synthetic Control Method (SCM). For each treated location, a synthetic control group is estimated with elastic-net[5] using a large donor pool of untreated locations during pre-treatment years. We estimate counter-factual house price growth in each treated location using the synthetic control group. The Average Treatment Effect on the Treated (ATT) is equal to the observed house price growth in the treated location minus the estimated counter-factual house price growth. Confidence intervals, for the ATTs, are estimated using the Jackknife+ method.[6]

We study five treated locations: Greater Toronto, Greater Vancouver, Sydney, Melbourne, and New Zealand. The initial treatment in each treated location was an FBT. Two of the five treated locations (Greater Toronto and New Zealand) received exactly one treatment. The other three locations (Greater Vancouver, Sydney, and Melbourne) received follow-up taxes (both FBTs and VTs) after the initial FBT. In all cases, we estimate the synthetic control group until one quarter before the first treatment. Following Abadie and Gardeazabal (2003) the SCM can be used to study the impact of multiple sequential treatments (which occur in three of our five treated locations). The caveat is that in these

[1] Favilukis, Ludvigson, and Van Nieuwerburgh (2017); Gallin (2006); Piazzesi and Schneider (2016)

[2] Chi, LaPoint, and Lin (2020); Favilukis, Mabille, and Van Nieuwerburgh (2019); Pavlov and Somerville (2020); Somerville, Wang, and Yang (2020)

[3] Chao and Eden (2015); Chinco and Mayer (2015); Favilukis and Van Nieuwerburgh (2021); Li, Shen, and Zhang (2020); Sakong (2021)

[4] Duca, Muellbauer, and Murphy (2020) note the limited literature on attempts to control house prices through taxes.

[5] Zou and Hastie (2005); Doudchenko and Imbens (2016)

[6] Barber, Candes, Ramdas, and Tibshirani (2021); Ben-Michael, Feller, and Rothstein (2018)

cases the ATT_t, after the first treatment, should be interpreted as the cumulative effect of treatments, for the treated location, up to time t.

A key advantage of the SCM over traditional Difference-in-Differences (DiD) is that the SCM uses a data-driven approach to construct a more attractive control group (Athey and Imbens, 2017). We compare the SCM estimates with more traditional DiD estimates in the Appendix and find the DiD treatment effects are qualitatively similar while the SCM has much better pre-trends.

We find that a 15% FBT in Greater Toronto, passed in 2017, lowered house price growth up to 28% relative to the counterfactual. Greater Toronto house price growth remained 12% below the counterfactual (if there were no tax) three years later in the first quarter of 2020 (the end of our sample). A 15% FBT in Greater Vancouver, passed in 2016, lowered house price growth up to 34%, and a follow-up 20% FBT in 2018 lowered house price growth up to 44%. Greater Vancouver house price growth remained 30% below the counterfactual three and a half years after the initial tax (two years after the second tax) in the first quarter of 2020. These are quite large capitalization effects given that, in 2016 before the tax, 10% of property transfer taxes in Vancouver involved foreign nationals (who tend to make more expensive purchases).[7]

A 4% FBT in Sydney, passed in 2016, had no statistically significant effect on house price growth, a follow-up 8% FBT in 2017 lowered house price growth up to 20% in early 2019. The effect on the rate of change disappeared at the end of 2019, while the lower level of prices persisted. A 3% FBT in Melbourne, passed in 2015, had no statistically significant effect on house price growth, a follow-up 7% FBT in 2016 (and 8% FBT in 2019) lowered house price growth up to 32% at the end of 2018. Finally, a ban on all foreign buyers for existing home purchases (but not for new homes) in New Zealand, passed in 2018, had no statistically significant effect on house price growth.

Our analysis in Australia indicates that the impact of FBTs on house price growth is a nonlinear function of the size of the tax. Both of the smallest taxes we study (a 3% FBT in Melbourne in 2015 and a 4% FBT in Sydney in 2016) had no statistically significant effect on house price growth.

We also study the impact of the FBTs on several alternate outcome variables: population growth, unemployment rates, and local GDP growth. In most cases, we find these alternative outcome variables were unaffected by the FBTs. In a few cases these variables

[7] CTV News

were slightly affected in ways that work against our main effect. For example, there was a small drop in the unemployment rate in Toronto after the FBT. Since lower unemployment is associated with a stronger economy (and higher house price growth) this works against us finding a drop in house price growth after the FBT. Together, the alternate outcome estimates indicate the impact of the FBTs on house price growth was likely direct and was not likely due to some indirect economic effect.

The key identifying assumption behind a causal interpretation of our estimates is that the FBTs were uncorrelated with unobserved contemporaneous events that would have affected house prices in treated locations differently from untreated locations. For example, suppose Vancouver passed an income tax hike at the same time as the FBT, while untreated locations did not pass such an income tax hike. Such a hypothetical income tax hike might depress the economy and put downward pressure on house price growth. In this case it would be impossible for us to separately identify the impact of the FBT from the income tax hike. We study the institutional settings and do not find such contemporaneous events. In addition, if such confounding events existed, they would likely affect the alternate outcome variables by raising unemployment rates or lowering GDP growth in treated locations. To the contrary, we find the alternative outcome variables were mostly unaffected or slightly affected in ways that would not confound our estimates.

In addition, we conduct placebo tests, by estimating a synthetic control group for each of our 345 untreated locations, for each treatment, using the other 344 locations as controls. We consistently find average out-of-sample placebo effects to be close to zero. Together, our results indicate that our estimates are likely causal effects. However, our analysis indicates that we must be cautious when inferring external validity. For example, while we find that a 15% FBT in Greater Toronto significantly lowered house price growth, we cannot conclude that in general a similar policy would have a similar effect in a different location at a different time. Our results indicate that the impact of an FBT passed in a new location will depend positively and nonlinearly on the size of the tax and the amount of local foreign buyer activity.

Finally, we investigate whether anticipation effects confound our estimates by backdating the intervention as suggested by Abadie (2021). To test this possibility, we refit the outcomes for each treated location up to two quarters before the treatment and present the results in the appendix. These estimates are qualitatively and quantitatively similar to our main estimates (fit up to one quarter pre-treatment) except for Sydney's first FBT which is now slightly statistically significant.

2 Institutional Setting

The institutional details behind the various taxes we study are discussed in this section and summarized in Table 1. Rapidly rising housing prices in recent decades have been studied extensively from the US,[8] to China,[9] leading to the idea of supply-constrained "superstar cities"[10] where land use regulations play a significant role in constraining supply amid increasing demand for real estate, which has become globalized.

Canada has particularly been marked by rapidly rising housing prices in recent decades (with a far smaller decline than the US during the global financial crisis). Canada's housing market has some of the highest price-to-rent and price-to-income ratios in the OECD,[11] leading some to argue that there are housing affordability problems in Canada.

Toronto and Vancouver, two out of Canada's three largest urban areas, have experienced the greatest housing price increases in past decades. These two metropolitan areas also experienced some of the highest amount of immigration, particularly from China. In addition, real estate markets in these cities have experienced significant amounts of non-resident foreign buying. Many of such parcels are also vacant and not rented.

Australia has also witnessed a decades long real estate boom in addition to a three decade long economic boom which saw no contraction until the COVID-19 pandemic, potentially due to economic growth spillovers in neighboring China.[12]

In response, in the 2010s, nearly all Australian states and territories (except for Northern Territory) as well as Toronto and Vancouver have introduced legislation that seeks to improve housing affordability by taxing foreign buyers. New Zealand passed a ban on foreign buyers for existing home purchases (but not for new homes).[13] The Canadian federal government is also planning to implement a nationwide annual vacancy tax of 1%, for foreign buyers, starting January 1, 2022 as of this writing.[14] On March 29, 2022, Ontario

[8] Glaeser, Gyourko, and Saks (2005); Glaeser, Gyourko, and Saiz (2008)
[9] Rogoff and Yang (2020)
[10] Gyourko, Mayer, and Sinai (2013)
[11] `https://data.oecd.org/price/housing-prices.htm`
[12] news.com.au, Bloomberg News, South China Morning Post
[13] Effectively an ∞% FBT for existing home purchases.
[14] Financial Post

announced it would amend the previous FBT to increase the tax rate from 15% to 20%, extend the tax to apply throughout the entire province, and eliminate rebates for international students and foreign nationals working in Ontario. The revised tax became effective the next day, March 30, 2022.[15] None of the taxes passed in 2022 are included in our study since there is not enough post-treatment data to estimate the treatment effects.

An FBT is essentially a one-time tax on the sale price of a real estate parcel purchased by a non-resident buyer. These taxes are in addition to any other duties that a local resident would pay if they were to purchase the same real estate parcel. For example, suppose a foreign buyer were to purchase a home with a list price of $200,000 in a jurisdiction with a 10% FBT. They would pay the full price $200,000 to the seller, in addition to paying $20,000 to the taxing authority, in addition to other local transfer duties.

2.1 British Columbia

The British Columbia Provincial Government, led by Premier Christy Clark, on July 25th, 2016 announced their proposal for a 15% FBT as part of Bill 28 to improve housing affordability in the Greater Vancouver Regional District ("Metro Vancouver"), a political unit encompassing the city of Vancouver and neighboring municipalities. The tax was introduced "as part of legislation aimed at addressing low vacancy rates and high real estate prices in southern British Columbia."[16] The FBT was quickly passed into law and went into effect shortly after on August 2nd, 2016. Residential property purchased by a foreign entity (either a foreign national or a foreign corporation) in Metro Vancouver became subject to a tax equal to 15% of the sale price of the property. This amount was in addition to a baseline Property Transfer Tax that already existed for all real estate transactions. The City of Vancouver implemented a vacancy tax of 1.25% per year on January 1, 2017 "to return empty or under-used properties to use as long-term rental homes for people who live and work in Vancouver."[17]

The FBT tax rate was later increased to its current 20% rate on February 21st, 2018, and the scope of the tax was expanded to include additional regions nearby (Capital Regional District, Fraser Valley Regional District, Regional District of Central Okanagan, and Regional District of Nanaimo). In July 2021, the British Columbia Court of Appeal rejected

[15] EY Canada
[16] CBC News
[17] vancouver.ca

arguments that the FBT is unlawful and discriminatory finding it constitutional.[18]

2.2 Ontario

The Non-Resident Speculation Tax (NRST) was introduced to the public by the Liberal Ontario provincial government of Premier Kathleen Wynne[19] on April 20th, 2017, and came into effect the following day. The NRST applies a 15% tax rate to the sale price of any home sold within the Greater Golden Horseshoe (GGH) region of Ontario, which includes Greater Toronto and surrounding cities. There is also a rebate for homebuyers who become residents within a limited time period after the purchase. The use of Toronto herein refers to the Greater Toronto Area.

2.3 Australia

In 2015, states and territories in Australia began implementing FBTs (often called "Foreign Purchaser Duty Surcharges") in addition to the usual duty payable on the purchase by residents. Australia considers foreign buyers people who are not citizens of Australia or New Zealand. Largely the motivation behind foreign buyers taxes in Australia was to improve housing affordability amid rising housing prices. Some like Western Australia Treasurer Ben Wyatt argued that it was "fair" for foreign owners of real estate to pay a surcharge because they benefited from services and infrastructure.[20] As of today, New South Wales (including Sydney) charges an 8% FBT, and Victoria (including Melbourne) also charges an 8% FBT.

2.4 New Zealand

In New Zealand, the Overseas Investment Amendment Act (OIAA) of 2018 was passed on October 15th, 2018. and came into effect a week later on October 22nd, 2018.[21] This law change bans most non-resident foreigners (besides Australians and Singaporeans) from buying existing houses by classifying them as sensitive land. Notably, the OIAA does not prevent foreigners from buying newly built houses.

[18] Vancouver Sun
[19] Ontario Ministry of Finance
[20] www.internationalinvestment.net
[21] New Zealand Parliament

3 Data, Empirical Strategy, and Identification

3.1 Data

This section discusses the various datasets used in the paper which are listed in Table 2. The main outcome variable we study is house price growth. The alternative outcome variables we study are population growth, unemployment rates, and local GDP growth.

