Zoning Reform in New Zealand:

A Staggered Estimation Framework

James Hancock

Supervised by Dr Ryan Greenaway-McGrevy

11/11/24

A *dissertation* presented in

partial fulfilment of the requirements for the Degree of

Bachelor of Commerce (Honours)

in Economics

of The University of Auckland

Abstract

This research examines the effects of zoning reform on housing construction across several regions in New Zealand, using a staggered synthetic control method to estimate impacts in Auckland, Lower Hutt, and Christchurch. The analysis applies a novel multiple treated unit framework, leveraging temporal variations in reform implementation to create synthetic control units for each city. Findings indicate significant increases in housing consents following reform, peaking at an additional 7.5 consents per thousand residents in the second year post-reform, with effects diminishing over time. Notably, Christchurch's reform, focused on earthquake recovery, demonstrated distinct impacts compared to Auckland and Lower Hutt, where reform aimed to address urban density constraints. These results are robust to a range of checks, including placebo tests and alternative donor pools. Conservative inference testing finds the majority of treatment effects at an aggregate level to be statistically significant, with mixed results at the individual level. The study underscores the effectiveness of zoning reform in boosting housing supply but also highlights the complexities in isolating individual policy impacts within the reform suite. This framework offers valuable insights for policymakers aiming to address housing shortages through land use regulation adjustments, emphasizing the need for ongoing research to refine inference methods and manage treatment heterogeneity in staggered adoption analyses.

Contents

1 Introduction

Few other sectors command the same political, economic, media, social and cultural attention as housing. Supply and demand for shelter jointly determine and are determined by where people want to live, work, meet, socialise and be. Owning your own home is still viewed as a financial and social milestone, even as housing affordability has spiralled downwards across the world (Wetzstein 2017).

It is of no surprise, then, that significant attention has been placed on ways to alleviate these housing pressures. Micro-economic theory suggests that, ceteris paribus, the increase in prices must come from a net increase in demand or a decrease in supply. Proposed solutions to housing shortages have typically been split down these lines, with emphasis placed on the role of immigration, tax relief, state housing - and land use regulations.

Land use regulations (LURs) are the legal frameworks and policies that control how land is developed and used. These regulations may include zoning laws and environmental or heritage protections. They can also direct the allocation of land for various purposes, including residential, commercial, industrial, and agricultural uses.

New Zealand has not avoided the pitfalls of stringent LURs. The New Zealand Productivity Commission (2015) identified restrictive land use rules as a key contributor to the housing crisis, with supply unresponsive to changes in demand. This characteristic holds particularly true in New Zealand's urban environments including Auckland, Wellington and Christchurch (Ministry of Housing and Urban Development 2024). In response, successive central governments adopted varying degrees of pro-supply housing policy aimed at loosening these restrictions to allow new private and public developments.

These policies were not implemented uniformly, however. Only five territorial authorities (TAs) - the local authorities that govern planning rules in their respective regions - have adopted zoning reform, each in different clusters. Christchurch City, Waimakariri District and Selywn District instituted reform following the 2010 and 2011 earthquakes; Auckland City up-zoned from 2013 onwards, and Lower Hutt City began reform in 2018.

While researchers have estimated the effect of these separate cases of zoning reform on housing construction (Greenaway-McGrevy 2023, Maltman & Greenaway-McGrevy 2024), there has been no attempt to quantify the impact within a multiple treated unit framework.

In this dissertation I use the staggered adoption synthetic control method (SCM) proposed by Cattaneo et al. (2023) to estimate the cumulative impact of zoning reform on housing construction in areas that enacted reform. I find significant statistical and economic impacts for the first five years, peaking at an additional 7.5 consents per thousand residents in the second year following reform relative to the synthetic control estimate. Year six post adoption is both the last year available to analyse and the only one without a statistically significant impact. Individual territorial authority findings show a range of treatment effects in terms of size and magnitude, with a clear division between the Christchurch region and Auckland/Lower Hutt. The synthetic control estimates are robust to a range of model specifications, but inference can differ notably. This finding suggests that further research on the nascent area of inference in staggered adoption frameworks is necessary. I also establish a range of potential and actual limitations of the framework for future analysis, including treatment heterogeneity, donor pool selection, and potential spillover effects.

Proposed first by Abadie & Gardeazabal (2003), the synthetic control method estimates the effect of an intervention by constructing a weighted combination of untreated units to create a "synthetic" version of the treated unit, representing what would have happened without the intervention. This synthetic control acts as a comparison group, and the difference in outcomes between the treated unit and the synthetic control after the intervention provides an estimate of the treatment effect. In this empirical application, the SCM creates synthetic versions of treated units (Auckland, Lower Hutt, and the Christchurch region) to estimate the impact of zoning reform on housing consents relative to if no zoning reform had occurred.

Auckland's zoning reform has been studied extensively as one of the few empirical examples of widespread zoning reform (Greenaway-McGrevy 2023, Cheung et al. 2024). Greenaway-McGrevy (2023) creates a synthetic Auckland using information from other 'Functional Urban Areas' - or commuting zones, of which New Zealand has 53 - to estimate the impact of zoning reform relative to what would have happened without zoning reform. The author finds that reform led to an 82 per cent increase in consents during the 2016-2022 period.

Maltman & Greenaway-McGrevy (2024) is the only current academic attempt to quantify the causal impact of this suite of zoning reforms on housing construction in Lower Hutt. Following the same approach as Greenaway-McGrevy (2023), the authors form a synthetic control of Lower Hutt by matching non-treated territorial authorities on common housing market outcomes. Unlike its predecessor, it matches on two sets of housing outcomes, and implements a much larger 'battery' of robustness checks. They also correct for the impact of spillover effects, or development diverted away from neighbouring areas towards Lower Hutt due to reform. Accounting for these effects, the analysis indicates that the reforms resulted in a 12 to 17 per cent growth in housing consents in Lower Hutt during the 2018-2023 period.

There has been no effort to estimate empirically the effect of zoning reform on the Christchurch region post-2011. The monitoring report produced by the Christchurch Earthquake Recovery Authority, (Canterbury Earthquake Recovery Authority (CERA) 2015), finds that from February 2011 to December 2014, 16,800 new houses received building permits in the greater Christchurch region, and that Selwyn and Waimakariri had a 193 and 220 per cent increase in new residential units consented compared to the decade years before the earthquakes. CERA does not explore what portion of this increase can be attributed to zoning reform, rather than the continuation of pre-earthquake consenting patterns.

My study contributes to existing literature in three ways. The dissertation develops the first multiple-unit framework for the continued analysis of zoning reform in New Zealand that allows for the staggered adoption of reform, albeit with a range of limitations. Secondly, this analysis adds another point of evidence - the experience of Christchurch and surrounding areas - to the effectiveness of zoning reform as a method of addressing housing shortages. Finally, it explores a new range of inference methods and robustness checks to combine and largely replicate results from existing literature, namely Greenaway-McGrevy (2023) and Maltman & Greenaway-McGrevy (2024).

This thesis is structured in the following manner. Section One has established the motivation for zoning reform and the intentions of this dissertation, Section Two examines the context and implementation of the zoning reforms introduced in Auckland and Lower Hutt, as well as those of the Christchurch region. I then outline my method in Section Three to estimate the impact of this reform on housing construction. I give a brief overview of data in Section Four. Section Five presents the main findings of the two implemented models. In Section Six I run inference on my results, before conducting a series of robustness checks. In Section Seven I review limitations of my modelling approach, and discuss avenues for future research. Section Eight concludes the dissertation.

2 Institutional Background

This section reviews the institutional background for zoning reform in New Zealand, as well as the key legislation that enabled reform in each of the examined areas. Figure 1 shows a map of New Zealand, with areas that reformed outlined in red.

2.1 New Zealand context

West (2024) outlines the evolution of supply-side reform to the housing sector in New Zealand. Three key pieces of legislation have guided the national approach to land use regulation reform: the *NPS-UDC* (2016), the *NPS-UD* (2020), and the *MDRS* (2021).

The 2016 *National Policy Statement on Urban Development Capacity* (NPS-UDC) required local government to undertake regular assessment of their capacity for urban growth and

Figure 1: Map of New Zealand highlighting territorial authorities that adopted zoning reform. Author's edit of Stats NZ (2023).

business development. Councils then had to ensure "that at any one time there is sufficient housing... land development capacity" in the short (next three years), medium (between three and ten years), and long term (between ten and thirty years) (Ministry for the Environment and Ministry of Business, Innovation and Employment 2016, p.11) . The *NPS-UDC* does not restrict where this development capacity might come from, stating it might come from “intensifying existing urban areas, [or] by releasing land in greenfield areas” (Ministry for the Environment and Ministry of Business, Innovation and Employment 2016, p.3).

The 2020 *National Policy Statement on Urban Development* (NPS-UD), instructed councils to relax LURs by enabling density around mass rapid transit lines and in walkable city centres, as well as to remove minimum parking requirements. It was later followed by the 2021 *Medium Density Residential Standards* (MDRS), which mandated that local authorities permit up to three homes, each up to three storeys high, on any section of land. Councils have delayed the implementation of both policies[1], with only Upper Hutt, Lower Hutt and Wellington operationalising the *MDRS*.

These national pieces of legislation required enactment by local government, with varying levels of resistance (West 2024). The varying council responses form the basis of the framework identified in Section 3.1, which allows for estimation of zoning reform impacts in cases of staggered adoption.

I now briefly outline the zoning history and legislative changes of each council that have enacted reforms.

2.2 Auckland City

Auckland is New Zealand’s most populous region, with its population estimated at 1.798 million as of June 2024 (Stats NZ 2024*a*). Despite its population count, it is also the second smallest by land area at 4,941 km^2 (Stats NZ 2020).

Before 2010 Auckland’s local government comprised seven separate city and district councils, each with their own approach to the regulation of land use. The seven councils were merged in November 2010 into a new unitary authority, or ‘super city’, which necessitated a unified district plan. The main aim of the plan was to use the merger as an opportunity to create a

[1]The *MDRS* is being reworked by the current government to allow councils to opt-out provided they zone for 30 years of growth in their next district plan review (Brunskill 2023).

set of consistent planning rules directed towards sustainable development (Duguid & Chan 2014). An initial *Proposed Auckland Unitary Plan* (PAUP) signalled the first stage of land use reforms. Its main emphasis was on restricting sprawl and intensifying around growth centres and high-frequency transit lines (Auckland Council 2013).

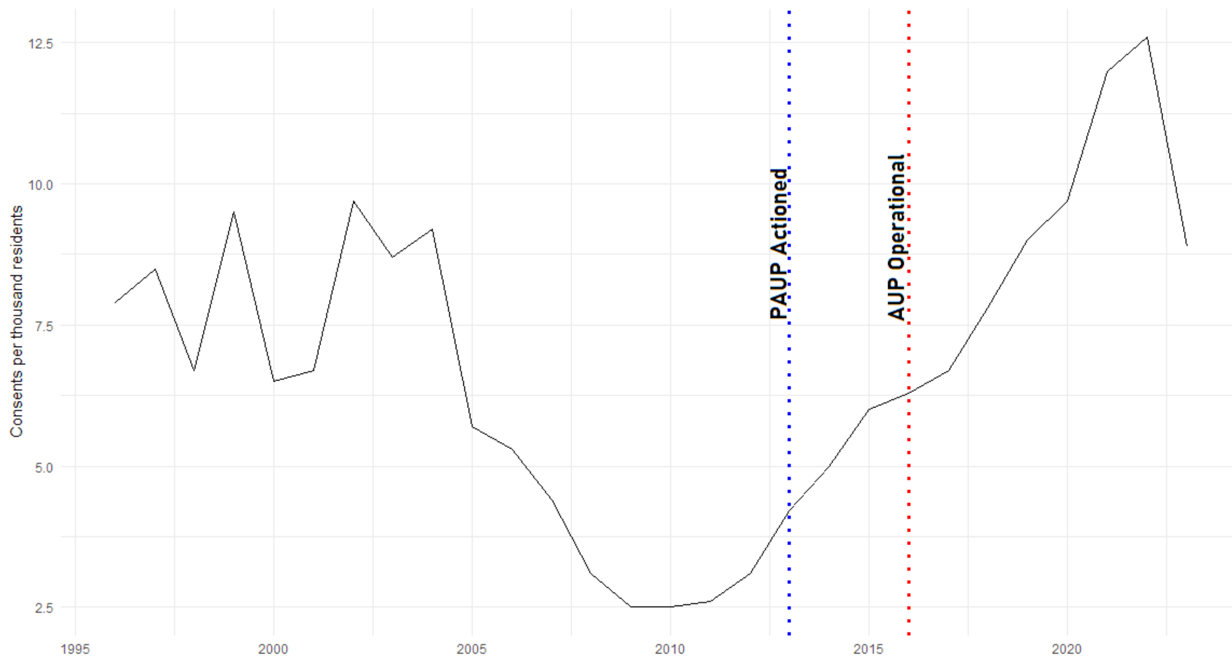

Figure 2: Auckland consents per thousand residents. Source: Stats NZ (2024*b*)

Reform began in September 2013 with the inclusionary zoning programme entitled 'Special Housing Areas'. Developers who allocated ten percent of their developments to affordable housing received expedited permitting processes - and were assessed under the looser *PAUP* regulations. Buildings in the new *Mixed Housing Urban* zone could be built up to 12 metres, with Terraced Housing Apartments up to 16 metres (Auckland Council 2013, Greenaway-McGrevy & Jones 2023). These looser LURs remained consistent in the final *Auckland Unitary Plan* (AUP). The final date of introduction for the *AUP* was November 2016, with nearly $\sim$ 75% of Auckland's residential land area up-zoned by eliminating single-family zoning (Greenaway-McGrevy & Jones 2023).