The house price series in this paper correct for the changing quality and characteristics of houses being sold at any point in time by estimating price changes with repeat-sales methods, hedonic regression methods, stratification methods, sale price-appraisal ratio (SPAR) methods, or hybrid methods.[22] The Canadian home price data, from CREA, is based on a hybrid model that merges repeat-sales and hedonic price approaches.[23] The Australian home price data, from the Australian Bureau of Statistics (ABS), is constructed with property sales data supplied to the ABS by CoreLogic using stratification techniques.[24] American home price data are from Zillow,[25] international home data are from the Dallas Fed International House Price Database and the BIS.[26]

House price data are available at two frequencies (quarterly (mostly) and monthly) with varying start dates. We convert all monthly indexes to quarterly and compute year over year house price growth for each quarter. Our combined house price sample has 41 quarters, beginning quarter one 2010 and ending quarter one 2020 (before the COVID pandemic). In the main analysis there are 345 untreated donor locations (all partially treated locations were removed from the donor pool) and five treated locations: Vancouver, Toronto, Sydney, Melbourne, and New Zealand. This paper uses donor locations at the city and country level as in Abadie and Gardeazabal (2003), this does not violate any of the identifying assumptions given in Section 3.3. In the heterogeneity analysis, we compare the effect in fourteen treated locations in Ontario as well as two regional districts in British Columbia.

This paper also studies the mechanism by examining the impact of the FBTs on various economic variables in treated locations. Table 2 lists data sources for population, unemployment rate (UR), and GDP data. The population data sample is at an annual frequency. We construct population growth rates and begin the sample in 2005. The local GDP data

[22] Dorsey, Hu, Mayer, and Wang (2010); Mack, Martínez-García, et al. (2011)
[23] CREA HPI Methodology
[24] ABS Methodology
[25] Zillow Methodology
[26] The Dallas Fed and BIS combine house price indexes from various sources, often from central banks.

is annual and the country GDP data is quarterly. All data are converted to annual growth rates and begin 2001. The UR data are at monthly and quarterly frequencies. We convert all UR data to quarterly unemployment rates and begin the sample in 2011 (the first year with Canadian data).

3.2 Empirical Strategy

This section discusses the empirical strategy. This paper estimates the impact of FBTs on house price growth in treated locations using the Synthetic Control Method (SCM).[27] We use methods described in Doudchenko and Imbens (2016) and Ben-Michael et al. (2018), who generalize the classic SCM to allow nonconvex weights and a permanent additive difference between the treated and control units. They show that this powerful generalization nests many existing approaches as special cases including classical difference-in-differences and matching methods.

The following example can illustrate the advantage of using more general models (with possibly non-convex weights and an intercept) over traditional SCM.[28] Suppose we have data $y_{i,t}$ for three cities $i \in \{A, B, C\}$. If city A's data is in the convex hull of B and C, for example $y_{A,t} = \frac{1}{2}y_{B,t} + \frac{1}{2}y_{C,t}$, then city B's data cannot be in the convex hull of A and C, $y_{B,t} = 2y_{A,t} - y_{C,t}$. A model such as elastic net does not a priori assume a treated city is in the convex hull of the donors, but rather lets the data speak for itself.

Let $Y_{z,t}$ denote the outcome variable for treated location z in year t and $Y_{j,t}$ the outcome for untreated location j. Let $Y_{z,pre}$ and $Y_{j,pre}$ be vectors of the outcome variables in the pre-treatment years. Let $Y_{C,pre}$ be a matrix of predictors whose columns consist of $Y_{j,pre}$ (outcome variables for all control location in the pre-treatment years) as well as other control variables (an intercept, linear, and quadratic time trends). The Doudchenko and Imbens (2016) estimator minimizes the distance between the treated outcome and an affine combination of the untreated outcome for the pre-treatment period, regularized[29] by the elastic-net (en) penalty (Zou and Hastie, 2005):

$$(\hat{\mu}^{en}, \hat{\omega}^{en}) = \underset{\mu,\omega}{\operatorname{argmin}} \left\| Y_{z,pre} - \mu - Y_{C,pre} \cdot \omega \right\|_2^2 + \lambda \cdot \left(\alpha \|\omega\|_1 + (1-\alpha) \|\omega\|_2 \right)$$

[27] Abadie and Gardeazabal (2003); Abadie, Diamond, and Hainmueller (2015); Abadie (2021); Zevelev (2020)

[28] We thank Guido Imbens for this example.

[29] Following Zou and Hastie (2005) the intercept μ is not regularized.

The parameter $\lambda \geq 0$ determines the amount of regularization, and $\alpha \in [0,1]$ determines the type. The case $\alpha = 1$ corresponds to a LASSO penalty function, which captures a preference for parsimony via a small number of nonzero weights. The case $\alpha = 0$ corresponds to a Ridge penalty function, which captures a preference for smaller weights. Doudchenko and Imbens (2016) propose a cross-validation procedure to tune the regularization parameters λ and α that minimize the average mean squared prediction error for all untreated units.

In general, elastic-net cross-validation techniques tune the hyperparameters λ and α in ways that favor parsimonious and sparse models. Importantly, these techniques split the pre-treatment data into training and testing samples, and validate by maximizing the fit on unseen test data that was not used for training. SCM techniques are often motivated by the idea that there are common unobserved factors that explain the treated and untreated outcomes.[30] Regularized and cross-validated techniques provide a data-driven approach to recover the optimal control group, based on out-of-sample fit.

These estimates give the counterfactual outcome for treated location z if it did not receive the treatment as a function of the control locations:

$$\hat{Y}_{z,t}(0) = \hat{\mu}^{en} + Y_{C,t}\hat{\omega}^{en},$$

where $Y_{C,t}$ is a row vector consisting of outcomes for the control locations and potentially other controls (linear and quadratic time trends etc) in year t. The identifying assumption is that the relationship between the treated and control outcome variables, given by $\hat{\mu}^{en}$ and $\hat{\omega}^{en}$, would have remained the same in absence of the treatment. We further discuss, and defend, the identifying assumptions below. The estimated treatment effect[31] for location z is the gap (i.e., difference) between the observed and counterfactual outcome:

$$\hat{\eta}_{z,t} = Y_{z,t} - \hat{Y}_{z,t}(0).$$

This paper estimates 95% confidence intervals (for the ATT) using the Jackknife+ method.[32] The idea behind this approach is that the absolute residual during pre-treatment periods can be used to estimate the unobserved residual post-treatment. Barber et al. (2021) explain a naive approach is to use in-sample prediction errors $R_t^{Naive} = \left| Y_{z,t} - \hat{Y}_{z,t} \right|$

[30] Goetzmann and Wachter (1995); Goetzmann, Spiegel, and Wachter (1998)
[31] In our paper the estimand is the Average Treatment Effect on the Treated (ATT).
[32] Barber et al. (2021); Ben-Michael et al. (2018)

for $t \in \{1, \ldots, T_0\}$ where T_0 is the last pre-treatment period. R_t^{Naive} is likely a bad estimator for $R_{T_0+1} = \left| Y_{z,T_0+1} - \hat{Y}_{z,T_0+1} \right|$ to the extent that in-sample errors tend to be smaller than out-of-sample errors.

Barber et al. (2021) propose using leave-one-out (LOO) pre-treatment residuals instead. For example, fit the model on all pre-treatment data except observation t, and use this model to predict the outcome for the omitted observation $\hat{Y}_{z,t}^{-t}$. Then compute LOO residuals during the pre-treatment period $R_t^{LOO} = \left| Y_{z,t} - \hat{Y}_{z,t}^{-t} \right|$ for $t \in \{1, \ldots, T_0\}$. In addition to using LOO residuals (as in the Jackknife), the Jackknife+ method also uses LOO predictions for $Y_{z,T_0+1}(0)$, denoted $\hat{Y}_{z,T_0+1}^{-t}(0)$, to construct prediction intervals that account for potential regression algorithm instability.[33] This gives us T_0 estimates of the upper/lower bounds of the prediction interval $\left\{ \hat{Y}_{z,T_0+1}^{-t}(0) \pm R_t^{LOO} \right\}_{t=1}^{T_0}$. The Jackknife+ prediction interval, for $Y_{z,T_0+1}(0)$, at a given target level $1-\alpha$, is given by:

$$\hat{C}_{T_0,\alpha}^{\text{jackknife+}}\left(Y_{z,T_0+1}(0)\right) = \left[\hat{q}_{T_0,\alpha}^{-} \left\{ \hat{Y}_{z,T_0+1}^{-t}(0) - R_t^{LOO} \right\}, \hat{q}_{T_0,\alpha}^{+} \left\{ \hat{Y}_{z,T_0+1}^{-t}(0) + R_t^{LOO} \right\} \right]$$

where $\hat{q}_{T_0,\alpha}^{-}\{x\}$ is the α empirical quantile of the vector x, and $\hat{q}_{T_0,\alpha}^{+}$ is the $1-\alpha$ empirical quantile. $\hat{C}_{T_0,\alpha}^{\text{jackknife+}}$ can guarantee finite sample predictive coverage under no assumptions other than exchangeability of the training and test data.

Following Abadie and Gardeazabal (2003) the SCM can be used to study the impact of multiple sequential treatments. The only caveat is that in these cases $\hat{\eta}_{z,t}$ should be interpreted as the cumulative effect of all treatments, in location z, up to time t.

3.3 Identification

There are several levels of interpretation of the estimates in this paper with different assumptions. First, a non-causal interpretation of our estimates, is that house price growth in treated locations *fell by $\hat{ATT}_t$ after the treatment*. The only assumption required for the non-causal interpretation is that the data are measured properly. All of our house price index data come from official sources that correct for the changing compositions of homes sold over time, as explained in Section 3.1, so we believe the non-causal interpretation is likely correct.

Second, a causal interpretation of our estimates, that *the FBTs caused house price*

[33] If things are stable, omitting observation t should have little effect, and $\hat{Y}_{z,T_0+1}^{-t}(0) \approx \hat{Y}_{z,T_0+1}(0)$.

growth in treated locations to fall by $\hat{ATT}_t$, requires additional, stronger assumptions than the non-causal interpretation. The identifying assumption for a causal interpretation (using the SCM estimator) is that the relationship between the outcome and predictors before the treatment, would have remained the same after the treatment, if not for the treatment. The SCM estimator faces similar threats to identification as the Difference-in-Differences (DiD) estimator. One threat to identification is any contemporaneous unobserved variable that affects treated outcomes differently than predictors. For example, suppose there was an income tax hike in Vancouver at the same time as the FBT, but not in control locations. To the extent that the income tax hike would lower house price growth, we could not separately identify the impact of the FBTs from the income tax hike. We study the institutional settings and do not find such contemporaneous events. In addition, if such confounding events existed, they would likely affect the alternate outcome variables by raising unemployment rates or lowering GDP growth in treated locations. To the contrary, we find the alternative outcome variables were mostly unaffected or slightly affected in ways that do not confound our estimates. We believe such a contemporaneous confounding variable is unlikely.

Another threat to identification is a contemporaneous variable that affects predictors differently than treated outcomes. This includes the possibility of spillover effects. For example, suppose the Vancouver FBT caused potential Vancouver foreign homebuyers to prefer homes in untreated locations such as Montreal or London. This could bias our estimates to the extent that Montreal or London house price growth are selected by the elastic-net as significant pre-treatment predictors of Vancouver house price growth. To examine these possibilities we conduct extensive placebo tests, by estimating a synthetic control group for each of our 345 untreated locations, for each treatment, using the other 344 locations plus linear and quadratic time trends as predictors.[34] We consistently find average out-of-sample placebo effects to be close to zero.

Another set of threats to identification are due to model mis-specification pre-treatment. For example if we are missing important predictors, use the wrong functional form, or overfit the model. This is an area where the elastic-net SCM estimator shines. We examine pre-trends and consistently find good pre-treatment fit. In addition, we use a conservative cross-validation technique (Doudchenko and Imbens, 2016) to tune the elastic-net hyperparameters α, λ to avoid overfitting. Another type of mis-specification is due to anticipation.

[34] This approach is valid to the extent that the Vancouver FBT did not cause a change in global housing markets large enough to significantly affect all 345 untreated locations in our sample.

For example, there may have been a rise in demand by foreign buyers to buy Vancouver housing before the FBT is implemented. To test this possibility, we refit the outcomes for each treated location up to two quarters before the treatment and present the results in the appendix. These estimates are qualitatively and quantitatively similar to our main estimates (fit up to one quarter pre-treatment) except for Sydney's first FBT which is now slightly statistically significant. Together, our results indicate that our estimates are likely causal effects.

Third, even if the causal effect is correctly identified, additional assumptions are required to identify the mechanism. The identifying assumption behind a direct mechanism is that the FBTs did not have a significant effect on alternative variables which affect house price growth. For example, suppose the FBTs (and not some other unobserved variable) depressed the Vancouver economy, then it is not clear to what extent house price growth fell due to the direct demand mechanism versus the indirect economic effects. We conduct falsification tests that study the impact of the FBTs on population growth, local GDP growth, and unemployment rates. We find the alternative outcome variables were mostly unaffected or slightly affected in ways that would attenuate the main effect.

A final layer of interpretation is external validity, to what extent do we expect the effects we estimated in our sample will hold for other FBTs, in other places, passed at other times. We take this opportunity to remind the reader that our estimand is the Average Treatment Effect on the Treated (ATT) and not the unconditional Average Treatment Effect (ATE). Hence, while we believe we estimate the causal effects of the taxes in Vancouver, Toronto, Sydney, Melbourne, and New Zealand, we do not make strong claims about external validity. However, our analysis in Australia indicates that the impact of FBTs on house price growth is a nonlinear function of the size of the tax. Both of the smallest taxes in our main sample (a 3% FBT in Melbourne in 2015 and a 4% FBT in Sydney in 2016) had no statistically significant effect on house price growth.

4 Estimates

Figure 1 panel A, plots house price growth in Canada's three biggest metros: Toronto, Montreal, and Vancouver. The untreated city Montreal, the second largest metro in Canada, is included for reference. A few things are immediately clear: 1 there are large drops in house price growth following the enactment of the taxes, 2 the taxes in Vancouver did not appear to affect house price growth in Montreal and Toronto, 3 the tax in Toronto did not appear to affect house price growth in Montreal and Vancouver. It is important to note that these raw data figures do not tell us the causal impact of the corresponding taxes. For that we estimate the counterfactual in the section below.

Figure 1 panel B, plots house price growth in Sydney, Melbourne, and New Zealand. A few observations: 1 the Australian locations each passed multiple FBTs, 2 the impact of the FBTs on house price growth is not obvious. We will need to estimate the counterfactual outcomes to determine if there was an effect.

4.1 Impact of Foreign Buyer Taxes on House Price Growth

This section presents the main results, the impact of foreign buyer taxes on house price growth in Vancouver, Toronto, Sydney, Melbourne, and New Zealand. For each of the five main treated locations, we present three figures: 1 ATTs with 95% confidence intervals, 2 observed outcomes and estimated counterfactuals, 3 placebo ATTs for 345 untreated locations. For each treated location, we also provide a table with ATTs and 95% confidence intervals. For each treated location the (non-zero) weights can be seen in appendix Figure 12.

The first panel of Figure 2 plots the ATTs for Greater Toronto and Table B.3 presents the corresponding estimates. A 15% FBT in Greater Toronto, passed and effective in the second quarter of 2017, lowered house price growth up to 28% in the second quarter of 2018, a year later. Greater Toronto house price growth remained 12% below the counterfactual (if there was no tax) three years later in the first quarter of 2020 (the end of our sample). The first panel of Figure 3 plots the observed and counterfactual house price growth for Greater Toronto. Our model predicts that, after the treatment, Toronto house price growth would have remained almost as high as pre-treatment levels, falling from 30% to 20% toward the end of the sample. The first panel of Figure 4 plots placebo ATTs for 345 untreated locations in gray and the average placebo ATT, each quarter, in green. It is comforting to see that the average placebo ATT is close to zero. Greater Toronto appears to have the biggest negative effect in the sample, with the exception of Serbia 2019-2020. The first panel of Figure 10 plots the ATTs for Greater Toronto except now the pre-treatment sample ends two quarters before the first FBT. The estimates are qualitatively and quantitatively similar to those in the main analysis (Figure 2) where the pre-treatment sample ends one quarter before the first FBT. Hence it is unlikely that anticipation effects played an important role in Greater Toronto.

The second panel of Figure 2 plots the ATTs for Greater Vancouver and Table B.3 presents the corresponding estimates. A 15% FBT in Greater Vancouver, passed and effective in the third quarter of 2016, lowered house price growth up to 34% in the second quarter of 2017, three quarters later. Note that a 1.25% vacancy tax was passed in the first quarter of 2017, but it is not clear if it had much effect as the ATT continued to grow for only one more quarter. A follow-up 20% FBT, passed and effective in the first quarter of 2018, lowered house price growth up to 44% in the second quarter of 2019, five quarters later. Greater Vancouver house price growth remained 30% below the counterfactual (if there was no tax) three and a half years after the initial tax (two years after the second tax) in the first quarter of 2020. The second panel of Figure 3 plots the observed and counterfactual house price growth for Greater Vancouver. Our model predicts that, after the treatment, Vancouver house price growth would have remained similar to pre-treatment levels, between from 30% and 40%. The second panel of Figure 4 plots placebo ATTs for 345 untreated locations in gray and the average placebo ATT, each quarter, in green. It is comforting to see that the average placebo ATT is close to zero. Greater Vancouver has the biggest negative effect for most of the sample. The second panel of Figure 10 plots the ATTs for Greater Vancouver except now the pre-treatment sample ends two quarters before the first FBT. The estimates are qualitatively and quantitatively similar to those in the main analysis (Figure 2) where the pre-treatment sample ends one quarter before the first FBT.

The third panel of Figure 2 plots the ATTs for Sydney and Table B.3 presents the corresponding estimates. A 4% FBT in Sydney, passed in the second quarter of 2016, had no statistically significant effect on house price growth. A follow-up 8% FBT in the third quarter of 2017, and a 2% annual vacancy tax in the first quarter of 2018, lowered house price growth up to 20% in the first quarter of 2019, six quarters after the 8% FBT and one year after the 2% annual vacancy tax. The effect disappeared at the end of 2019. The third panel of Figure 3 plots the observed and counterfactual house price growth for Sydney. Our model predicts that, after the treatment, Sydney house price growth would have slowed relative to pre-treatment levels, and have been much less volatile, but would have remained positive. The third panel of Figure 4 plots placebo ATTs for 345 untreated locations in gray and the average placebo ATT, each quarter, in green. It is comforting to see that the average placebo ATT is close to zero. Sydney had one of the biggest negative effects in 2018 and 2019. The third panel of Figure 10 plots the ATTs for Sydney except now the pre-treatment sample ends two quarters before the first FBT. After 2018, the estimates are qualitatively and quantitatively similar to those in the main analysis (Figure 2) where the pre-treatment sample ends one quarter before the first FBT. However, Sydney's first FBT, passed in 2016, is now statistically significant.

The fourth panel of Figure 2 plots the ATTs for Melbourne and Table B.3 presents the corresponding estimates. A 3% FBT in Melbourne, passed in the third quarter of 2015, had no statistically significant effect on house price growth. A follow-up 7% FBT in the third quarter of 2016, combined with a 1.5% vacancy tax in the first quarter of 2017, lowered house price growth up to 32% in the fourth quarter of 2018, two years and one quarter after the 7% FBT and one year and three quarters after the 1.5% vacancy tax. Two additional small increases, a follow-up 8% FBT (up from 7%) in the third quarter of 2019, and a follow-up 2% (up from 1.5%) annual vacancy tax in the first quarter of 2020, didn't seem to have much effect. The effect attenuated to 10% by the first quarter of 2020, the end of our sample. The fourth panel of Figure 3 plots the observed and counterfactual house price growth for Melbourne. Our model predicts that, after the first FBT (3%) Melbourne house price growth would have been the same as observed, after the second FBT (7%) Melbourne house price growth would have been much higher than observed, reaching a peak of 40% in the second quarter of 2017, and falling to 20% in the first quarter of 2020. The fourth panel of Figure 4 plots placebo ATTs for 345 untreated locations in gray and the average placebo ATT, each quarter, in green. It is comforting to see that the average placebo ATT is close to zero. Melbourne had one of the biggest negative effects, second only to Russia, after the second FBT (7%) in 2016. The fourth panel of Figure 10 plots the ATTs for Melbourne except now the pre-treatment sample ends two quarters before the first FBT. The estimates are qualitatively and quantitatively similar to those in the main analysis (Figure 2) where the pre-treatment sample ends one quarter before the first FBT.

The fifth panel of Figure 2 plots the ATTs for New Zealand, and Table B.3 presents the corresponding estimates. A ban on all foreign buyers for existing home purchases (but not for new homes), passed in the third quarter of 2018, reduced house price growth 2.5%, but the effect was not statistically significant. The fifth panel of Figure 3 plots the observed and counterfactual house price growth for New Zealand. Our model predicts that, after the ban, New Zealand house price growth would have remained approximately the same as before at roughly 5%, about twice as high as the observed house price growth of 2.5%. The fifth panel of Figure 4 plots placebo ATTs for 345 untreated locations in gray and the average placebo ATT, each quarter, in green. It is comforting to see that the average placebo ATT is close to zero. New Zealand did not have one of the biggest negative effects in this sample, over this time period. The fifth panel of Figure 10 plots the ATTs for New Zealand except now the pre-treatment sample ends two quarters before the first FBT. The estimates are qualitatively and quantitatively similar to those in the main analysis (Figure 2) where the pre-treatment sample ends one quarter before the first FBT, except now the ATTs are (barely) statistically significant in 2019.

4.2 Heterogeneity

This section explores if and in what ways the effect differed across different locations which received the same treatment. We study heterogeneity in jurisdictions where multiple locations, that we can observe, received the same treatment: fourteen locations in the Greater Golden horseshoe region of Ontario and two locations in British Columbia.

The first panel of Figure 5 plots the ATTs for fourteen locations (cities, districts, and Census Areas) in the Greater Golden horseshoe region of Ontario. All fourteen treated locations received the same treatment, 15% FBT, in the second quarter of 2017. The pre-trends for the fourteen treated locations are pretty good, and the ATTs are negative but vary widely. We combine the treated locations into three categories and plot the ATTs (averaged within each category) in the second panel of Figure 5. We find the smallest effect in the first category, denoted "small ATT", consisting of three locations: Guelph, Peterborough, and Simcoe. House price growth in these locations fell at most 10%. The "medium ATT" category consists of six locations: Brantford, Cambridge, Hamilton, Kawartha Lakes, Kitchener, and Niagara. House price growth in these locations fell at most 20% (in 2018), and the effect attenuated to 8% in the first quarter of 2020. The "big ATT" category consists of five locations: Barrie, Greater Toronto (from the main analysis), Mississauga, Northumberland, and Oakville. House price growth in these locations fell at most 30% (in 2018), and the effect attenuated to 16% in the first quarter of 2020. The general pattern we observe from this analysis is not surprising: less densely populated locations (Guelph, Peterborough, and Simcoe), which are less globally connected with presumably less global demand, are less affected by the same FBTs.

In addition, we calculate the immigrant share for each treated Ontario location in 2016 (before the FBT) inspired by Pavlov and Somerville (2020) who argue that this is a good measure of foreign buyer activity. Figure 11 presents a scatter plot of immigrant share against the minimum ATT for treated locations. Locations with the greatest share of immigrants (notably Greater Toronto, Mississauga, and Oakville), presumably with the most foreign buyer activity, also had the biggest ATTs.

Next Figure 6 compares the impact of the 2018 20% FBT on house price growth in Fraser Valley and Victoria. Fraser Valley saw up to a -30% ATT, while Victoria saw up to an -18% ATT. Note that Victoria (and not Fraser Valley) also had a separate "Speculation and Vacancy Tax" of 0.5% on December 31, 2018 raised to 2% on December 31, 2019. This provides more evidence of considerable heterogeneity as the effect in Fraser Valley was almost twice as big as in Victoria.