Figure 2 depicts annual consents per thousand residents in Auckland since 1996. After averaging around 7.5 consents per thousand residents from 1996 to the mid-2000s, there is a sharp decline before holding constant at 2.5 consents until the early 2010s. There appears to be a considerable increase with the introduction of the *PAUP* and *AUP*.

2.3 Lower Hutt City

Forming part of the Wellington region at the base of the North Island, Lower Hutt is New Zealand's sixth largest city. In June 2024 its estimated population was 114,500 (Stats NZ 2024*a*). It is bordered by Wellington City, Upper Hutt, Porirua and South Wairarapa, and is governed at a local government level by the Hutt City Council territorial authority. Approximately 10% of Lower Hutt's commuter population travels to Wellington City for work (Greater Wellington Regional Council 2020). Like the rest of the Wellington region, Lower Hutt has experienced significant increases in house prices in recent years (Scoop 2016, Hutt City Council 2017).

Before reform, Lower Hutt had relatively low housing density. According to Hutt City Council (2017), 83% of Lower Hutt's housing consisted of detached homes up to two storeys within the *General Residential Activity Area*. As in Auckland, this housing faced additional non-floor area restrictions, including minimum parking requirements, setbacks, and outdoor space mandates. Lower Hutt is also geographically limited, bounded by three river valleys: the Hutt river/Te Awa Kairangi, Wainuiomata and Orongorongo. This geographically severely restricted opportunities for new greenfield developments. A 2018-2021 investigation into greenfield and infill housing in Lower Hutt found that greenfield housing had a feasible development capacity of just 871 dwellings, mainly in Wainuiomata and Western Hills, 8 km and 4 km respectively from the Hutt City central business district (CBD) (Hutt City Council 2022*b*). There were no remaining greenfield sections in Central and South Lower Hutt, the areas of greatest demand. To meet the projected growth in population, and demands of the *NPS-UDC*, Hutt City Council estimated it needed between 6,105 and 11,256 new dwellings for the period 2017 to 2047 (Wellington Regional Leadership Committee 2023). As a result, from the late 2010s, Hutt City Council instituted a range of zoning reforms through revisions in its district plan.

The first phase was Plan Change 39 ('Transport'), which amended the residential parking standards by lowering the general requirement from two parking spaces per dwelling to one. Plan Change 39 was later superseded by the removal of minimum car parking spaces as required under the *NPS-UD*, which had immediate effect (MRCagney 2022).

The second phase of zoning reform was District Plan Change 43 ('Residential and Suburban Mixed Use'). Partially operational in April 2020, this reform put in place the *NPS-UD*, namely the greater intensification of the wider city. Hutt City Council re-zoned eight areas of city land into two zones: the *Medium Density Residential Activity Area* and the *Suburban Mixed Use Activity Area*. The former allowed building up to 10 metres, and the latter up to 12 metres. Planning controls in the *General Residential Activity Area* were also relaxed to accommodate medium density housing development, including terraced and clustered dwellings up to 8 metres in height, on sites greater than 1400 square metres.

District Plan Change 56 ('Enabling Intensification in Residential and Commercial Areas'), a more significant set of zoning reforms, closely followed in September 2023[2]. Its main aim was to incorporate the *Medium Density Residential Standards*, and to give effect to the walkable catchment development section of the *NPS-UD* (Hutt City Council 2022*a*). Plan Change 56 streamlined the four existing zones into just two areas: *High Density Residential Activity Area* and *Medium Density Residential Activity Area*. Development up to six storeys was now possible within 1200-metre walkable catchments of city centres and 800 metres of railway stations. Concurrently, the varied height limits of 8, 12, and 20 metres for commercial and community health facilities were unified into a single 22-metre height limit across the *General Business Activity Area* and *Community Health Activity Area*.

[2]Partly operational from 18 August 2022.

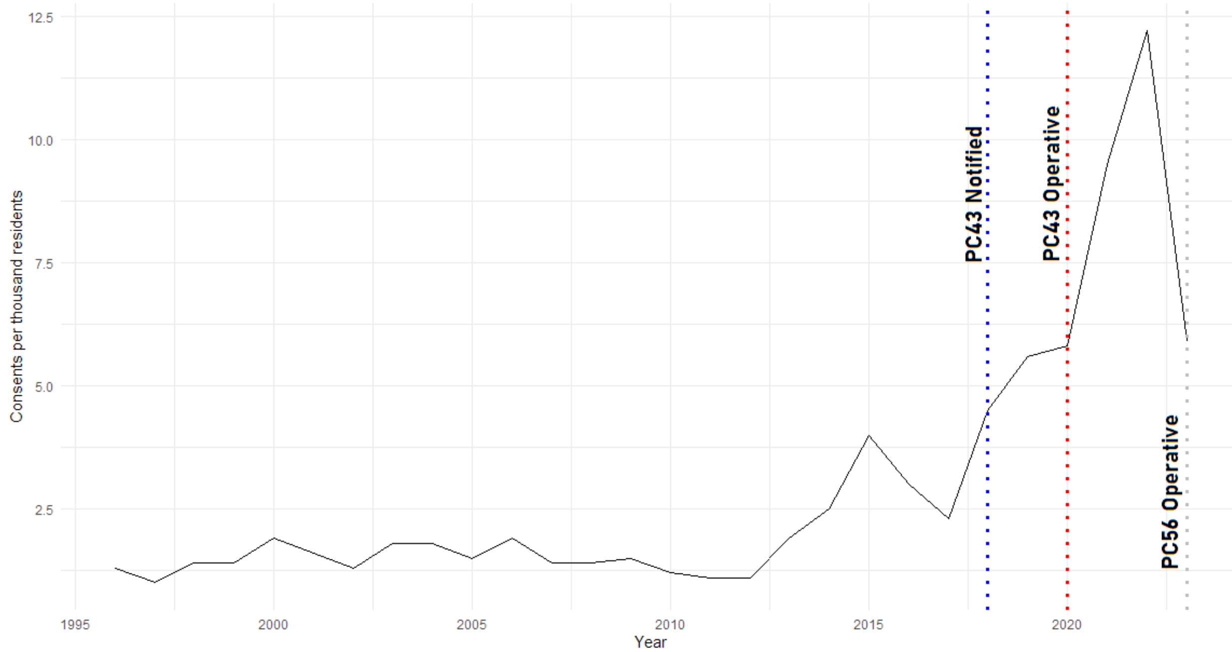

Figure 3: Lower Hutt consents per thousand residents. Source: Stats NZ (2024*b*)

Figure 3 exhibits annual consents per thousand residents in Lower Hutt since 1996. There is far less variability than in the Auckland consent plot, with consents per thousand residents staying constant at approximately 1.5 from 1996 through to the early 2010s. There is a short increase and dip before a steep rise with the notification of Plan Change 43 to a peak of nearly 12.5 consents per thousand residents.

2.4 Christchurch region

Christchurch is New Zealand's second largest city, with an estimated population of 403,000 (Stats NZ 2024*a*). Geographically it is far less dense and flatter than both Auckland and Lower Hutt, with the Port Hills being the only major set change of elevation in the city.

It is governed by the Christchurch City Council territorial authority, and closely neighbours both Waimakariri and Selwyn Districts. Together the three authorities oversee the area stretching from the Rakaia river in the west, to Banks Peninsula in the east, including much of the Canterbury plains.

Selwyn and Waimakariri are both predominantly rural, home to a combined 155,000 people in 8,600 km^2 total area (Stats NZ 2024*a*). Approximately 36.9% and 27.9% of the respective populations live in the district's largest towns: Rolleston (pop. 31,600) and Rangiora (pop. 19,400).

In September 2010 and February 2011, the Christchurch, Waimakariri, and Selwyn Districts were hit by two catastrophic earthquakes. Darfield, 40 km west of Christchurch, was the epicentre of the September quake. The February earthquake was smaller in magnitude but only 5 km underground and 6.7 km south east of the main CBD (Survey 2011). Christchurch's housing supply was severely damaged in terms of both quantity and quality. Of the 184,000 total dwellings, 167,000 had earthquake-related damage claims, with more than 25,000 deemed uninhabitable or significantly impaired (Brunsdon 2019). House and rent prices surged across the city, even as the area experienced a loss in households (LaCour-Little & Staer 2016, Colbert et al. 2022). A prolonged insurance claim and repair process further worsened the situation (Robinson 2013).

Following the severe change in housing and infrastructure stock post-earthquake, there was unanimous agreement between local and central government on the need for a change to land use planning in the area (Theelen 2014). Prior LURs were unclear and numerous, with 103 different zones - 40 types for residential alone - present in the two operative district plans: the *Christchurch City Plan* and the *Banks Peninsula District Plan*. A review in January 2011 of the district plan found it was of "unknown efficiency in controlling the growth of the city" but had resulted in "increased density of development" (Theelen 2014, p.44). The plan's main approach for zoning for projected demand was to rezone several small land parcels along the perimeter of the existing urban area to support urban development; a far cry from a planning system that prioritised intensification (Theelen 2014, p.45). Christchurch was also quickly expanding, with the amount of undeveloped vacant land decreasing from 1272 hectares to 801 hectares, or 37%, between 1999 and 2009[3].

[3] Theelen (2014) cites a Christchurch City Council document called the *Urban Growth Indicator Sheets*

Broadly speaking, Christchurch's zoning approach pre-earthquakes allowed some density in the centre, but decreased density quickly with distance from the CBD. Figure 4 shows the location of new residential units consented since 1991 by area unit; the consenting drop off just a short distance from the CBD is clear.

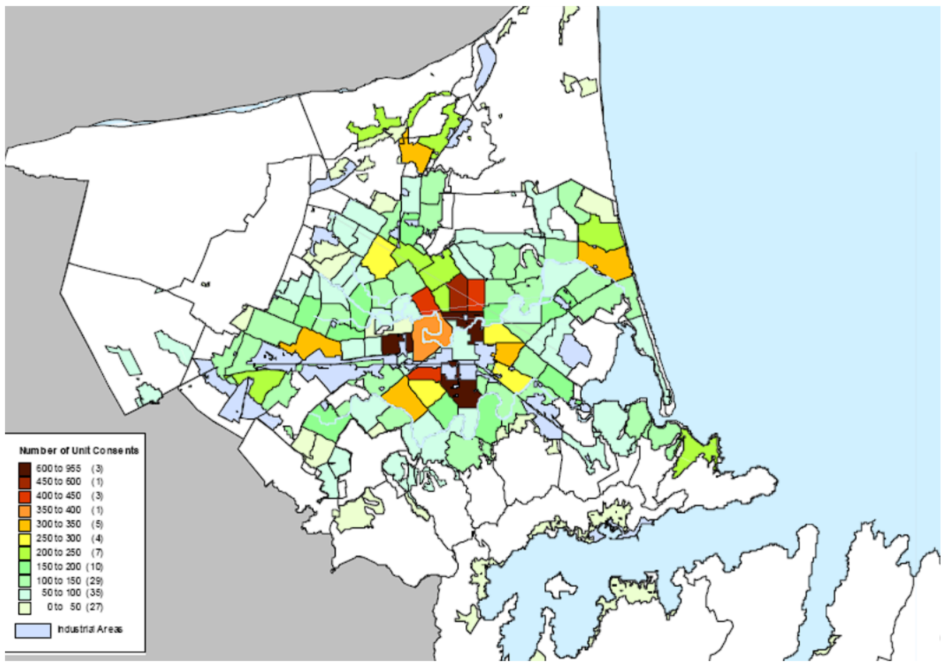

Figure 4: Christchurch, location of new residential units since 1991 by area unit. Source: Theelen (2014)

Three key pieces of legislation were passed following the earthquakes as part of the recovery process. These were the *Land Use Recovery Plan* (LURP), the *Central City Recovery Plan*, and the *Replacement District Plan* (RDP) [4]

The *LURP* modified the existing *Canterbury Regional Policy Statement 2013* to give directions on how land use and associated regulations should change in the wake of the earthquakes. Alongside the *Replacement District Plan 2014*, it was the main earthquake recovery document. It focused on the "consolidation and intensification of urban areas, [avoiding] unplanned expansion." (Environment Canterbury Regional Council 2013, p. 50). This in-

from September 2009, which is no longer available. This statistic aligns with Manaaki Whenua Landcare Research (2021).

[4] The first two of these were enabled by the general *Canterbury Earthquake Recovery Act 2011* (CERA).

tensification was to take place in a range of areas but mainly the Christchurch CBD and surrounding neighbourhoods, with Selwyn and Waimakariri experiencing less intensification. Three main residential zones were created: *Residential Suburban* (single or two storey detached or semi-detached houses), *Residential Suburban Density Transition* (mixed density, one to three storeys), and *Residential Medium Density* (two to three storey buildings including semi-detached and terraced housing, low-rise apartments) (Panel 2015). Given the mix of focus on infill, brownfield and greenfield land, the term 'out-zoning' (as opposed to the clearer 'up-zoning' approach of Auckland and Lower Hutt) has been applied to describe the earthquake response (West 2024).