In general, detailed foreign buyer share data is not available for most treated locations. However, British Columbia began collecting related data before the first FBT. According to British Columbia Property Transfer Tax data, two months before the first Vancouver FBT, foreign buyers made 12.3% of home purchases in Vancouver. After the first FBT was implemented in August 2016 (but before the second FBT), the average monthly foreign buyer share was 3.4%. After the second FBT in February 2018, the average monthly foreign buyer share was 2.1%. This indicates the Vancouver FBTs had a large direct effect on home purchases by foreigners.

4.3 Impact of Foreign Buyer Taxes on Alternative Outcomes

In this section we study the impact of the FBTs on three alternative outcomes: population growth, GDP growth, and unemployment rates.

Figure 7 presents the estimates of the impact of the FBTs on population growth for the main locations we study. There was no statistically significant effect for Toronto, Vancouver, and New Zealand. Sydney had a slight positive effect (.25%) in 2017 and Melbourne had a positive effect up to 1% during the post-treatment period. In general population growth is associated with higher housing demand, and thus with higher house price growth. Hence the small positive effects we observe in Sydney and Melbourne would, if anything, work against the drop in house price growth we find.

Figure 8 presents the estimates of the impact of the FBTs on local GDP growth for the main locations we study. There was no statistically significant effect for Ontario (Toronto), British Columbia (Vancouver), Victoria (Melbourne), and New Zealand. New South Wales (Sydney) had a positive effect (1.7%) in 2016. In general GDP growth is associated with higher housing demand, and thus with higher house price growth. Hence the positive effect we observe in New South Wales (Sydney) would, if anything, work against the drop in house price growth we find.

Figure 9 presents the estimates of the impact of the FBTs on unemployment rates for the main locations we study. There was no statistically significant effect for Vancouver and New Zealand. Toronto, New South Wales (Sydney), and Victoria (Melbourne) had slight drops in the unemployment rate. A few of the locations saw a rise in the unemployment rate at the beginning of the COVID pandemic in the first quarter of 2020. In general unemployment is associated with lower housing demand, and thus with lower house price growth. Hence the small negative effects we observe in Toronto, New South Wales (Sydney), and Victoria (Melbourne), would, if anything, work against the drop in house price growth we find.

5 Conclusion

Several locations around the world have passed FBTs in an attempt to slow down house price growth. We find the taxes passed in Toronto, Vancouver, Sydney, and Melbourne had significantly negative and persistent effects. New Zealand's ban on foreign buyers, for existing housing, had no statistically significant effect on house price growth.

There are four levels of identification and interpretation of the results. First, all of our datasets use house price indexes which adjust for the quality of home sold over time. Hence we believe the non-causal interpretation is very safe. Second, we have studied the institutional setting behind the law changes and conducted validation tests examining pre-trends and falsification tests examining the impacts of the FBTs on alternative outcome variables. Taken together, we believe our estimates can be interpreted as causal effects.

Third, the alternative outcome estimates find that the FBTs did not have significant effects on the local economies and populations. This indicates that the mechanism behind our effect is likely direct. Finally, we find heterogeneity in the treatment effects. Consequently, we do not make strong claims about external validity except that we expect larger taxes, passed in locations with more foreign buyer activity to have bigger negative effects on house price growth.

References

Abadie, Alberto, 2021, Using synthetic controls: Feasibility, data requirements, and methodological aspects, *Journal of Economic Literature* 59, 391–425.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, 2015, Comparative politics and the synthetic control method, *American Journal of Political Science* 59, 495–510.

Abadie, Alberto, and Javier Gardeazabal, 2003, The economic costs of conflict: A case study of the basque country, *American economic review* 93, 113–132.

Athey, Susan, and Guido W Imbens, 2017, The state of applied econometrics: Causality and policy evaluation, *Journal of Economic Perspectives* 31, 3–32.

Barber, Rina Foygel, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani, 2021, Predictive inference with the jackknife+, *The Annals of Statistics* 49, 486–507.

Ben-Michael, Eli, Avi Feller, and Jesse Rothstein, 2018, The augmented synthetic control method, *arXiv preprint arXiv:1811.04170* .

Chao, Chi-Chur, and SH Eden, 2015, Housing markets with foreign buyers, *The Journal of Real Estate Finance and Economics* 50, 207–218.

Chi, Chun-Che, Cameron LaPoint, and Ming-Jen Lin, 2020, Flip or flop? tobin taxes in the real estate market, *Tobin Taxes in the Real Estate Market (December 23, 2020)* .

Chinco, Alex, and Christopher Mayer, 2015, Misinformed speculators and mispricing in the housing market, *Review of Financial Studies* 29, 486–522.

Dorsey, Robert E, Haixin Hu, Walter J Mayer, and Hui-chen Wang, 2010, Hedonic versus repeat-sales housing price indexes for measuring the recent boom-bust cycle, *Journal of Housing Economics* 19, 75–93.

Doudchenko, Nikolay, and Guido W Imbens, 2016, Balancing, regression, difference-in-differences and synthetic control methods: A synthesis, *National Bureau of Economic Research Working Paper* 22791.

Duca, John V, John Muellbauer, and Anthony Murphy, 2020, What drives house price cycles? international experience and policy issues, *Journal of Economic Literature* .

Favilukis, Jack, Sydney C Ludvigson, and Stijn Van Nieuwerburgh, 2017, The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium, *Journal of Political Economy* 125, 140–223.

Favilukis, Jack, Pierre Mabille, and Stijn Van Nieuwerburgh, 2019, Affordable housing and city welfare, Technical report, National Bureau of Economic Research.

Favilukis, Jack, and Stijn Van Nieuwerburgh, 2021, Out-of-town home buyers and city welfare, *The Journal of Finance* .

Gallin, Joshua, 2006, The long-run relationship between house prices and income: evidence from local housing markets, *Real Estate Economics* 34, 417–438.

Glaeser, Edward L, Joseph Gyourko, and Albert Saiz, 2008, Housing supply and housing bubbles, *Journal of urban Economics* 64, 198–217.

Glaeser, Edward L, Joseph Gyourko, and Raven E Saks, 2005, Why have housing prices gone up?, *American Economic Review* 95, 329–333.

Goetzmann, William N, Matthew Spiegel, and Susan M Wachter, 1998, Do cities and suburbs cluster?, *Cityscape* 193–203.

Goetzmann, William N, and Susan M Wachter, 1995, Clustering methods for real estate portfolios, *Real Estate Economics* 23, 271–310.

Gyourko, Joseph, Christopher Mayer, and Todd Sinai, 2013, Superstar cities, *American Economic Journal: Economic Policy* 5, 167–99.

Li, Zhimin, Leslie Sheng Shen, and Calvin Zhang, 2020, Capital flows, asset prices, and the real economy: A china shock in the us real estate market .

Mack, Adrienne, Enrique Martínez-García, et al., 2011, A cross-country quarterly database of real house prices: a methodological note, *Globalization and Monetary Policy Institute Working Paper* 99.

Pavlov, Andrey, and Tsur Somerville, 2020, Immigration, capital flows and housing prices, *Real Estate Economics* 48, 915–949.

Piazzesi, Monika, and Martin Schneider, 2016, Housing and macroeconomics, in *Handbook of Macroeconomics*, volume 2, 1547–1640 (Elsevier).

Rogoff, Kenneth S, and Yuanchen Yang, 2020, Peak china housing, Technical report, National Bureau of Economic Research.

Sakong, Jung, 2021, Rich buyers and rental spillovers: Evidence from chinese buyers in us housing markets, *Available at SSRN 3814058* .

Somerville, Tsur, Long Wang, and Yang Yang, 2020, Using purchase restrictions to cool housing markets: A within-market analysis, *Journal of Urban Economics* 115, 103189.

Zevelev, Albert Alex, 2020, Does Collateral Value Affect Asset Prices? Evidence from a Natural Experiment in Texas, *The Review of Financial Studies* hhaa117.

Zou, Hui, and Trevor Hastie, 2005, Regularization and variable selection via the elastic net, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67, 301–320.

A Appendix: Figures

A.1 House Price Growth

Figure 1: House Price Growth

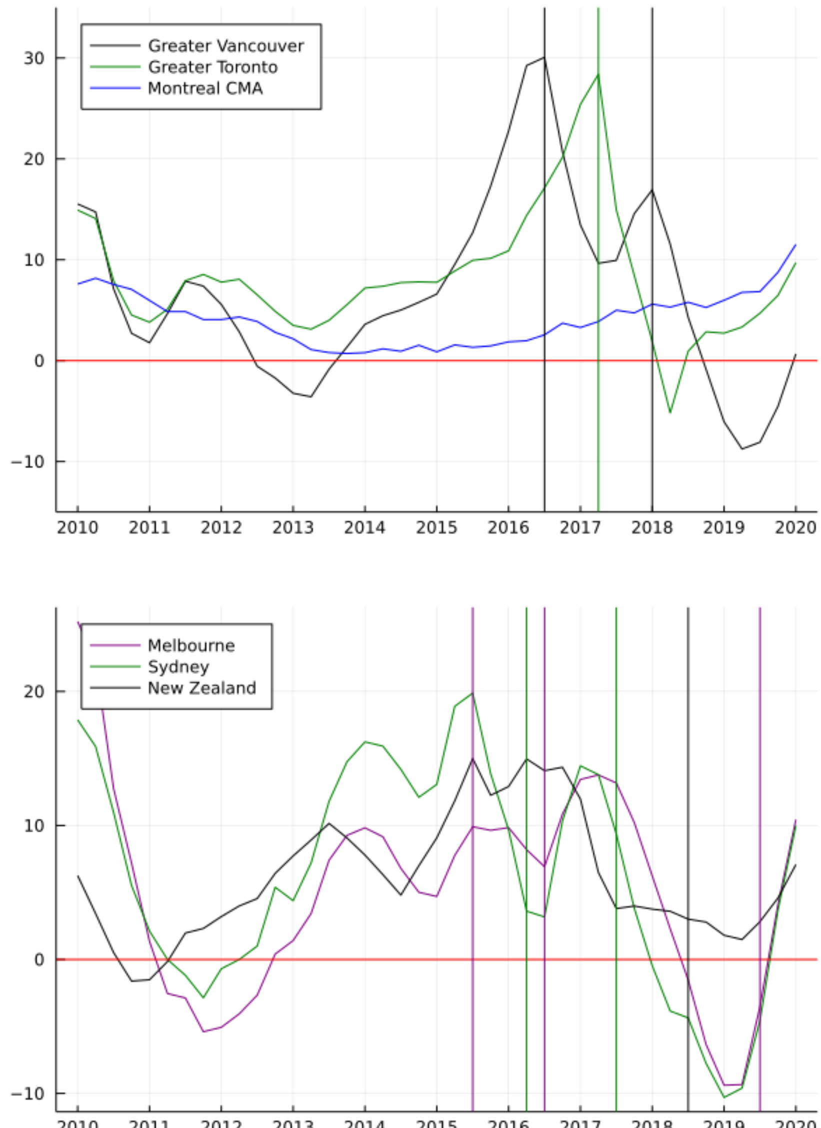

Note.
This figure plots year over year house price growth, for each quarter, for locations described in the legend. Data sources can be found in Table 2. Vertical bars, in the same color as the treated location, correspond to quarters Foreign Buyer Taxes were implemented.

A.2 ATTs with Confidence Intervals

Figure 2: ATTs with Confidence Intervals

Note.
This figure plots point estimates for ATTs and 95% confidence intervals for treated locations described in the legend. ATTs are estimated using the SCM and confidence intervals are estimated using Jackknife+ as explained in Section 3.2. Vertical bars, in red, correspond to quarters before Foreign Buyer Taxes were implemented. Dashed vertical bars correspond to vacancy taxes. Data sources can be found in Table 2.

A.3 Outcomes and Counterfactuals

Figure 3: Outcomes and Counterfactuals

Note.
This figure plots observed outcomes (house price growth) along with estimates for synthetic (counterfactual) outcomes for treated locations described in the legend. Synthetic outcomes are estimated using the SCM as explained in Section 3.2. Vertical bars, in red, correspond to quarters before Foreign Buyer Taxes were implemented. Dashed vertical bars correspond to vacancy taxes. Data sources can be found in Table 2.

A.4 ATT with Placebo Effects

Figure 4: ATT WITH PLACEBO EFFECTS

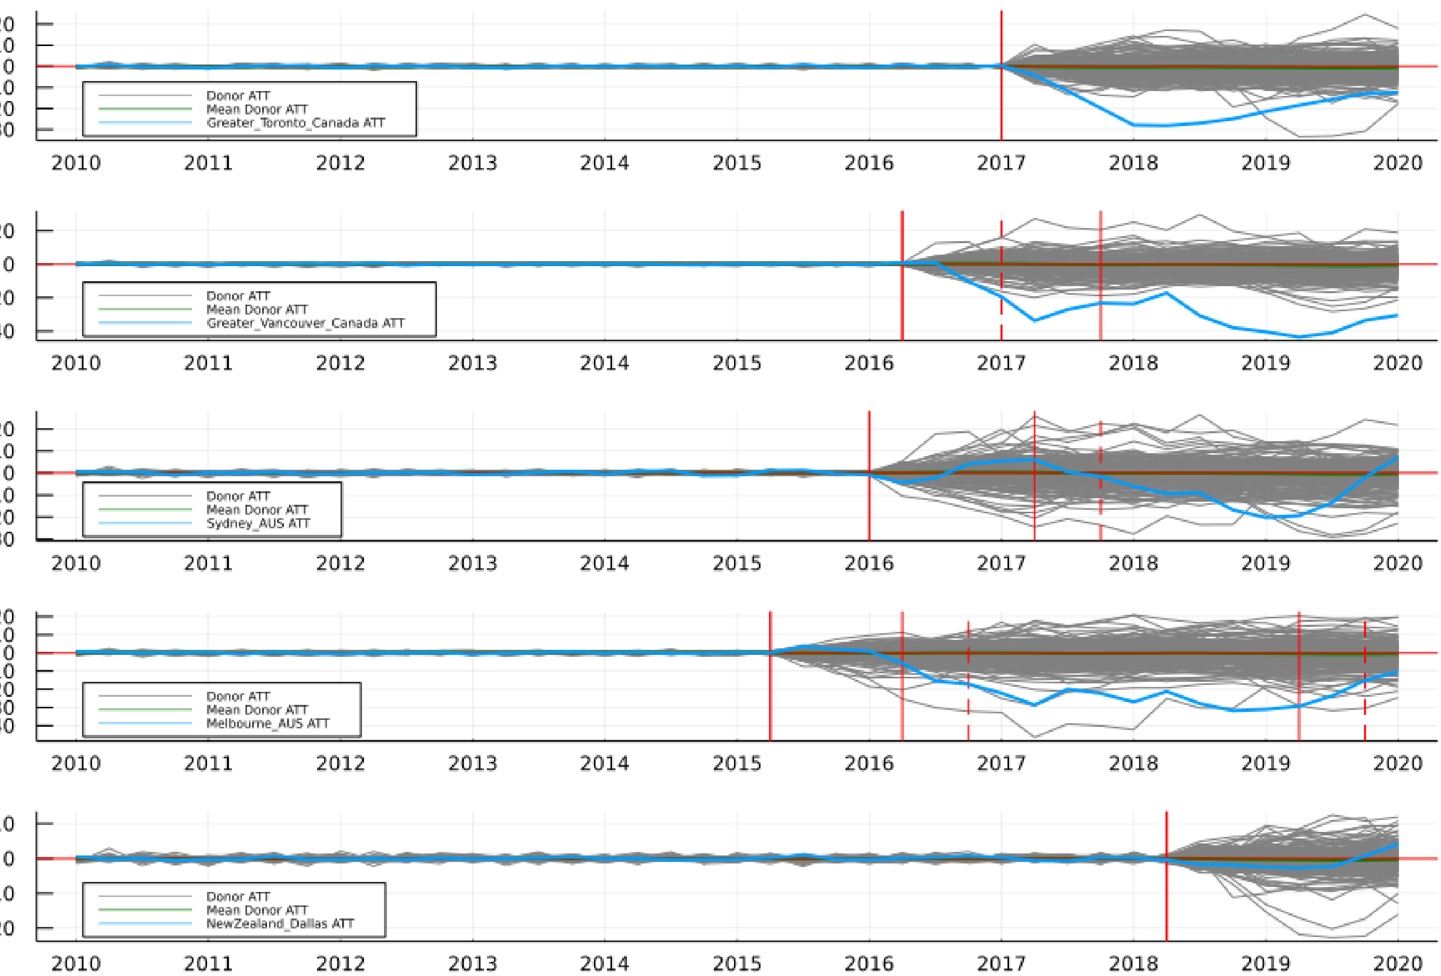

Note.
This figure plots point estimates for ATTs for treated locations (in blue) described in the legend along with ATTs for each untreated location (in gray). ATTs are estimated using the SCM as explained in Section 3.2. Vertical bars, in red, correspond to quarters before Foreign Buyer Taxes were implemented. Dashed vertical bars correspond to vacancy taxes. Data sources can be found in Table 2.

A.5 Heterogeneous Treatment Effect Ontario

Figure 5: Heterogeneous Treatment Effect Ontario

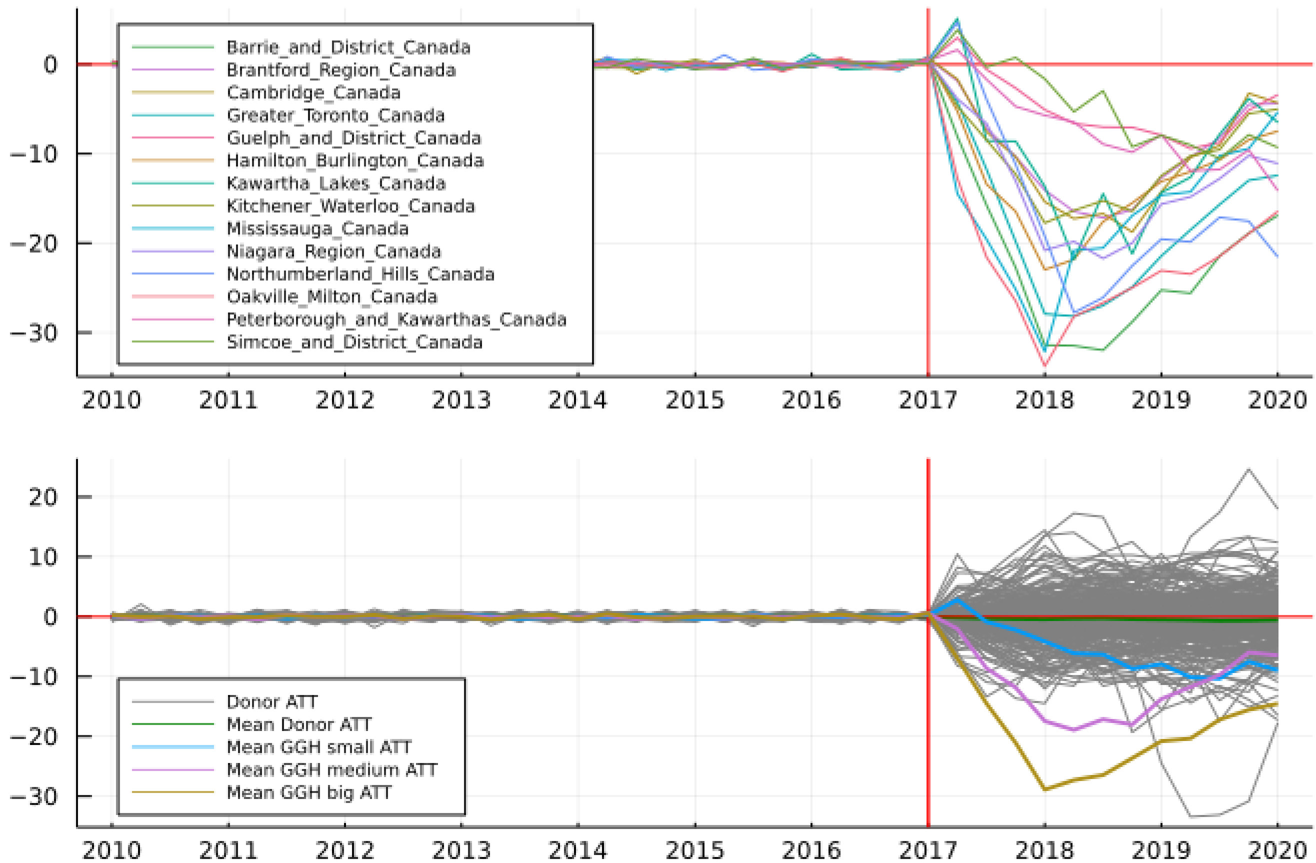

Note.
The top figure plots point estimates for ATTs for treated locations, in Ontario, described in the legend. The bottom figure plots average ATTs across three of Ontario's treated groups as explained in Section 4.2. ATTs for each untreated location are plotted in gray. ATTs are estimated using the SCM as explained in Section 3.2. Vertical bars, in red, correspond to quarters before Foreign Buyer Taxes were implemented. Data sources can be found in Table 2.

A.6 Heterogeneous Treatment Effect British Columbia

Figure 6: Heterogeneous Treatment Effect British Columbia

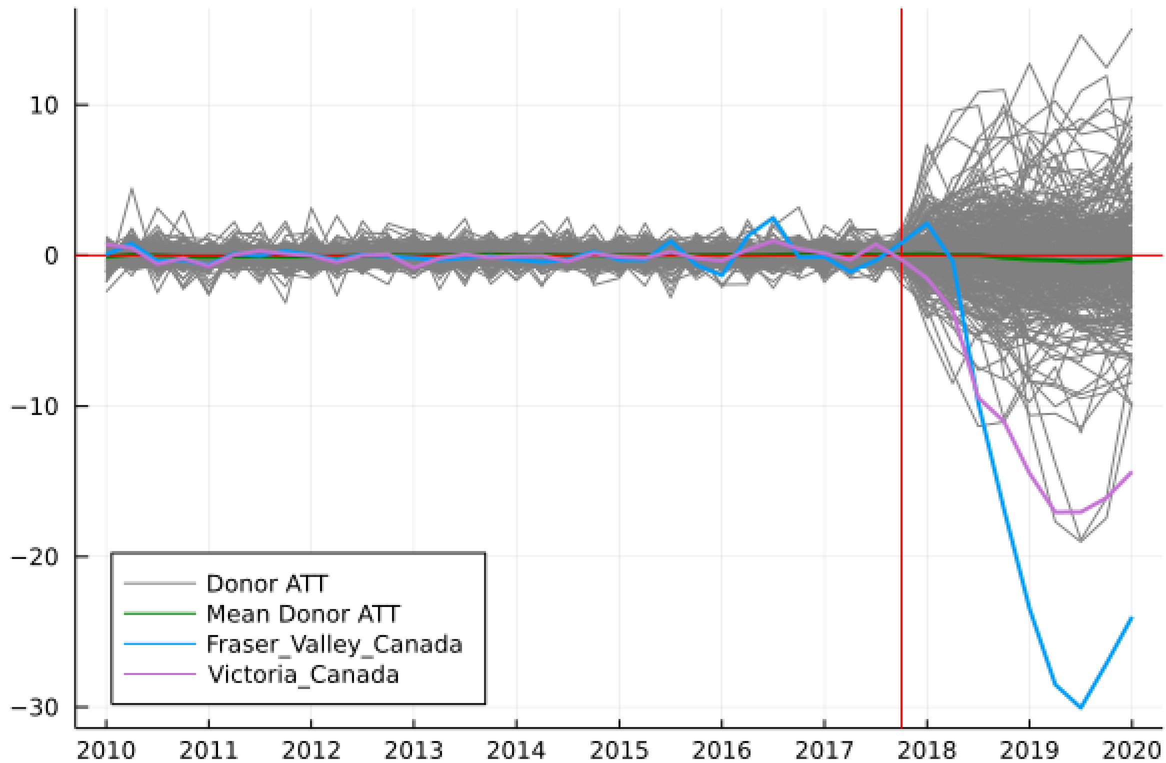

Note.
This figure plots point estimates for ATTs for two treated locations, in British Columbia, described in the legend. ATTs for each untreated location are plotted in gray. ATTs are estimated using the SCM as explained in Section 3.2. Vertical bars, in red, correspond to quarters before Foreign Buyer Taxes were implemented. Data sources can be found in Table 2.