Although the *LURP* was gazetted in December 2013, there is no clear single date for the beginning of zoning reform in Christchurch. *CERA* gave the Minister for Canterbury Earthquake Recovery (Rt. Hon. Gerry Brownlee) the power to pause, modify, or withdraw any document related to resource management in the affected area. This discretionary power included the re-zoning of greenfield land for residential purposes in November 2011 (Theelen 2014). This land totalled 336 hectares, and was quickly followed in May 2012 by land in north west Belfast (93 hectares) (Christchurch City Council 2015). To account for the range of potential policy start times, I set treatment time equal to the first full year following the initial zoning reform: 2012.

In addition, Christchurch has since removed minimum parking requirements as required under the *NPS-UD* in February 2022, and is reviewing feedback on their proposed implementation of the *MDRS*, Plan Change 14 ('Housing and Business Choice') (Christchurch City Council 2024).

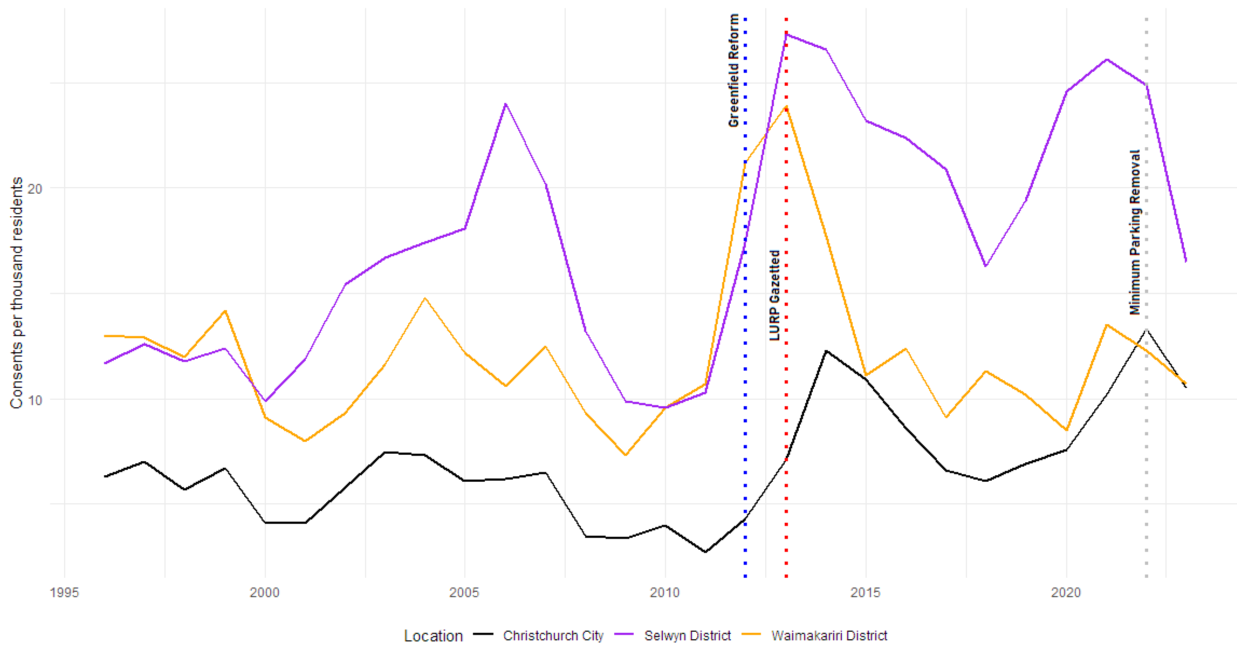

Figure 5: Christchurch region consents per thousand residents. Source: Stats NZ (2024*b*)

Figure 5 displays annual consents per thousand residents in the three Christchurch region territorial authorities since 1996. Although there is a clear level difference between the three areas, they all have similar consenting patterns. Selwyn appears to have the greatest variability. There is a sharp decrease in consents per thousand residents after the global financial crisis with an equally large increase from 2012 onwards with the gazetting of the $LURP$ and onset of greenfield reform. This spike is temporary but somewhat persistent, with average consents for Selwyn and Christchurch - but not Waimakariri - still higher than the pre-earthquake average.

3 Methodology

This section outlines the causal inference framework, as well as the synthetic control method in the multiple treated unit case. I draw heavily from Cattaneo et al. (2023); those familiar with this setup may wish to skip to the next section.

3.1 Causal inference framework

Consider a set of N units observed over T time periods, with units indexed by $i = 1, 2, \ldots, N$, and time periods indexed by $t = 1, ..., T$. Each unit may be exposed to a treatment starting at different time periods. Let T_i represent the treatment adoption time for unit i; for untreated units, this is denoted by $T_i = \infty$. N_1 and N_0 represent respectively eventually treated units and never treated units. For convenience, units are ordered according to the time of adoption: $T_1 \leq T_2 \leq ... \leq T_{N_1}$. Following the intervention period, we only observe the treated outcome for the treated units, and the untreated outcome for the untreated. To capture the true impact of the intervention requires estimating the *counterfactual* outcome for the treated units: that is, if unit i had not received treatment, what would have occurred. To do so, we form a 'donor pool' of potential control units from the remaining $N - N_1 = N_0$ units that are not exposed to the treatment. The donor pool does not vary by the choice of the causal predictands below. Subsections 4.2 and 5.1 detail this process.

Let $Y_{it}(s)$ represent the potential outcome of interest for unit i at time t, if unit i had received treatment in period s, for $s = 1, ..., T, \infty$. This gives us the following equation for the observed outcomes:

$$Y_it = Y_{it}(\infty)\mathbb{1}(t < T_i) + Y_{it}(T_i)\mathbb{1}(t \geq T_i).$$

Of the range of causal predictands proposed by Cattaneo et al. (2023), I am interested in the following list. These capture both 1) the desired results of usual synthetic control analysis and 2) the new predictands that can be calculated with a multiple treated unit approach.

1. For a specific treated unit and $k \geq 0$, let τ_{ik} denote the individual treatment effect for unit i in period $T_i + k$. This represents the effect k periods after treatment initiation and is used to compute additional predictands. τ_{ik} is the usual predictand of interest in synthetic control analysis.

$$\tau_{ik} \coloneqq Y_{i(T_i+k)}(T_i) - Y_{i(T_i+k)}(\infty). \tag{1}$$

2. Average treatment effect on units treated at time s_0, k periods after treatment adoption. Given that only the Christchurch territorial authorities are treated at the same time, this predictand is equivalent to $\tau_{\cdot k}$ for the Christchurch-Waimakariri-Selwyn model.

$$\tau_{\cdot k, s_0} \coloneqq \frac{1}{|i : T_i = s_0|} \sum_{i:T_i=s_0} \tau_{ik}.$$

3. Average treatment effect on the treated, k periods after treatment adoption:

$$\tau_{\cdot k} \coloneqq \frac{1}{N_1} \sum_{i=1}^{N_1} \tau_{ik}.$$

In the staggered adoption model I restrict the number of periods after treatment to the smallest of all treated units, denoted K. Thus $\tau_{\cdot k}$ can only be considered for $k = 0, .., K$. The most recent treated unit is typically Lower Hutt, although the exact value of K depends on the model specification. Setting all $T_i = \infty$ for $i \geq 2$ reduces this framework to the classic single treated unit case[5].

3.2 Synthetic control method (SCM)

This subsection details the application of the synthetic control method to the multiple treated unit framework. The synthetic control is a weighted combination of units from the donor pool. This combination is set up to best match the characteristics of the treated units. Establishing the method requires an understanding of the potential donor pool, what variables are being matched on, and how the weighted combination is chosen. I outline the donor pool and matching variable selection method in Sections 4.1 and 4.2.

[5]For further discussion on how these predictands are constructed and estimated, refer to Cattaneo et al. (2021) and Cattaneo et al. (2022).

For instance, in this application, $\mathbf{A}^{[i]}$ contains the de-meaned housing consents per thousand residents and the suite of other matching variables of a treated TA i during the pre-treatment period. $\mathbf{B}^{[i]}$ contains the same information except for the donor pool used to match $\mathbf{A}^{[i]}$. For each matching variable, $\mathbf{C}^{[i]}$ contains an intercept and a linear time trend.

$$\mathbf{A}^{[i]} = \begin{bmatrix} \mathbf{A}_1^{[i]} \\ \vdots \\ \mathbf{A}_M^{[i]} \end{bmatrix}, \quad \mathbf{B}^{[i]} = \begin{bmatrix} \mathbf{B}_1^{[i]} \\ \vdots \\ \mathbf{B}_M^{[i]} \end{bmatrix}, \quad \mathbf{C}^{[i]} = \begin{bmatrix} \mathbf{C}_1^{[i]} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{C}_2^{[i]} & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{C}_M^{[i]} \end{bmatrix}$$

For each unit i, there is a matrix $\mathbf{A}$ representing the (pre-treatment) set of matching variables belonging to the treated unit i. A corresponding matrix $\mathbf{B}$ has the same information for the donor pools, also over the pre-treatment time span. Finally, for each treated i there is a matrix $\mathbf{C}$, that, for each matching variable, contains an intercept and a linear time trend.

The SCM's aim is to find a vector of weights $\mathbf{w} = (\mathbf{w}^{[1]\prime}, ..., \mathbf{w}^{[N_1]\prime})' \in \mathcal{W} \subseteq \mathbb{R}^{JN_1}$ across the M matching variables and a vector of coefficients $\mathbf{r} = (\mathbf{r}^{[1]\prime}, \dots, \mathbf{r}^{[N_1]\prime})' \in \mathcal{R} \subseteq \mathbb{R}^{KM}$, such that the linear combination of $\mathbf{B}_l$ and $\mathbf{C}_l$ approximates $\mathbf{A}_l$ as closely as possible for all $1 \leq l \leq M$. In principle, the SCM tries to create a 'psuedo' or 'synthetic' treated unit out of the donor pool, that closely resembles the set of matching variables (including the pre-treatment outcomes). Once this synthetic control is created, it is used as the counterfactual relative to which we measure the impact of treatment.

This match is found by the following minimisation problem:

$$\hat{\boldsymbol{\beta}} := (\hat{\mathbf{w}}', \hat{\mathbf{r}}')' \in \underset{\mathbf{w}\in\mathcal{W}, \mathbf{r}\in\mathcal{R}}{\operatorname{argmin}} (\mathbf{A} - \mathbf{B}\mathbf{w} - \mathbf{C}\mathbf{r})'\mathbf{V}(\mathbf{A} - \mathbf{B}\mathbf{w} - \mathbf{C}\mathbf{r}) \tag{2}$$

where

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \\ \vdots \\ \mathbf{A}_M \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \\ \vdots \\ \mathbf{B}_M \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} \mathbf{C}_1 & 0 & \cdots & 0 \\ 0 & \mathbf{C}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{C}_M \end{bmatrix}$$

Here $\mathbf{V}$ refers to the weighting matrix discussed in subsubsection 3.2.1.

The predicted counterfactual outcome of each treated unit is then constructed by

$$\hat{Y}_{it}(\infty) := \mathbf{p}_t^{[i]} \hat{\boldsymbol{\beta}}^{[i]},$$

where

$$\mathbf{p}_t^{[i]} = (\mathbf{x}_t^{[i]\prime}, \mathbf{g}_t^{[i]\prime})', \ i = 1, ..., N_1, \ t > T_i$$

where $\mathbf{x}_t^{[i]}$ represents the vector of predictors for the control units at time t that are utilized to estimate the counterfactual for the treated unit i, and $\mathbf{g}_t^{[i]}$ denotes the vector of predictors associated with the additional control variables defined in $\mathbf{C}^{[i]}$.

By substituting this estimate for $\hat{Y}_{it}(\infty)$ into equation 1, we get our range of *estimated* causal predictands.

In both the single and multiple treated unit cases there are a range of possible restrictions on the feasibility sets $\mathcal{W}$ and $\mathcal{R}$. Different restrictions on the feasibility sets correspond to different ways to minimise the 'distance' between the set of matching variables (including pre-treatment values of the outcome variable) of the treated units, and those of the synthetic control. Rather than choosing an ad-hoc set of restrictions on the weights, I select two popular types. These are:

1. Simplex. This type is the most common SCM formulation, and the most intuitive. According to Abadie et al. (2010) and Abadie et al. (2015), the simplex constraint is

the best at maintaining interpretability and ensuring over-fitting is avoided.

$||\mathbf{w}||_1 = 1, w_j \geq 0, j = 1, ..., J$

2. Lasso. This relaxes the requirement that the weights must sum to one and be non-negative, allowing for negative weights that satisfy the given L1 norm (usually Q = 1). It is better to suited to cases when pre-treatment fit of other options (such as Simplex constraints) is poor (Cattaneo et al. 2023).

   $||\mathbf{w}||_1 \leq Q$

Both Greenaway-McGrevy (2023) and Maltman & Greenaway-McGrevy (2024) follow Abadie et al. (2010) in restricting $\mathbf{w}$ to $w_i \in [0, 1]$ and $\sum_{i=2}^{N+1} w_i = 1$; equivalent to the Simplex approach.