A.7 Alternative Outcomes: Population Growth

Figure 7: ALTERNATIVE OUTCOMES: POPULATION GROWTH

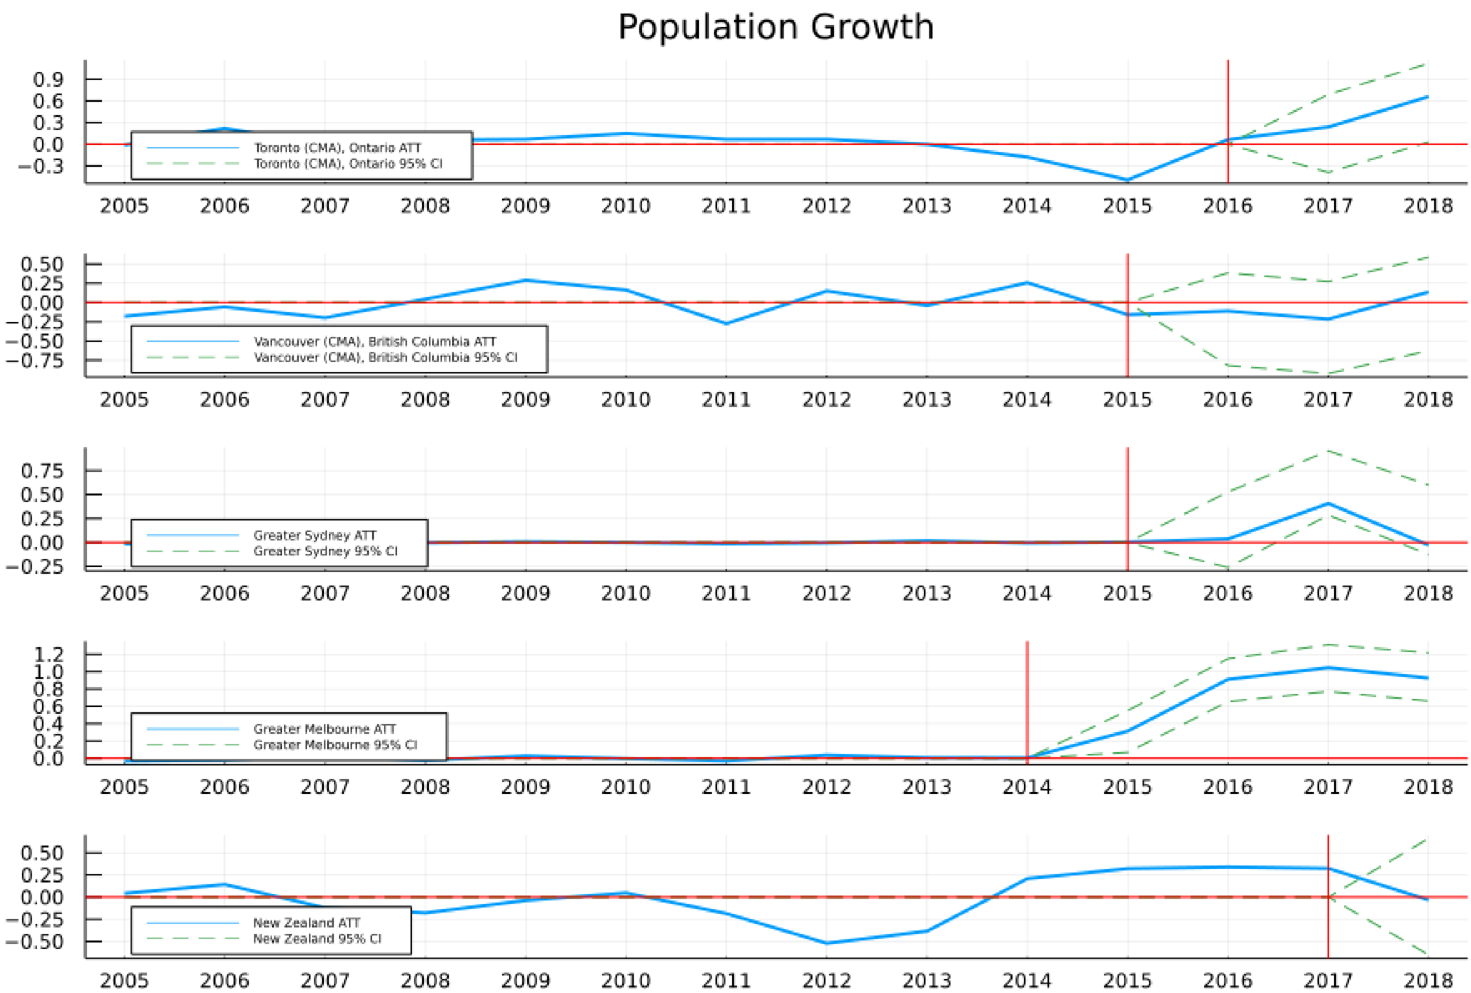

Note.
This figure plots point estimates for ATTs and 95% confidence intervals for treated locations described in the legend. ATTs are estimated using the SCM and confidence intervals are estimated using Jackknife+ as explained in Section 3.2. Vertical bars, in red, correspond to quarters before Foreign Buyer Taxes were implemented. Data sources can be found in Table 2.

A.8 Alternative Outcomes: GDP Growth

Figure 8: Alternative Outcomes: GDP Growth

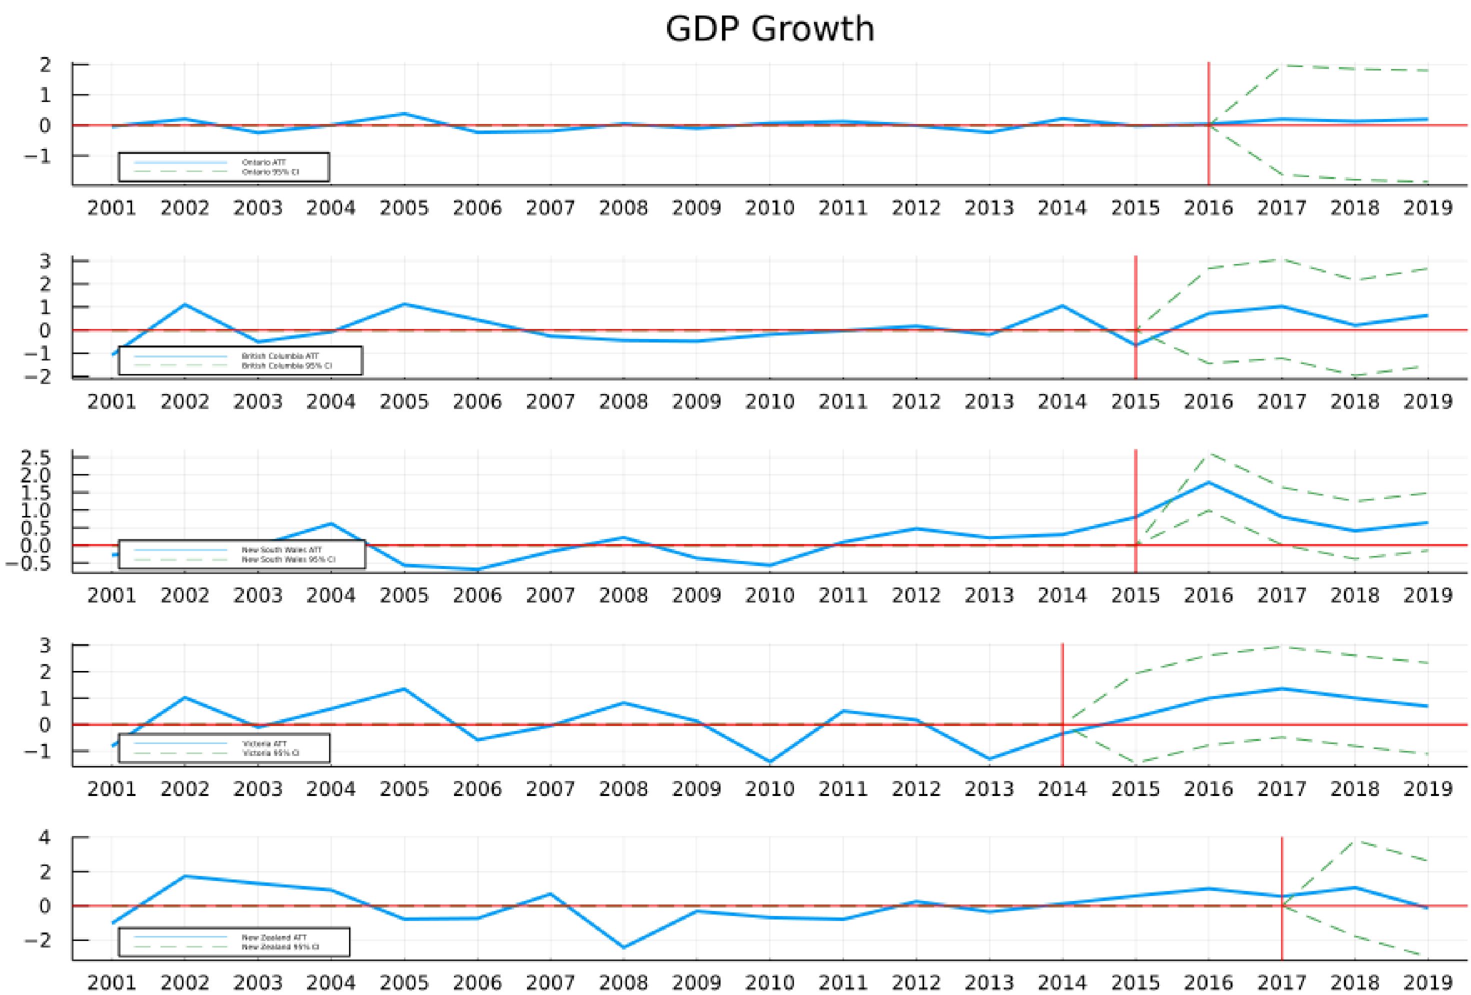

Note.
This figure plots point estimates for ATTs and 95% confidence intervals for treated locations described in the legend. ATTs are estimated using the SCM and confidence intervals are estimated using Jackknife+ as explained in Section 3.2. Vertical bars, in red, correspond to quarters before Foreign Buyer Taxes were implemented. Data sources can be found in Table 2.