3.2.1 Weighting matrix

Multiple treated units give us two ways of calculating equation 2. We may first allow the synthetic control weights to vary by treated unit, then average them. Otherwise, we could pool the treated units into a single treated unit, then proceed as normal in using the synthetic control method in the single treated unit case. This is represented by the choice of $\mathbf{V}$, the weighting matrix present in the objective function in equation 2.

This weighting matrix will determine how the sum of squared residuals are minimised. Choice of matrix is an evolving area, with a range of methods offered by both Ben-Michael et al. (2021) and Cattaneo et al. (2023). Broadly they all fall along the following spectrum: either completely separate, completely pooled, or some mixture of the two. 'Separate' weighting corresponds to the case where the weighting matrix is the identity matrix ($\mathbf{V} = \mathbf{I}$), and the minimised objective function is:

$$\sum_{i=1}^{N_1} \sum_{l=1}^{M} \sum_{t=1}^{T_0} (a_{t,l}^{[i]} - \mathbf{b}_{t,l}^{[i]\prime} \mathbf{w}^{[i]} - \mathbf{c}_{t,l}^{[i]\prime} \mathbf{r}_l^{[i]})^2. \tag{3}$$

This optimises the separate fit for each treated unit; minimising the sum of squared pre-treatment residuals for each individual treated unit.

A 'pooled' weight matrix is useful when the goal is to minimise the pooled fit (or the sum of squared averaged errors across treated units) of the treated units. With $\mathbf{V} = \frac{1}{N_!^2}\mathbf{1_{N_1}}\mathbf{1'_{N_1}} \bigotimes \mathbf{I_{T_0M}}$, where $\bigotimes$ signifies the Kronecker product operator, the minimised objective function is:

$$\sum_{l=1}^{M}\sum_{t=1}^{T_0}(\frac{1}{N_1^2}\sum_{i=1}^{N_1} a_{t,l}^{[i]} - \mathbf{b}_{t,l}^{[i]\prime}\mathbf{w}^{[i]} - \mathbf{c}_{t,l}^{[i]\prime}\mathbf{r}_{l}^{[i]})^2. \tag{4}$$

Ben-Michael et al. (2021) also allow one to supply a weighting matrix. This option could be chosen to prioritise weighting time periods closer to reform, or to favour the fit of certain treated units over others. This analysis uses a separate weighting matrix for simplicity, and to prioritise the fits of each individual synthetic control.

3.3 Assumptions

The causal inference and synthetic control environment outlined above embody a range of implicit assumptions. Briefly, these are:

1. **No interference**: The treatment status of unit i does not affect the outcomes of other units. This is equivalent to a 'Stable Unit Treatment Value Assumption' which requires that the outcome for a given unit depends solely on the treatment assigned to that unit, without being influenced by treatments assigned to other units: no spill-over effects' assumption.

2. **Parallel trends in the absence of treatment**: In the absence of treatment, the treated unit and its synthetic control should follow similar trends.

3. **No time-varying unobservable confounders**: It is assumed that any unobserved factors influencing the outcomes do not vary across treated and control units over time.

4. **Permanent treatment**: Once a unit is treated, it cannot be ‘untreated’.

5. **No anticipation**: Pre-intervention time outcomes are equal to those if there had been no intervention. If there is a risk of intervention, Abadie et al. (2010) suggests moving treatment time forward to when the policy might feasibly impact outcomes.

4 Data

Data on new housing consents per thousand residents are compiled by Statistics New Zealand and are available by region and territorial authority (Stats NZ 2024*b*). Per thousand figures are used to better compare construction between regions of different sizes. Treatment occurs at the territorial authority level, data at this level are used: a region may be composed of areas that did and did not implement zoning reform. For instance, the Wellington region includes Lower Hutt (which has been up-zoned) and Masterton (which has not). I discuss the issue of treatment level further in Section 6.

There are 67 territorial authorities in New Zealand, comprising 13 city councils and 54 district councils. City councils tend to govern urban areas with larger populations, while district councils govern mixed rural and urban areas with smaller populations. Alongside handling local services (such as water supply, sewage, local roads and community facilities), they also manage land use consents and local land regulation. Time series data is available from April 1990 through to May 2024, filtering for new dwellings; each dwelling in a housing complex or apartment building is tallied individually, with territorial authority areas as at 2021.

Although monthly data are available, I elect to use data aggregated to the calendar year for three reasons. One, to match the approaches of past literature on the topic (Maltman & Greenaway-McGrevy 2024). Two, to address the scale of the outcome variables. Even when using housing consents per thousand residents, many of the potential donors are small TAs, with few housing consents issued in every month for analysis. Aggregating to an annual scale still gives enough time periods to establish a strong pre-treatment fit, while keeping the donor pool as large as possible, in accordance with Abadie et al. (2015). Three, to avoid issues of varying seasonality between units. It is possible that different territorial authorities experience different monthly construction patterns (such as winter conditions being more impactful in colder climates), and the proposed framework is unable to capture such variability.

In line with Ferman & Pinto (2021), I improve pre-treatment fit by normalising the outcome variable. To do so, I subtract the pre-treatment mean housing consent per thousand residents by territorial authorities from all consent values for that given territorial authority. Results depicted are de-meaned unless explicitly mentioned.

4.1 Donor pool

In all analyses I exclude the Chatham Islands territorial authority. These islands are 650km off mainland New Zealand with a resident population of just 730, by far the smallest and most remote council. Its size and low rate of consenting render it inappropriate as a potential donor.

Spillover effects are a concern in donor pool construction (McClelland & Gault 2017). This occurs when there is a risk that areas nearby to the treated unit receive flow-on effects from the treatment, such as diverted housing production away from neighbouring Upper Hutt to Lower Hutt. The inclusion of spillover-affected units in the donor pool will result in biased estimates, as the synthetic control's outcomes will be lower than if the treated unit had never received treatment.

Auckland is bordered by Waikato and Kaipara Districts, and I also exclude these from the analysis. Similarly for Lower Hutt, I exclude Wellington City, South Wairarapa, Upper Hutt, and Porirua.

The three up-zoned South Island TAs (Christchurch, Selwyn, Waimakariri) are neighboured by Hurunui, Westland, and Ashburton. I exclude Hurunui and Ashburton only from the donor pool. Westland is the least dense TA, in both population per kilometre-squared and dwellings per capita, and is geographically separated from the earthquake-affected TAs by a significant mountain range. As a result, it is not at practical risk of spillover effects.

The final donor pool size is shown in Table 1.

Treated Units	**Donor Pool Size**
Auckland, Lower Hutt	55
Christchurch, Waimakariri, Selwyn	55
Auckland, Lower Hutt, Christchurch, Waimakariri, Selwyn	53

Table 1: Donor Pools

4.2 Matching variables

The primary variable used to match potential control units with the treated unit is the time series of the pre-treatment outcome variables (Abadie et al. 2010). Alongside this, it is common to include a set of additional variables on which to match, particularly in the case when there are insufficient pre-treatment outcome variables[6]. As housing consents are only tracked at a territorial level since 1996, we have between 17 and 22 pre-treatment values depending on the specific treated unit. Further matching variables will likely be needed to tighten the synthetic control fit.

The role of an expanded matching variable selection have been disputed in literature, with Doudchenko & Imbens (2016) downplaying the value of additional covariates when constructing the synthetic control. Pickett et al. (2022) questions this conclusion, asserting that they play a critical role in the case of short pre-treatment time series, and given how their ab-

[6] Also known in the staggered adoption literature as 'auxiliary covariates' and 'covariate features', see: Ben-Michael et al. (2021) and Cattaneo et al. (2023).

sence can lead to significant changes in point estimates. Matching variables should be good predictors of the outcome variable before and after treatment, with the goal of improving the match between treated units and their synthetic controls (McClelland & Gault 2017). For robustness I present results from two specifications of matching variable: one with only pre-treatment outcomes (the main specification), and one with pre-treatment outcomes and extra covariates (the expanded specification). Expanded specification results are displayed in Appendix C.

In large part I mirror Greenaway-McGrevy (2023) in selecting matching variables, with slight differences due to missing 2001 census data.

I acquire sub-national population projects from Stats NZ, looking at both 2013 and 2018 base projections, as well as actual population growth between census years 2006, 2013, and 2018 (calculated by taking the difference of logged population counts). I also include dwellings per capita in census years to capture the differences in dwelling composition by territorial authority. Finally, to account for the impact of income on housing choices I include the logged difference of median household income between census years. I note that data was only available for two censuses (2006 and 2013), meaning there is only one value for this matching variable. Table 2 presents the 2013 values for each of these matching variables for each of the treated units, as well as the national average.

Name	**Dwellings per capita**	**Increase in median household income, %**	**Population growth, %**	**Projected population growth, %**
Auckland	0.333	18.8	8.1	1.5
Lower Hutt	0.367	20.3	0.6	0.2
Christchurch	0.382	30.3	-2.0	0.8
Waimakariri	0.373	30.1	15.4	1.6
Selwyn	0.340	30.9	28.2	2.6
New Zealand	0.368	21.6	5.2	1.1

Table 2: Matching variables for each treated unit and the national average

The synthetic control method implemented by Cattaneo et al. (2022) accepts two values for each matching variable: a constant and time trend. When necessary (such as in the case of dwellings per capita) I supply the pre-treatment mean value of the matching variable, as suggested by McClelland & Gault (2017). An alternative approach (not explored here) would be to take only the final pre-treatment value. I later demonstrate that even the inclusion of matching variables makes little difference to the overall findings.

4.3 Implementation

Empirical estimation was carried out in the statistical programming language R using the code package `scpi` (Cattaneo et al. 2023). Code is available on github at `https://github.com/JamesHancock1/thesis_zoning_reform.git`. For more information on the software implementation of the staggered adoption method, consult Cattaneo et al. (2022).

5 Findings

All results in this section are from the main matching variable specification using the Simplex restrictions, with Appendix C containing the expanded specification results. Appendix A displays the results of the model when Auckland and Lower Hutt are the only treated units; the equivalent for the Christchurch region territorial authorities is shown in Appendix B. Finally, Appendix D contains results with the LASSO set of feasibility restrictions.

Table 3 presents the chosen weights for both the main and alternative specifications. As outlined in Section 1, the donor pool here consists of 57 potential donors, with territorial authorities at risk of spillover excluded. Either five or six donors are used for each treated unit, with weights ranging from 0.004 for Christchurch City (the Mackenzie District) to 0.593 for Lower Hutt (Timaru District). Twenty of all 57 available donors are used, with 15 from the North Island of New Zealand and five from the South. Hamilton City and the Tasman District are the top contributors to the synthetic controls, supplying 0.657 and 0.607 of the

Donor unit	Auckland	Lower Hutt	Christchurch	Waimakariri	Selwyn
Carterton District	-	-	-	-	0.461
Gisborne District	-	-	-	0.121	-
Hamilton City	-	-	0.359	0.298	-
Kapiti Coast District	0.206	-	-	-	-
Mackenzie District	-	-	0.004	0.067	0.091
Masterton District	0.146	-	-	-	-
Matamata-Piako District	-	-	-	-	0.077
Nelson City	-	-	0.050	-	-
New Plymouth District	-	-	0.122	-	-
Opitiki District	-	0.087	0.120	0.213	-
Queenstown-Lakes District	0.022	-	-	-	0.055
South-Waikato District	-	0.023	-	-	-
Stratford District	-	-	-	0.104	-
Tasman District	0.434	0.070	0.103	-	-
Taupo District	-	-	0.009	-	-
Tauranga City	0.191	-	0.078	0.196	-
Thames-Coromandel District	-	-	-	-	0.316
Timaru District	-	0.593	-	-	-
Waipa District	-	0.078	0.155	-	-
Wairoa District	-	0.149	-	-	-

Table 3: Main model specification donors

weight across the five treated units.

These weights are robust to matching variable choice, featuring almost identical weights to the expanded specification (see Appendix Table 6). The same is true for when the model is split into the two considered subgroups of Auckland-Lower Hutt and Christchurch-Selwyn-Waimakariri, as seen in Appendix Tables 4 and 5. The Lower Hutt composition is identical in donors (but not donor weights) to the main specification in Maltman & Greenaway-McGrevy (2024).

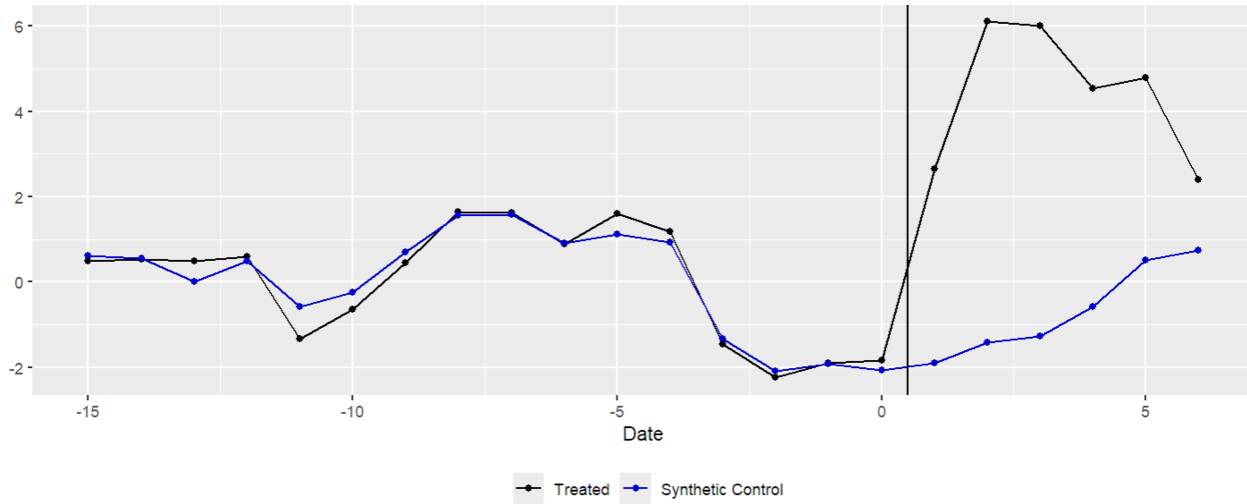

Figure 6: Average treatment effect on the treated, normalised consents per thousand residents

5.1 Aggregate effects

Figure 6 exhibits the effect of reform in each period after treatment averaged over all treated territorial authorities, with all five treated units included. As noted in Section 3.1, the time frame is limited to the shortest post-reform period: six years. There is a noticeable divergence in the paths of the actual consents versus the synthetic estimate, peaking in the second year following reform. The two paths begin to converge beginning in the fourth year, suggesting that there may be a limit to the ongoing impact of reform. The model fits well prior to treatment, particularly in the five years leading up to reform. Given this is an averaged result, the strong pre-treatment fit suggests that the individual synthetic controls could fit similarly well.

5.2 Individual unit effects

An initial visual inspection of pre-treatment fit is strong, with the synthetic control closely matching that of the actual consent pattern in both Figures 6 and 7. The only exception in the individual case is Waimakariri where the synthetic control fails to capture the variability in consents immediately pre-treatment.

Figure 7 presents the individual impact of reform on each individual treated unit over time, as well as the entire pre-treatment time series.

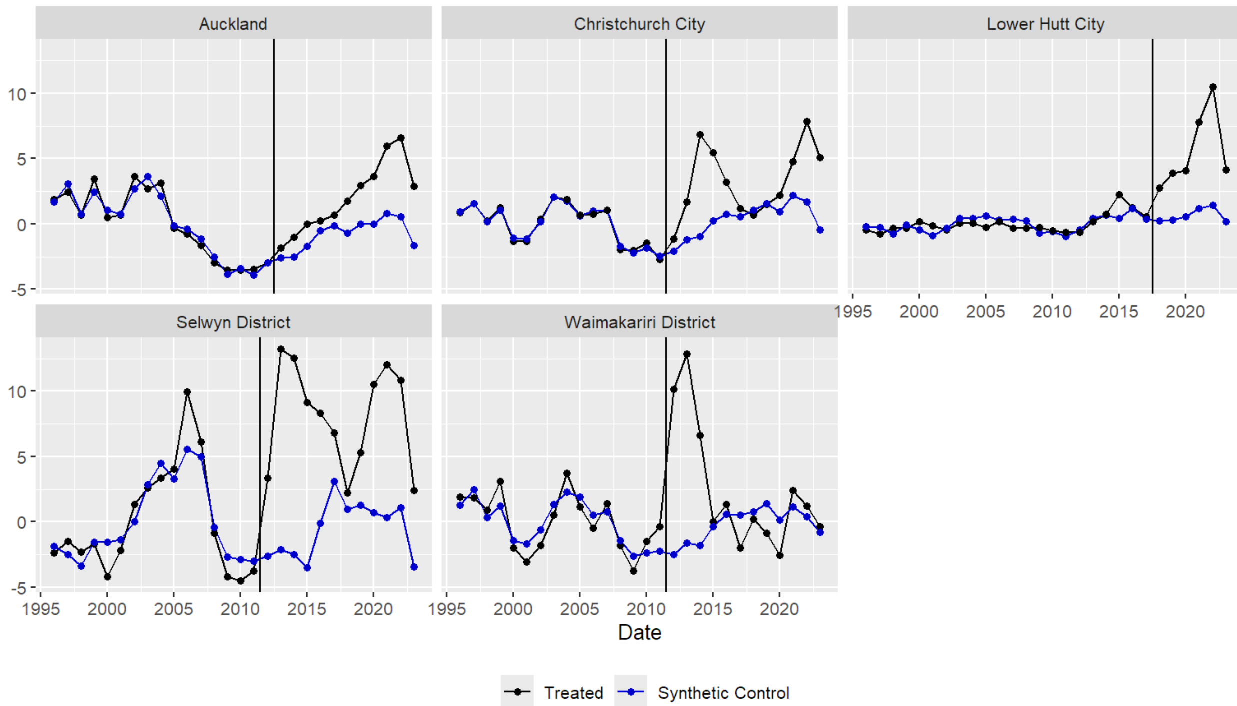

Figure 7: Individual synthetic and actual normalised consents per thousand residents in the five treated areas

In the three Christchurch region territorial authorities there is a clear separation between actual consents and the synthetic control estimates. This separation only appears to hold for the short and long terms, with a 'U' shape consent pattern evident in both Christchurch and Selwyn. There is a noticeable lull in impact from 2015 to 2020. Increases in consent variability could indicate the underlying housing market is becoming more elastic to changes in price, with developers able to react at a quicker rate. Waimakariri experiences a large initial increase before quickly dropping to a negligible difference post-2015. Waimakariri (and briefly, Christchurch) is the only territorial authority where actual consenting patterns drop below the synthetic control estimate for the same period, suggesting that some of the initial increase may have simply brought future consenting forward. Considering the immediate loss of housing stock, this is a desirable and expected outcome from earthquake-recovery focused zoning reform.

Although both Auckland and Lower Hutt appear to benefit from zoning reform, the results differ from those of the Christchurch region. The synthetic Auckland consenting trend closely matches that of actual consents, before diverging post-2016. With the *Auckland Unitary Plan* only becoming fully operational in 2016 this finding makes sense: the initial 2013-2016 period is marked by a level increase to consents due to early access afforded to some developers, before a structural change in 2016. The impact of reform peaks in 2022 at 6 consents per thousand residents more than the synthetic control estimate, which corresponds to consents per capita increasing due to up-zoning by approximately 105% per cent by 2022. Even though Greenaway-McGrevy (2023) use a slightly different treatment area, the study finds almost the same result over the given time span.

Lower Hutt presents the most striking evidence for the impact of zoning reform. Synthetic control estimates are largely flat from 2018 and onwards, while actual consents per thousand residents climb to a peak of 12.2 consents. I note that the Lower Hutt City results closely match those of Maltman & Greenaway-McGrevy (2024), which is to be expected given the overlap in donor pool.

Present in each sub figure is a drop in consents per thousand residents between 2022 and 2023, in both the actual values and synthetic control estimates. This decrease is due to a general drop in construction across New Zealand, with the sector facing uncertainty around future government contracts, weakening demand, increased building costs, and labour shortages (Ministry of Business, Innovation and Employment 2024*a*). The level gains from zoning reform persist for all treated units except Waimakariri (and to a lesser extent, Lower Hutt).

5.2.1 In time placebo check

I conduct an *in time* placebo check (Ben-Michael et al. 2021). This validation approach tests the method's reliability by artificially shifting the treatment period backwards in time. It creates “fake” (placebo) treatment periods during the pre-treatment phase, when no actual intervention occurred. I modify the treatment time index T_i by subtracting x periods $(T_j - x)$, effectively moving the treatment earlier. I then estimate treatment effects during these pre-intervention periods, where we should find no significant effects if our method is working correctly. Following Ben-Michael et al. (2021) I test one and two periods before the true treatment time. These periods are 2011 and 2012 for Auckland, 2016 and 2017 for Lower Hutt, and 2010 and 2011 for the Christchurch region TAs. Figure 8 presents these results. For both the one period and two periods before models, the synthetic control estimates immediately following (placebo) treatment are almost identical to the actual path of consents. The results show that the models match the pre-intervention data extremely well and, when tested with placebo treatments, find no false effects - exactly as we would expect when no real treatment has occurred.

Figure 8: Average treatment effect on the treated by time period

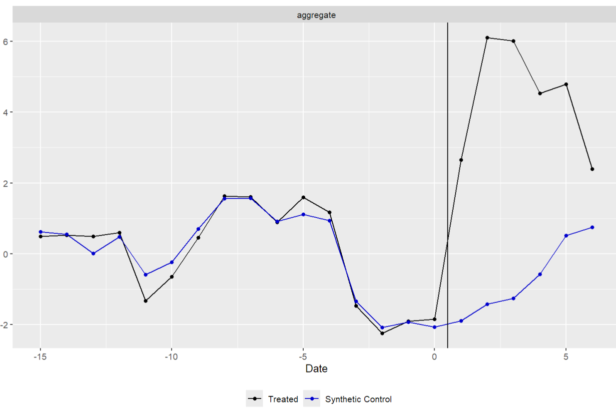

(a) Original treatment time

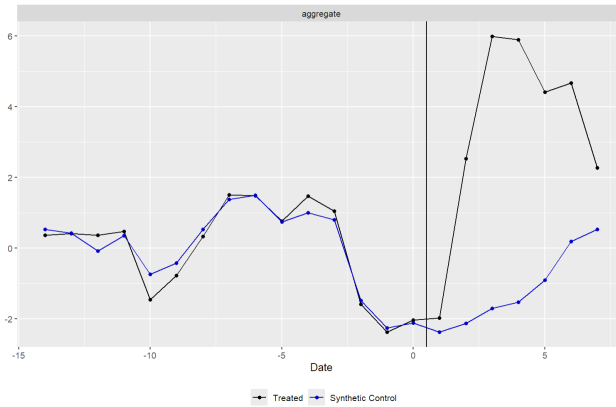

(b) One period before

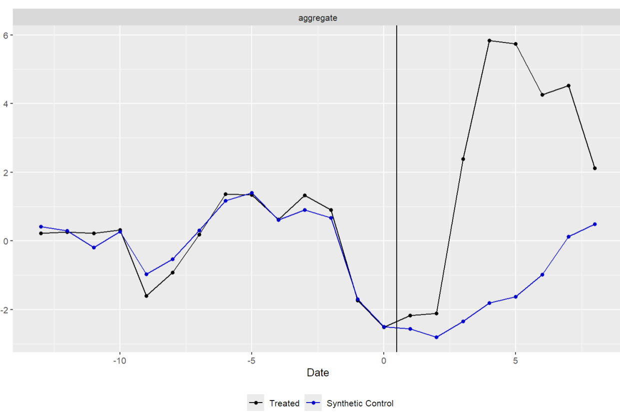

(c) Two periods before

5.3 Inference

Inference in the staggered adoption version of the SCM is a growing area of literature. I briefly outline typical SCM inference methods, before applying the framework developed in Cattaneo et al. (2021) and Cattaneo et al. (2023).

In a single treated unit SCM it is common to use design-based rank permutation test for significance. Such a test involves fitting a synthetic control for every member of the donor pool as a set of placebos, then calculating each of their respective prediction errors. For instance, Maltman & Greenaway-McGrevy (2024) compute the positive ratio of the Mean Squared Prediction Error between pre- and post-treatment using the formula below. These ratios are then ranked from smallest to largest; the larger the prediction error, the closer to 1 on the list it falls. If the treated unit is placed first of all n units, this means that if treatment had been assigned at random, there is only a $\frac{1}{n}$ chance of getting a ratio as large as that of the treated unit.

$$R^{+}(t_1, t_2) = \frac{1}{t_2 - t_1} \sum_{t=t_1}^{t_2} \left(\lfloor Y_{i,t} - Y_{i,t}^{N} \rfloor \right)^2$$

where $\lfloor x \rfloor = 0$ if $x > 0$ and $\lfloor x \rfloor = x$ otherwise.

Cattaneo et al. (2021) argue that, in the case of staggered adoption, this method provides insufficiently wide prediction intervals and propose the more conservative approach detailed below. Furthermore, it is unclear how best to implement it, given the range of treated units. While it is possible to estimate a synthetic control for every donor, it is unclear what the treatment time would be for each, as well as how many pseudo-treated units are included in each model estimation.