A.9 Alternative Outcomes: Unemployment Rate

Figure 9: Alternative Outcomes: Unemployment Rate

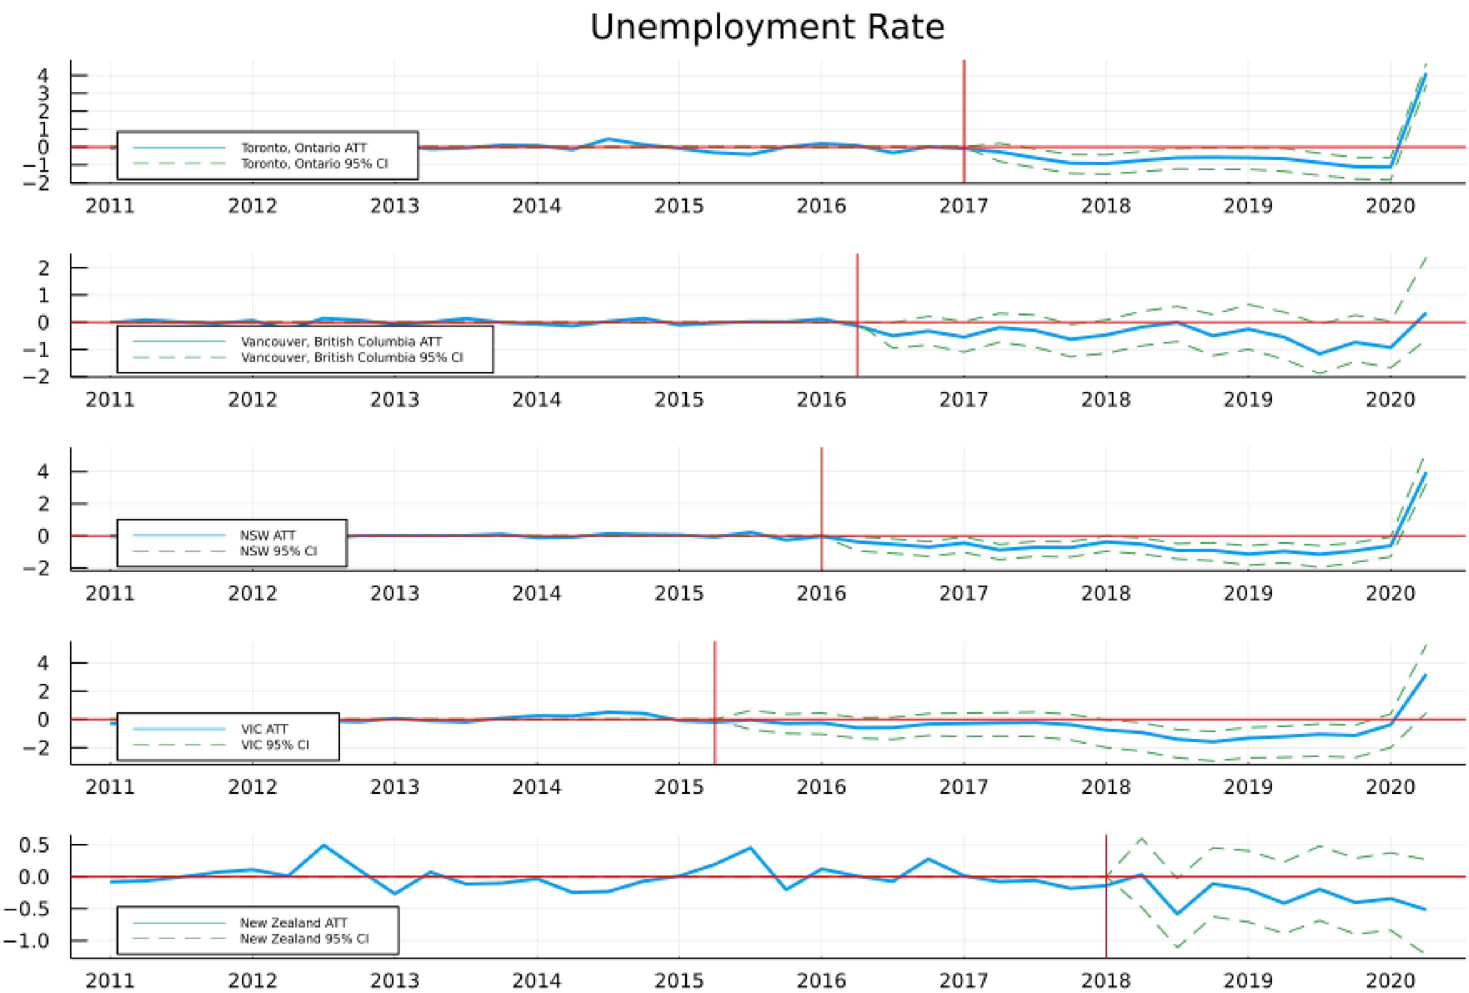

Note.
This figure plots point estimates for ATTs and 95% confidence intervals for treated locations described in the legend. ATTs are estimated using the SCM and confidence intervals are estimated using Jackknife+ as explained in Section 3.2. Vertical bars, in red, correspond to quarters before Foreign Buyer Taxes were implemented. Data sources can be found in Table 2.

B Appendix: Tables

B.1 Law Changes

Table 1: Taxes

Affected Area	Tax	Effective Date
	Canada	
Metro Vancouver Regional District	15% LTT	8/2/16 - 2/20/18
Metro Vancouver + 4 Regional Districts	20% LTT	2/21/18 -
City of Vancouver	1.25% VT	1/1/17 - 12/31/20
	3% VT	1/1/21 -
Greater Golden Horseshoe Region (GGH)	15% LTT	4/21/17 -
	Australia	
Victoria (Melbourne)	3% LTT	7/1/15 - 6/30/16
	7% LTT	7/1/16 - 6/30/19
	8% LTT	7/1/19 -
	1.5% VT	1/1/17 - 1/1/20
	2% VT	1/1/20 -
New South Wales (Sydney)	4% LTT	6/21/16 - 6/30/17
	8% LTT	7/1/17 -
	2% VT	12/31/17 -
	New Zealand	
New Zealand	Ban foreigners from buying existing houses	8/22/18 -

Note. This table lists Foreign Buyer Taxes (FBTs) across Canada, Australia, and New Zealand. All the listed FBTs are in the form of one-time Land Transfer Taxes (LTT). The Vacancy Taxes (VTs) are paid annually by owners of vacant homes.

B.2 Datasets

Table 2: Datasets

Provider	Dataset	Starts	Frequency	Locations
CREA	Canada home price data	2005 -	Monthly	Aggregate + 41 regions
ABS	Australian home price data	2003 -	Quarterly	8 Capital cities
Dallas Fed	International home price data	1975 -	Quarterly	25 countries
BIS	International home price data	2009 -	Quarterly	58 countries
Zillow	US home price data	1996 -	Monthly	200 cities
StatCan	Canada population data	2001 -	Annual	161 CMAs and CAs
ABS	Australia population data	2001 -	Annual	8 states and territories
OECD	International population data	2000 -	Annual	55 countries and regions
StatCan	Canada UR data	2011 -	Monthly	45 locations
ABS	Australia UR data	1978 -	Monthly	8 states and territories
OECD	International UR data	2010 -	Quarterly	42 countries and regions
StatCan	Canada GDP data	1999 -	Annual	13 provinces and territories
ABS	Australia GDP data	1990 -	Annual	8 states and territories
OECD	International GDP data	1960 -	Quarterly	56 countries and regions
StatCan	Immigration Share data	2016	Annual	14 regions

Note. This table lists sources for the different variables used in this paper. Census Metropolitan Areas (CMAs) are defined by Statistics Canada (StatsCan). Australian data are from the Australian Bureau of Statistics (ABS). International data are from the Organisation for Economic Co-operation and Development (OECD) and the Bank for International Settlements (BIS). In some locations (e.g. for states in Australia) local GDP is called Gross State Product (GSP).

B.3 Main Estimates

Greater Toronto

Time	Tax	ATT	CI lower	CI upper
01jun2017	15.0	-4.23	-8.58	-0.62
01sep2017	15.0	-11.75	-15.39	-8.47
01dec2017	15.0	-19.82	-23.5	-16.19
01mar2018	15.0	-27.91	-31.77	-23.93
01jun2018	15.0	-28.17	-31.46	-23.97
01sep2018	15.0	-26.98	-30.36	-22.96
01dec2018	15.0	-24.91	-28.61	-20.88
01mar2019	15.0	-21.43	-25.26	-17.49
01jun2019	15.0	-18.52	-23.35	-15.25
01sep2019	15.0	-15.68	-20.5	-12.76
01dec2019	15.0	-12.98	-17.53	-10.01
01mar2020	15.0	-12.42	-16.33	-8.65

Note.
This table presents point estimates for ATTs and 95% confidence intervals for the treated location described above the table. ATTs are estimated using the SCM and confidence intervals are estimated using Jackknife+ as explained in Section 3.2. The second column presents the tax for the quarter in the first column. Data sources can be found in Table 2.

Greater Vancouver

Time	Tax	ATT	CI lower	CI upper
01sep2016	15.0	0.8	-2.26	6.58
01dec2016	15.0	-10.31	-13.71	-3.48
01mar2017	15.0	-19.47	-23.45	-11.82
01jun2017	15.0	-33.75	-41.28	-25.04
01sep2017	15.0	-27.07	-33.4	-19.04
01dec2017	15.0	-23.27	-31.09	-15.61
01mar2018	20.0	-23.76	-32.07	-14.76
01jun2018	20.0	-17.21	-23.13	-10.28
01sep2018	20.0	-30.78	-37.6	-23.17
01dec2018	20.0	-37.95	-43.65	-28.69
01mar2019	20.0	-40.47	-45.58	-31.49
01jun2019	20.0	-43.52	-49.13	-34.11
01sep2019	20.0	-41.04	-46.31	-32.62
01dec2019	20.0	-33.57	-39.53	-26.51
01mar2020	20.0	-30.48	-37.04	-22.44

Sydney

Time	Tax	ATT	CI lower	CI upper
01jun2016	4.0	-4.34	-14.16	1.2
01sep2016	4.0	-2.26	-15.35	3.61
01dec2016	4.0	4.09	-11.0	9.52
01mar2017	4.0	5.55	-11.19	9.72
01jun2017	4.0	6.1	-13.65	10.69
01sep2017	8.0	0.47	-16.94	5.76
01dec2017	8.0	-1.57	-20.67	5.07
01mar2018	8.0	-6.09	-25.73	1.29
01jun2018	8.0	-9.45	-25.96	-2.1
01sep2018	8.0	-9.05	-26.87	-2.33
01dec2018	8.0	-16.75	-32.04	-10.47
01mar2019	8.0	-20.26	-32.62	-12.78
01jun2019	8.0	-19.41	-30.8	-12.59
01sep2019	8.0	-13.47	-25.51	-6.98
01dec2019	8.0	-2.03	-14.43	4.22
01mar2020	8.0	7.25	-6.48	12.67

Melbourne

Time	Tax	ATT	CI lower	CI upper
01sep2015	3.0	3.49	-1.05	7.12
01dec2015	3.0	2.03	-2.17	5.69
01mar2016	3.0	0.95	-3.45	4.72
01jun2016	3.0	-5.7	-9.97	-0.81
01sep2016	7.0	-15.48	-19.84	-8.67
01dec2016	7.0	-17.19	-22.05	-9.72
01mar2017	7.0	-22.15	-27.2	-13.53
01jun2017	7.0	-28.61	-35.31	-21.38
01sep2017	7.0	-19.93	-26.54	-13.94
01dec2017	7.0	-22.21	-29.8	-16.62
01mar2018	7.0	-26.99	-35.26	-21.19
01jun2018	7.0	-21.19	-28.43	-16.49
01sep2018	7.0	-27.87	-35.72	-22.68
01dec2018	7.0	-31.74	-39.05	-27.02
01mar2019	7.0	-31.03	-37.98	-27.05
01jun2019	7.0	-29.36	-36.11	-25.2
01sep2019	8.0	-23.68	-30.05	-19.39
01dec2019	8.0	-15.09	-21.76	-11.26
01mar2020	8.0	-10.23	-16.96	-5.95

New Zealand

Time	Tax	ATT	CI lower	CI upper
01sep2018	Inf	-1.61	-3.95	0.97
01dec2018	Inf	-1.83	-4.11	0.96
01mar2019	Inf	-2.42	-4.66	0.3
01jun2019	Inf	-2.65	-4.97	0.22
01sep2019	Inf	-2.25	-4.37	0.52
01dec2019	Inf	0.75	-1.58	3.31
01mar2020	Inf	4.21	1.98	6.88

C For Online Publication: Appendix

C.1 ATTs with Confidence Intervals, Fit Two Quarters Pre FBT

Figure 10: ATTs with Confidence Intervals, Fit Two Quarters Pre FBT

Note.
This figure plots point estimates for ATTs and 95% confidence intervals for treated locations described in the legend. ATTs are estimated using the SCM and confidence intervals are estimated using Jackknife+ as explained in Section 3.2. This is similar to Figure 2 except the outcomes are fit up to two quarters before the FBT. Grey vertical bars correspond to two quarters before FBTs were implemented. Red vertical bars correspond to the first quarter the FBTs were implemented. Data sources can be found in Table 2.