5.3.1 Alternate approach

For the reasons above I explore the statistical significance of the reform findings using the Cattaneo et al. (2021) framework. The authors conduct inference by first treating the (potential) outcome variables as random, then using finite-sample probability concentration methods to construct prediction intervals for the counterfactual estimates. They quantify two sources of uncertainty: (i) the in-sample error from estimating the weights of the synthetic control, and (ii) the out-of-sample error components from treatment effects. The authors then split up $Y_{it}(\infty)$ into these two components:

$$Y_{it}(\infty) = \mathbf{Y}'_{\mathcal{N}_t} \mathbf{w}_0^{[i]} + u_{it}, \quad i \in \mathcal{E}, \quad 1 \leq t \leq T, \tag{5}$$

where u_{it} out-of-sample error component, $\mathcal{N} = \{i : T_i = \infty\}$ represents the group of "never-reformed" territorial authorities, and $\mathcal{E} = \{i : T_i < \infty\}$ represents the group of treated units. For each area $i \in \mathcal{E}$ that has received treatment, we can separate the prediction of the counterfactual outcome as follows:

$$\hat{Y}_{i(T_i+k)}(\infty) - Y_{i(T_i+k)}(\infty) = \mathbf{Y}'_{\mathcal{N}(T_i+k)} \left(\hat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]} \right) - u_{i(T_i+k)}, \quad (6)$$

The first term on the right hand side of 6 reflects the in-sample uncertainty arising from the construction of the synthetic control weights based on pre-treatment data, while $u_i(T_{i+k})$ represents the out-of-sample uncertainty due to stochastic errors during a specific post-treatment period. In-sample errors are quantified using simulation, while out-of-sample errors are assumed to follow a sub-Gaussian distribution conditional on the outcomes of the treated units. Prediction interval bounds around the synthetic control estimate are set at 90% (Cattaneo et al. 2023).

Full error equations are detailed in Appendix F.

5.3.2 Inference results

Figure 9 displays the average impact of zoning reform on the treated units over time with 90% prediction intervals included. The synthetic control estimate is statistically significant in each period except the sixth. There is clear evidence of zoning reform having an effect on housing consents immediately following reform, with this effect weakening over time. Prediction intervals at the individual territorial authority level are depicted in Figure 10. An interesting split between Auckland - Lower Hutt and the Christchurch region forms, with the former having statistically significant results later following treatment while zoning reform was more impactful early on for the latter. This difference could be due to the differing nature of reform in the two groups. Christchurch's zoning changes involved up-zoning large amounts of underdeveloped greenfield land, comparatively easier to develop compared to the built-up urban landscapes of Auckland and Lower Hutt. This is just one possibility, with

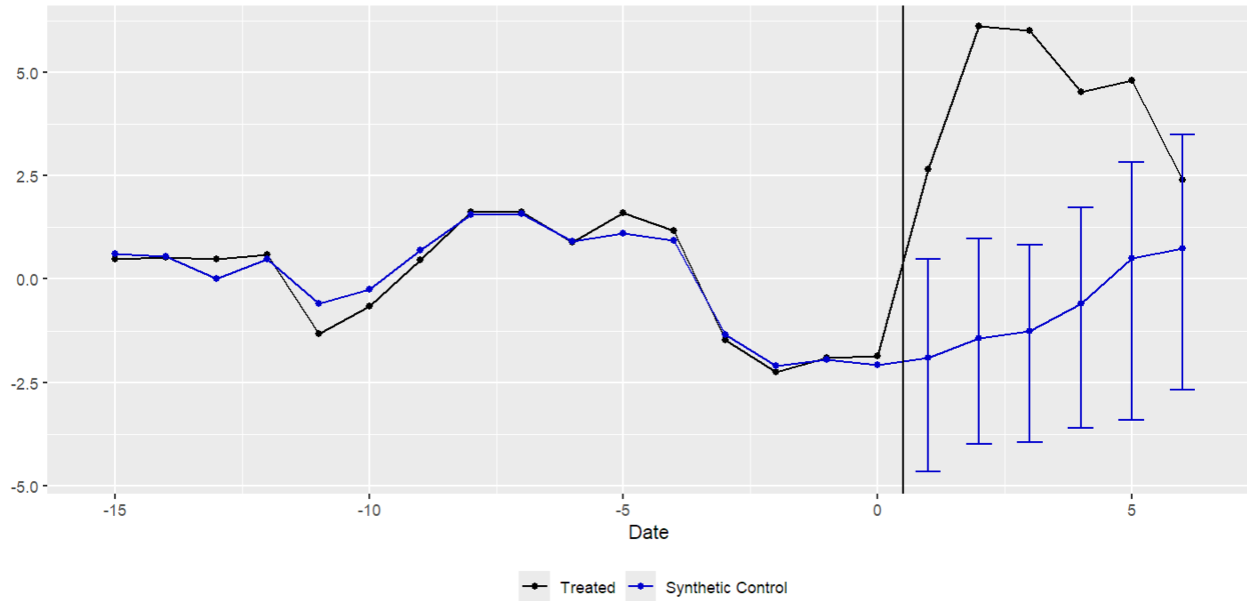

Figure 9: Average synthetic and actual normalised consents per thousand residents, with prediction intervals

more research needed assess potential reasons for the difference in treatment response times.

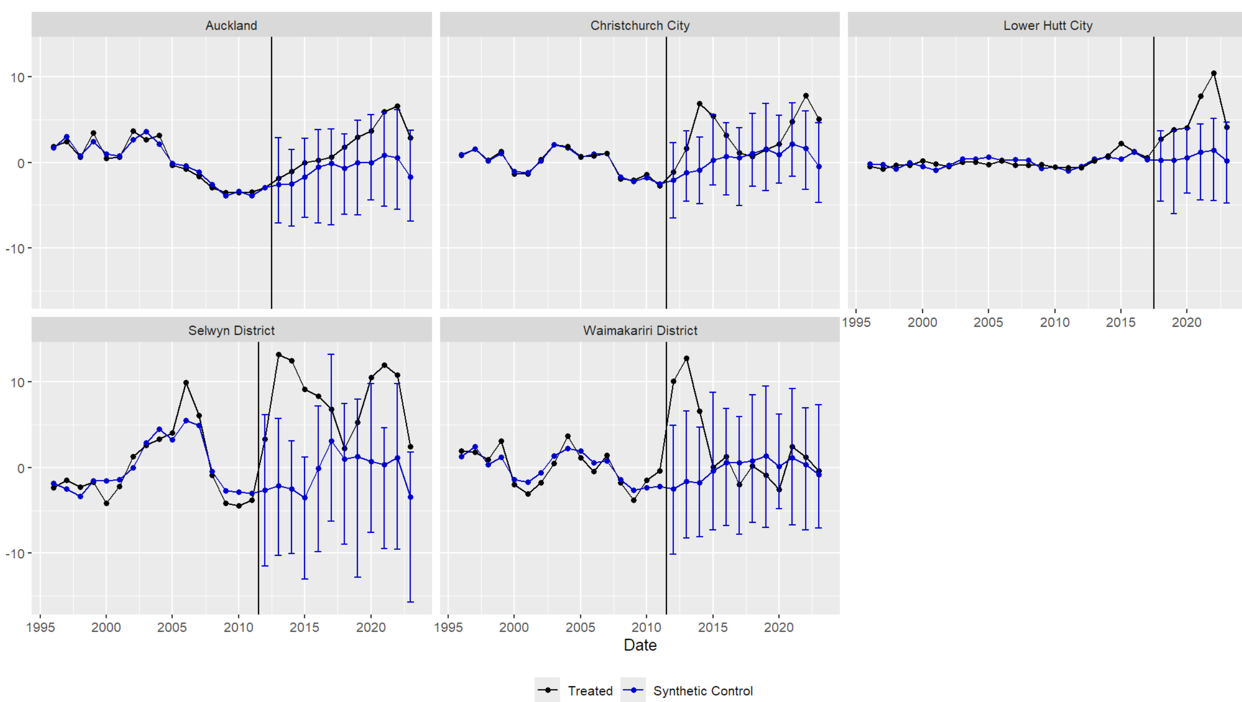

Figure 10: Individual synthetic and actual normalised consents per thousand residents, with prediction intervals

5.4 Inference results with LASSO constraints

Appendix Figures 16 and 17 depict normalised consents per thousand residents with the LASSO constraints on the two feasibility sets, as outlined in Section 3. The change in constraints produces pre-treatment estimates that are significantly closer to actual consenting patterns, yet with similar post-treatment estimates to the Simplex case. Improved pre-treatment fit comes at the cost of far wider prediction intervals in the post-treatment period, with decreased statistical significance of the Auckland, Christchurch, and later-period average results. The wider bounds are indicative of the trade-off with alternate feasibility set restrictions between pre-treatment fit and error boundaries, suggesting that the Simplex constraint (combined with normalising the outcome variable) is the preferred option when pre-treatment fit is sufficiently close.

5.4.1 'Leave one out' check

A final common robustness check is the 'leave one out' method. This check involves removing the most important donor of each treated unit and re-running the synthetic control, then inspecting to see if the estimated results change. These are Tasman District (for Auckland), Timaru District (for Lower Hutt) and Carterton District (for Selwyn), as well as Hamilton City (for both Christchurch and Waimakariri). Robust results should be largely unchanged, as shifts in predictions could indicate that the original estimates were being overly influenced by a single outlier donor. The new donor pool is shown in Appendix table 7. Appendix Figure 18 depicts the individual unit level estimates, with prediction intervals. A comparison of Appendix Figure 19 - the averaged impact of reform - to Figure 9 shows identical patterns. While most of the individual effects do not change, Auckland's results change notably, with reform now having a statistically significant impact during 2020-2022. This signifies that the multiple unit framework may have different levels of robustness for individual units, which should be carefully considered when deciding upon the donor pool. The original donor pool yielded results that differed in large regard from Greenaway-McGrevy (2023), while the restricted pool is much closer aligned. The bulk of trends are stable in both pre- and post-intervention periods, even when important TAs are removed from the donor pool, which indicates that my findings are robust and not driven by a particular territorial authority.

6 Discussion

In this section I provide a brief overview of the limitations of my estimation method, as well as a discussion on how future estimation methods may address these.

6.1 Continued supply of housing supply reform

There is no sign that supply-side reform has finished in New Zealand. The current National-led government continues to pass reforms that will impact the consenting process, including removing the need for a consent for small additional dwellings on a land parcel (Ministry of Business, Innovation and Employment 2024*b*). The *MDRS* and *NPS-UD* are still in effect[7], meaning local governments across the country will be required to institute reforms - despite ongoing delays. Although the results from this dissertation and other literature suggest reform will make inroads in addressing housing shortages, policy analysis will become more difficult as treatment adoption increases. The SCM - and other causal techniques - rely on a supply of untreated units to form a suitable control. Under the framework used in this dissertation, newly treated units are removed from the donor pool. As reform continues, the donor pool will eventually be unable to generate 'good enough' pre-treatment fits without resorting to extrapolation beyond the convex hull, as discussed in Section 3.2. Alternate methods that account for such issues will be needed. One possibility is to look elsewhere for donor units. There is precedence for this: Greenaway-McGrevy (2023) find that, when added, Australian cities and states do not form a significant part of the selected synthetic control. Another option is to allow for units that may be treated in the future to be used as

[7] Albeit with the *MDRS* shifting to be opt-in, as discussed in Section 2.

donors until they are eventually treated. Although there are recent frameworks that allow for this (See: Athey & Imbens (2022), Freedman et al. (2023)), there are none that also use an SCM.

6.2 Decomposing the impact of individual policies

Another issue with generalising zoning reform analysis is that, as demonstrated in Section 2, it can consist of a range of policies. Such policies are often introduced simultaneously or near enough together to make it difficult to decompose any increase in consenting patterns into the effects of individual policies. For instance, a council could remove requirements for a balcony, garage, or second set of stairs in a multi-storey building. It could reduce the necessary distance between the outer wall of a house and the property's fence. Although each will remove construction barriers, their impacts will likely differ. The heterogenous impacts of a range of policies cannot be considered by the proposed framework, unless the same set of policies are implemented in different territorial authorities. One way to address this is careful consideration of other data sources. For instance, a policymaker could examine the impact of the removal of the setback requirement by tracking the average setback by development, then estimating the increases in floor area this measure may have caused. If such a policy was introduced at the same time as removing minimum parking requirements, a measure like new dwelling consents would not be sufficient to track underlying trends.

Another way this could be mitigated empirically is in the choice of the weighting matrix $\mathbf{V}$. If a measure of zoning reform were developed, it could be used to categorise the type of treatment that each unit receives at a given time period. The treated units' contribution to the overall treatment effect could be weighted according to how much zoning reform they had received, rather than weighting equally or by minimising the pooled fit for the average of the treated units. Although this would likely result in a poorer fit, it would allow for an improvement in cross-treatment comparability.

6.3 Changing treatment status over time

Just as it is possible for a more liberal local government to relax zoning laws further (such as in Lower Hutt), it is also possible for conservative legislators to 'down-zone' previously up-zoned land. There are certainly prior New Zealand examples of this: during the 1970s and 1980s councils across the country down-zoned both near city centres and on the outskirts of their land, causing an estimated 69% rise in house prices relative to if there had been no (down-)zoning reform (New Zealand Infrastructure Commission 2022). However, given the tide of reform in favour of loosening LURs, a change in tack is unlikely.

As referenced in Section 3.3, permanent treatment is a core model assumption. If down-zoning were to occur, my proposed framework would not be able to account for it, outside of manually excluding such data points in treatment effect calculations. The impact of a one-off case of down-zoning could be studied through a classic synthetic control framework, but such an approach would have to carefully disentangle the impact of down-zoning directly from the lagged effect of prior LUR reforms.

6.4 Treatment scale

Section 2 discusses how zoning reform is enacted at a local government level, but enabled by central government. Perhaps more accurately, zoning reform often takes place at the individual dwelling and suburb level, often up to the judgement of council planning officials. This scale of reform is referred to as 'spot' up-zoning, and can have markedly different effects to zoning reform on wider scales (Gabbe 2019). Büchler & Lutz (2021) find a 10 per cent expansion in zoned development capacity in Zurich results in a 1.2 percentage point growth in the housing stock over a five-year period, but a Dong (2021) study of Portland spot up-zoning finds far more muted impacts on new dwellings. The initial post-earthquake up-zoning of particular suburbs in Christchurch is a further example of this policy, where select parts of a territorial authority were up-zoned while others were not.

A clear distinction has been made between spot up-zoning and between 'widespread' up-zoning, where entire regions experience reform. Auckland is one example, with its $\sim 75\%$ of residential land up-zoned, and so is Lower Hutt (even though as yet there are no public estimates of zoning reform coverage). Perhaps more interesting than estimates on a territorial authority scale are a micro-level analysis. This analysis could consider the impact of zoning reform on the likelihood of a dwelling on an up-zoned land parcel being developed versus one on a land parcel that has not experienced reform.

6.5 Feasibility set restrictions

There are two further typical options for the feasibility sets other than the Simplex and LASSO approaches described in Section 3. These are Ordinary Least Squares (OLS) and Ridge Regression, which pose further research options when pre-treatment fit is poor under the two options presented here.

One of the core assumptions of the SCM is that we are seeking to obtain estimates that fall in the convex hull of the donor pool; that is, we do not want to extrapolate outside the support of the data (Abadie et al. 2015). It is possible instead to use an Ordinary Least Squares approach to estimate the vector of weights by relaxing this constraint. This approach is equivalent to an unconstrained problem, and will likely produce weights that cause the synthetic control to match exactly the actual pre-treatment values (Cattaneo et al. 2023).

The second option is Ridge Regression. This method is suggested and implemented in the Augmented SCM from Ben-Michael et al. (2021). Again the simplex-type constraint is weakened (this time on the L2 norm), with Q a tuning parameter computed as follows:

$$(]||\mathbf{w}||_2 \leq Q = \frac{(J+KM)\hat{\sigma}_u^2}{||\hat{\mathbf{w}}_{OLS}||_2^2}$$

Each method can lead to over-fitting, leading to insufficient variability with which to form prediction intervals on the synthetic control estimates. During initial analysis this was an issue with both methods, but could potentially be fixed by 1) purposefully restricting the available donor pool, or 2) by increasing sparsity in the data by including fewer matching variables or pre-treatment time periods. Cattaneo et al. (2022), the software-implementation focused companion to their other papers, addresses this issue further.

6.6 Christchurch region inclusion

Alongside heterogeneous types of reform, there is also the concern of large exogenous shocks to the outcome variable. Abadie (2021) note that potential control units that experience such a shock should be excluded from the donor pool; Maltman & Greenaway-McGrevy (2024) remove Christchurch for this reason. There has been less attention placed on when the region with a shock is a treated unit. I justify the inclusion of the Christchurch region in two ways. Firstly, the contents of the $LURP$ and the actions of the Earthquake Recovery Minister are genuine examples of zoning reform, even if they are introduced for different reasons. Secondly, the proposed method still involves forming individual synthetic controls of each treated unit, and thus including Christchurch does little to impact the estimates of Auckland, or Lower Hutt. This outcome is shown explicitly in Appendix A, where results are robust to restricting the treated units to solely those two areas.

7 Conclusion

This research aimed to expand the empirical frameworks used to estimate the impact of zoning reform in New Zealand. I allow for multiple units that may be treated at different time periods, and show how the model may be updated with future examples of zoning reform. The new framework allows for a wider range of effects to be investigated, at the cost of computational complexity and donor pool restrictions.

I identify a strong positive impact of zoning reform on housing construction in the years immediately following reform, peaking at an additional 7.5 consents per thousand residents relative to the synthetic control estimate. These estimates are robust to different model specifications and optimisation techniques, with some differences arising in the levels of statistical significance. My analysis, using multiple statistical approaches and robustness checks, confirms Maltman & Greenaway-McGrevy (2024)'s findings regarding Lower Hutt housing consents and Greenaway-McGrevy (2023)'s conclusions on Auckland's up-zoning (albeit with looking at a slightly different treatment area). The impact of zoning reform in inducing extra construction in the Christchurch region, even if of a different nature to recent policy shifts, is also clear. These findings add another data-point to the debate on zoning reform as a method of addressing shortages in the construction of housing.

The local government-adaption of central government level policy has led to a wide range of responses to the push for zoning reform. This dissertation has exploited these differences as a natural experiment to see how three examples of zoning reform has impacted their associated rates of housing construction. As the push from central government to densify continues, more tools will be needed to be developed to assess the efficacy of zoning reform as a policy instrument. This framework is one such tool.

Appendices

A Additional Results: Auckland City and Lower Hutt City

Donor unit	Auckland	Lower Hutt
Kapiti Coast District	0.206	-
Masterton District	0.146	-
Opitiki District	-	0.087
Queenstown-Lakes District	0.022	-
South-Waikato District	-	0.023
Tasman District	0.434	0.067
Tauranga City	0.191	-
Timaru District	-	0.593
Waipa District	-	0.078
Wairoa District	-	0.149

Table 4: Auckland and Lower Hutt only donors

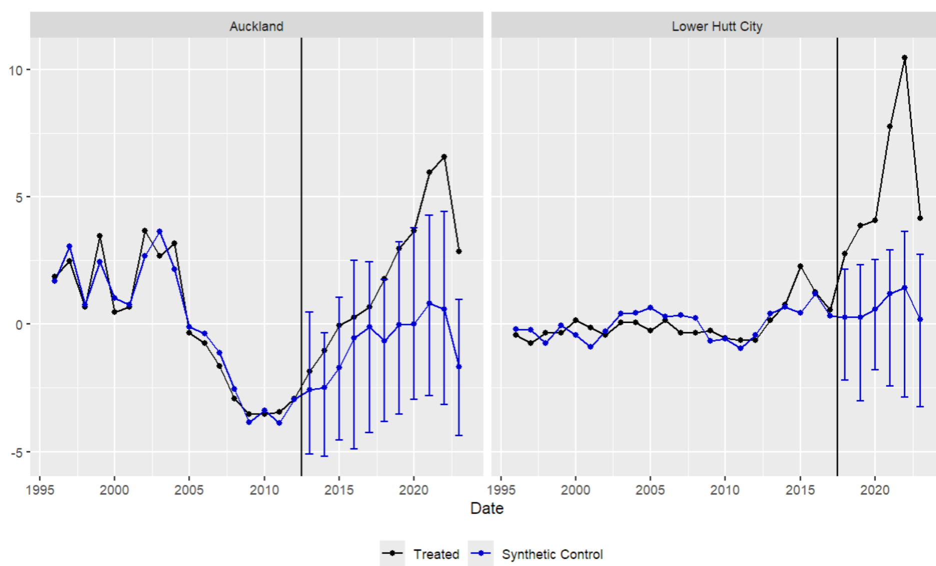

Figure 11: Auckland and Lower Hutt synthetic and actual normalised consents per thousand residents, with prediction intervals

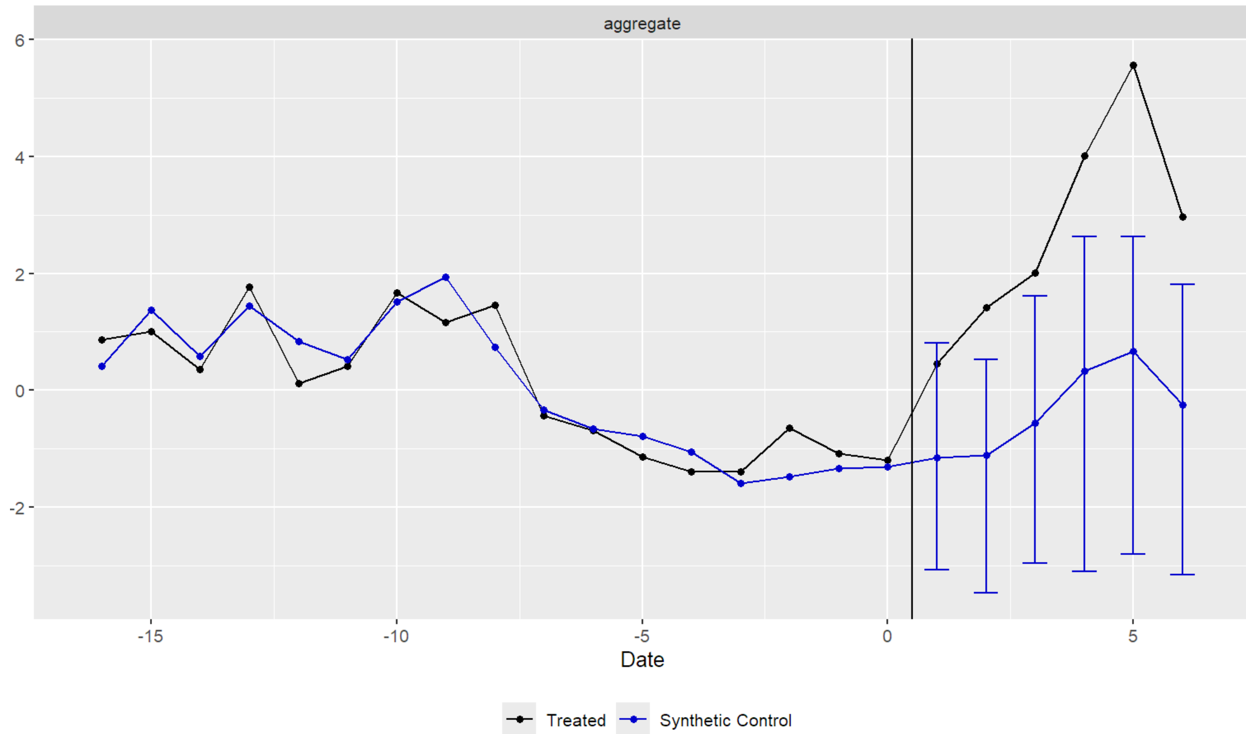

Figure 12: Average Auckland and Lower Hutt synthetic and actual normalised consents per thousand residents, with prediction intervals

B Additional Results: Christchurch City, Waimakariri District and Selwyn District

Donor unit	**Christchurch**	**Waimakariri**	**Selwyn**
Carterton District	-	-	0.461
Gisborne District	-	0.121	-
Hamilton City	0.359	0.298	-
Mackenzie District	0.004	0.067	0.091
Matamata-Piako District	-	-	0.077
Nelson City	0.050	-	-
New Plymouth District	0.122	-	-
Opitiki District	0.120	0.213	-
Queenstown-Lakes District	-	-	0.055
Stratford District	-	0.104	-
Tasman District	0.103	-	-
Taupo District	0.009	-	-
Tauranga City	0.078	0.196	-
Thames-Coromandel District	-	-	0.316
Waipa District	0.155	-	-

Table 5: Christchurch region only donors

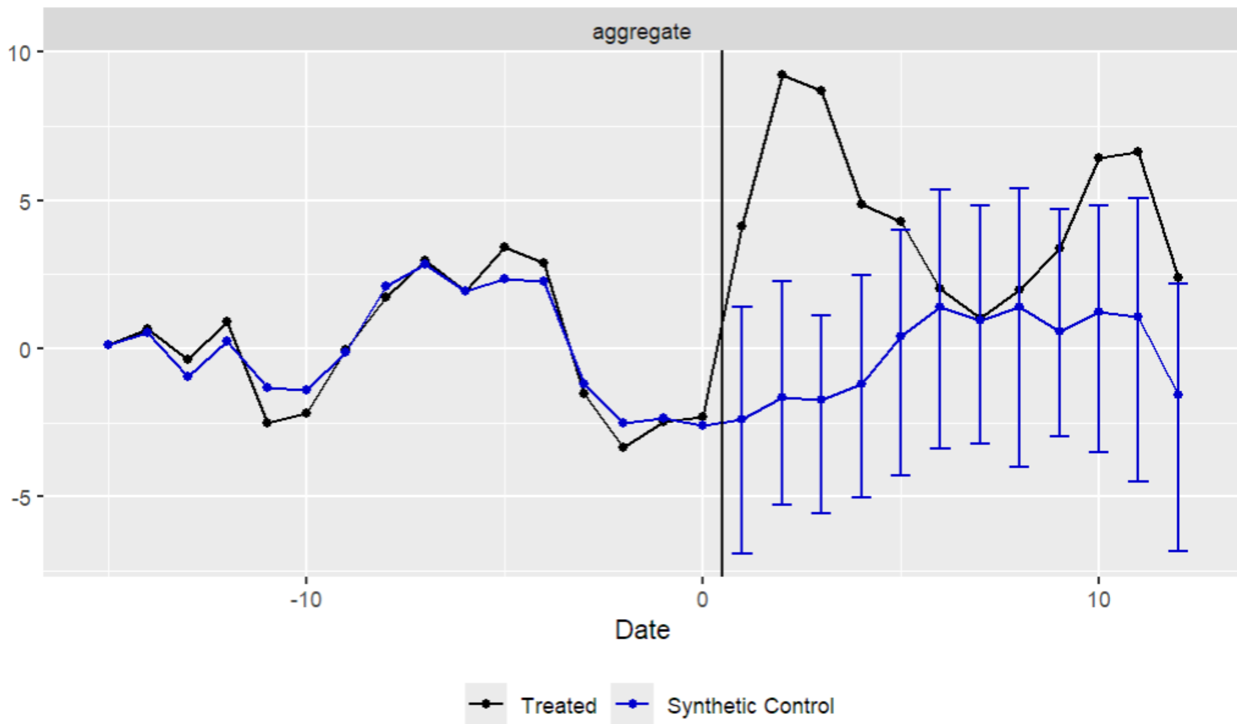

Figure 13: Average Christchurch region synthetic and actual normalised consents per thousand residents, with prediction intervals