C.2 Immigrant Share vs Minimum ATT in GGH

Figure 11: Immigrant Share vs Minimum ATT in GGH

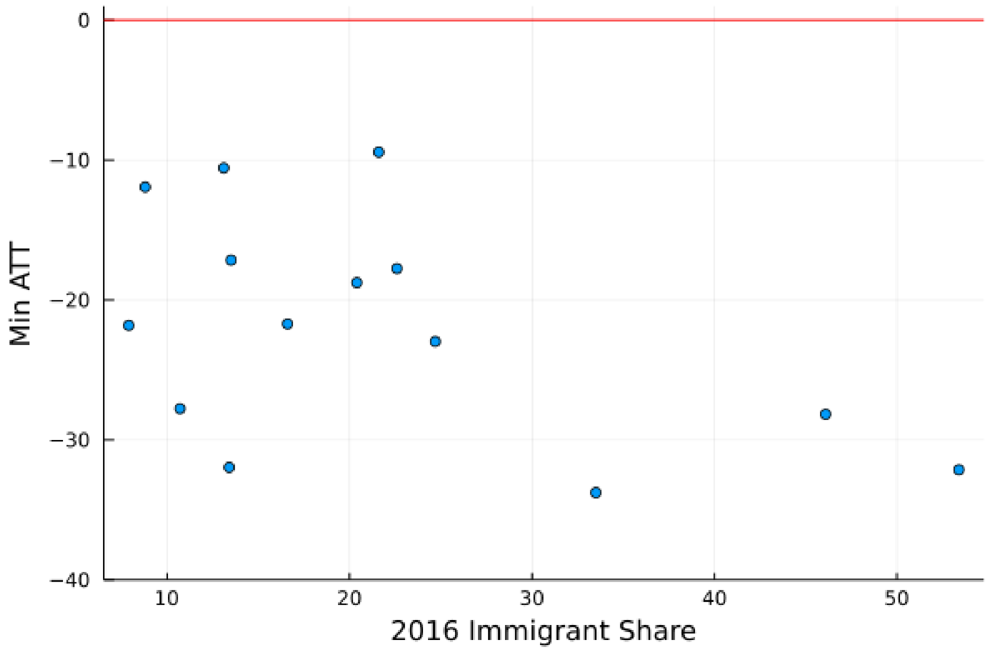

Note.
This figure plots minimum ATTs for 14 locations simultaneously treated in Ontario (Figure 5) against the immigrant share in 2016. The slope from a linear regression is −.3 and statistically significant at the 95% confidence level. Data sources can be found in Table 2.

C.3 Elastic-Net Weights from Main Estimates

Figure 12: Elastic-Net Weights from Main Estimates

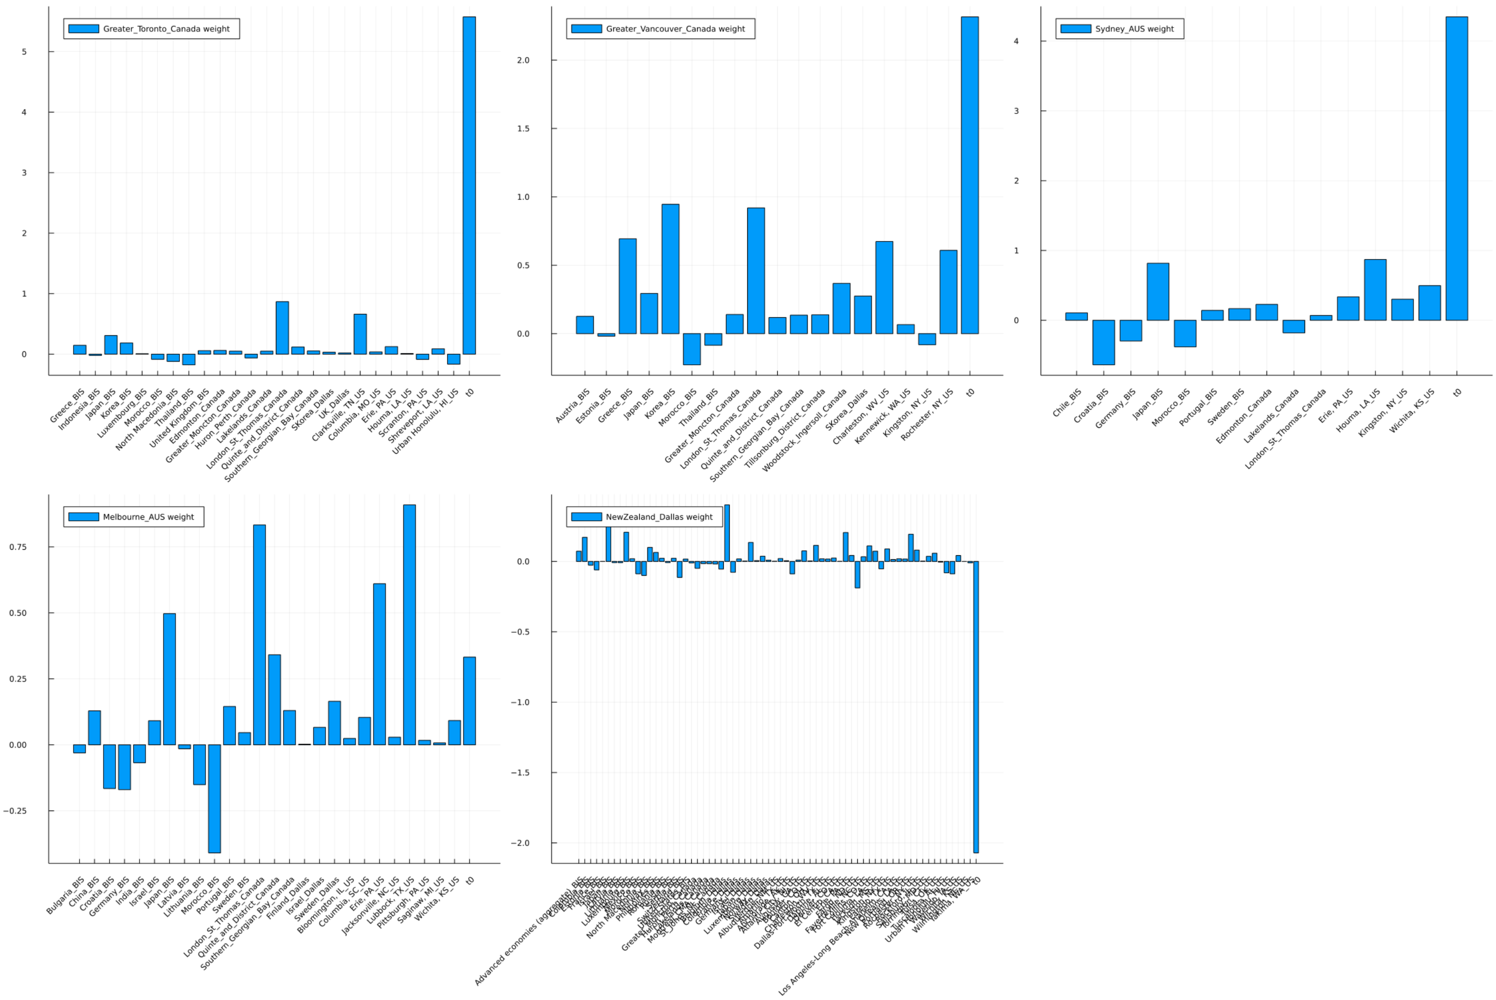

Note.
This figure plots the (non-zero) parameters fit by elastic-net, after cross-validation, for treated locations described in the legend. The method explained in Section 3.2. While there are 348 possible predictors (345 donor locations, an intercept denoted t0, a linear trend t1, a quadratic trend t2), a much smaller number is selected by elastic-net. Data sources can be found in Table 2.

C.4 Difference-in-Differences Estimates

Figure 13: Difference-in-Differences Estimates

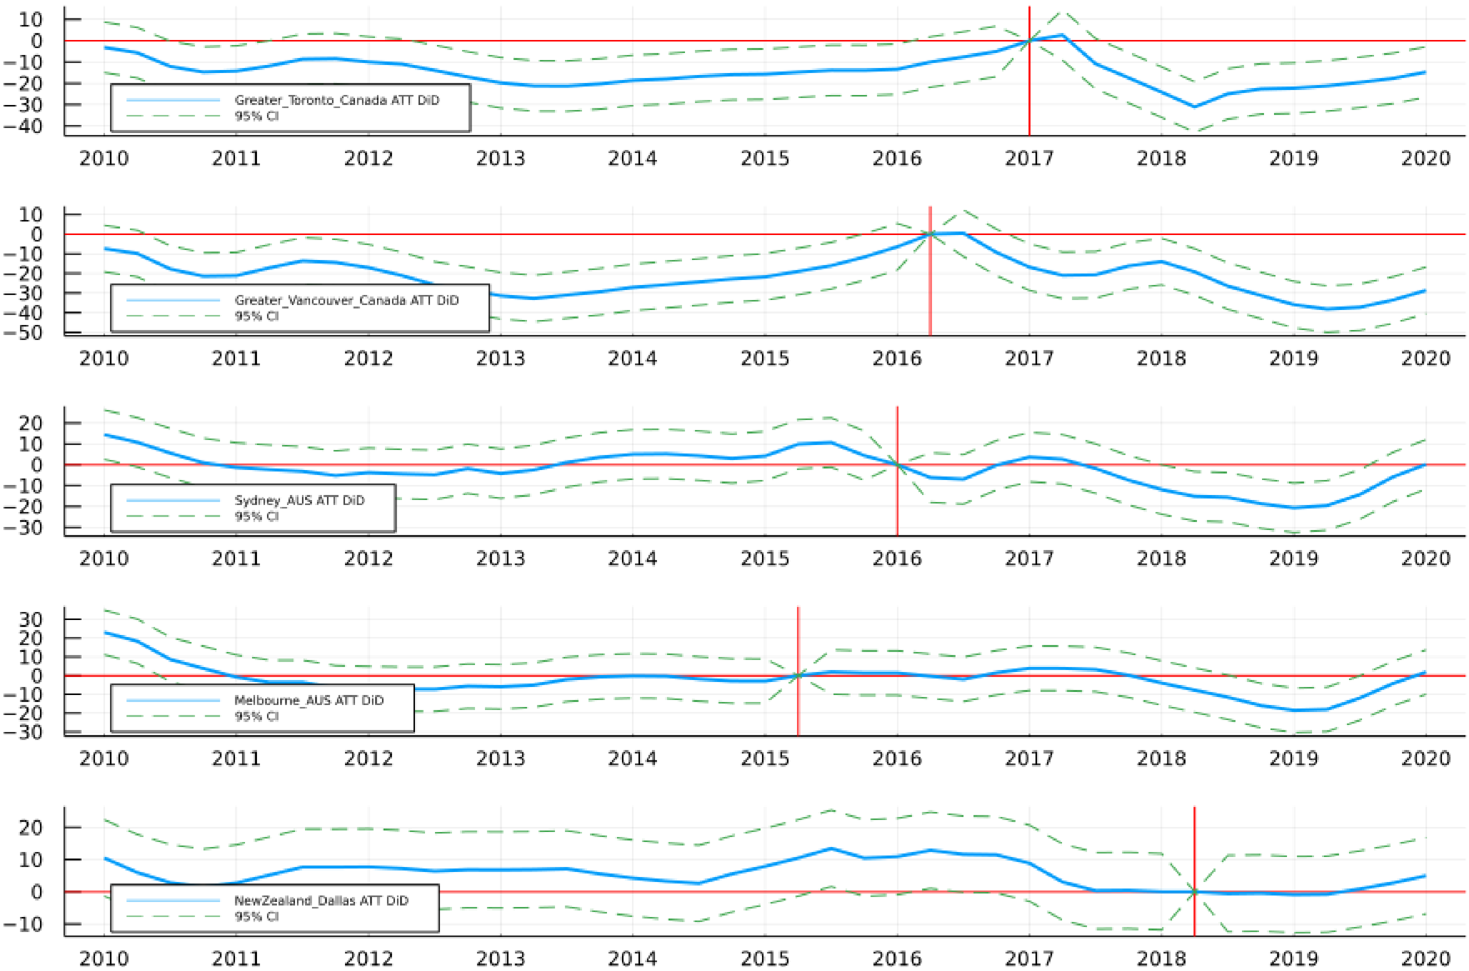

Note.
This figure plots point estimates for ATTs and 95% confidence intervals for treated locations described in the legend. The ATTs are estimated using the traditional two way fixed effects (difference-in-differences) estimator and can be compared with the SCM estimates in Figure 2. Vertical bars, in red, correspond to quarters before Foreign Buyer Taxes were implemented. Data sources can be found in Table 2.