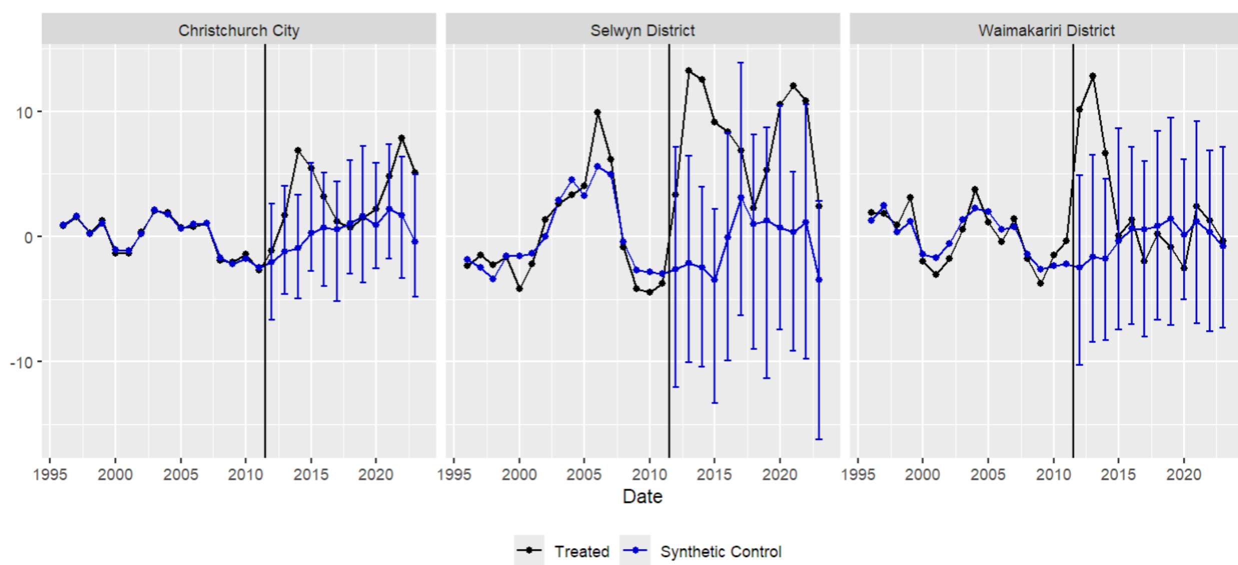

Figure 14: Christchurch region synthetic and actual normalised consents per thousand residents, with prediction intervals

C Additional Results: Expanded Set of Matching Variables

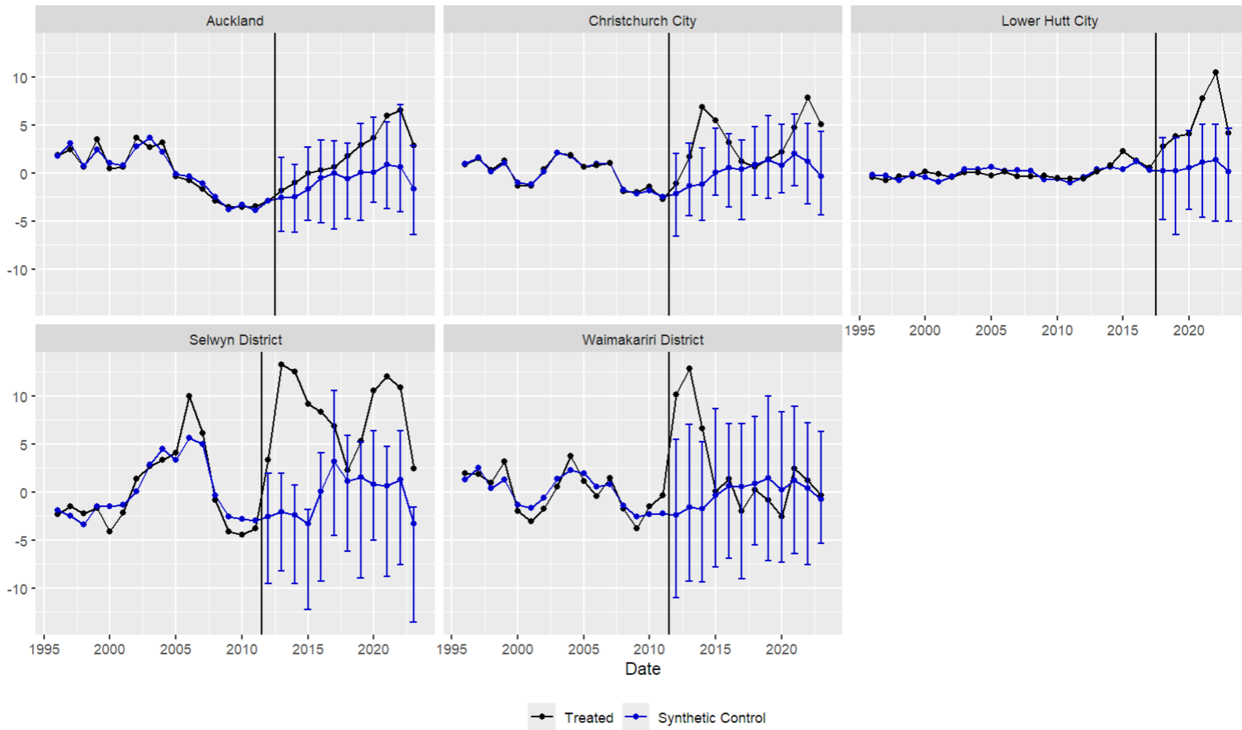

Figure 15: Individual synthetic and actual normalised consents per thousand residents, expanded Specification, with prediction intervals

Donor unit	**Auckland**	**Lower Hutt**	**Christchurch**	**Waimakariri**	**Selwyn**
Carterton District	-	-	-	-	0.448
Gisborne District	-	-	-	0.116	-
Hamilton City	-	-	0.362	0.305	-
Hastings District	-	-	0.035	-	-
Kapiti Coast District	0.200	-	-	-	-
Mackenzie District	-	-	0.002	0.067	0.092
Masterton District	0.149	-	-	-	-
Matamata-Piako District	-	-	-	-	0.097
Nelson City	-	-	0.046	-	-
New Plymouth District	-	-	0.128	-	-
Opitiki District	-	0.093	0.132	0.212	-
Queenstown-Lakes District	0.023	-	-	-	0.060
South-Waikato District	-	0.019	-	-	-
Stratford District	-	-	-	0.106	-
Tasman District	0.435	0.067	0.075	-	-
Taupo District	-	-	0.029	-	-
Tauranga City	0.192	-	0.081	0.194	-
Thames-Coromandel District	-	-	-	-	0.303
Timaru District	-	0.605	-	-	-
Waipa District	-	0.069	0.110	-	-
Wairoa District	-	0.147	-	-	-

Table 6: Expanded model specification donors

D Additional Results: LASSO Constraints

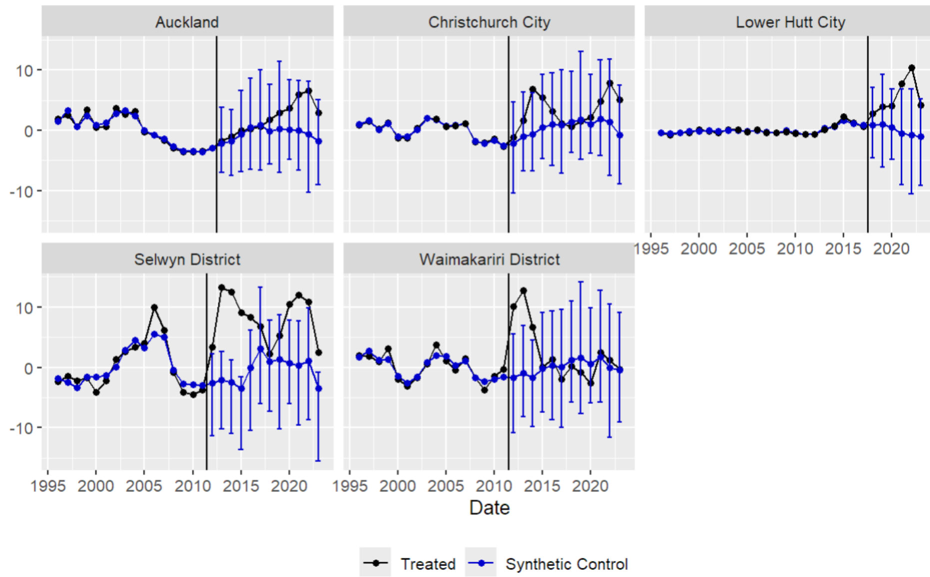

Figure 16: Individual synthetic and actual normalised consents per thousand residents, LASSO constraints, with prediction intervals

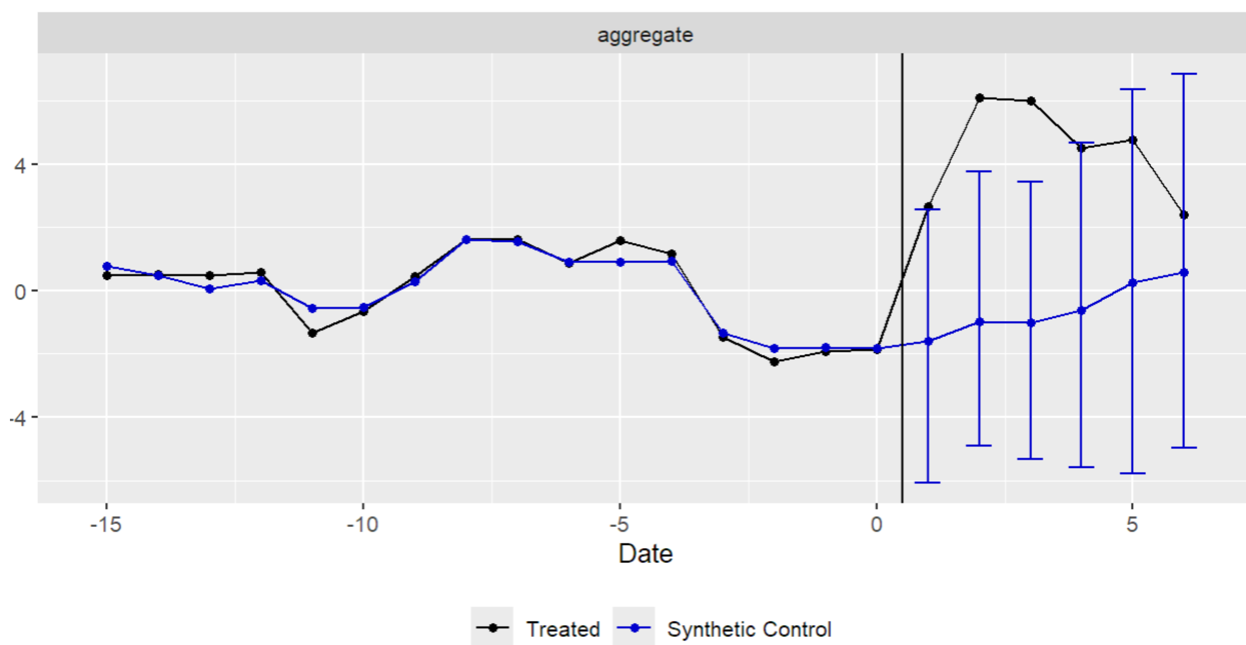

Figure 17: Average synthetic and actual normalised consents per thousand residents, LASSO constraints, with prediction intervals

E 'Leave One Out' Robustness Check Results

Donor unit	Auckland	Lower Hutt	Christchurch	Waimakariri	Selwyn
Central Otago District	-	-	-	-	0.095
Dunedin City	-	0.043	0.035	-	-
Gisborne District	-	-	-	0.131	-
Gore District	-	0.225	-	-	-
Hastings District	-	-	0.219	0.121	-
Hastings District	-	-	-	0.026	-
Kapiti Coast District	-	-	-	-	0.185
Kapiti Coast District	0.168	-	0.015	-	-
Mackenzie District	-	-	0.004	0.085	0.054
Matamata-Piako District	-	-	-	-	0.473
Nelson City	-	-	0.106	-	-
Opitiki District	0.296	0.043	0.133	0.301	-
Queenstown-Lakes District	0.107	-	-	-	0.104
South-Waikato District	-	0.023	-	-	-
Stratford District	-	-	-	0.166	-
Taupo District	-	-	0.007	-	-
Tauranga City	0.232	-	0.156	0.291	-
Thames-Coromandel District	-	-	-	-	0.089
Waipa District	-	0.257	0.244	-	-
Wairoa District	0.197	0.432	-	-	-

Table 7: 'Leave one out' check donors

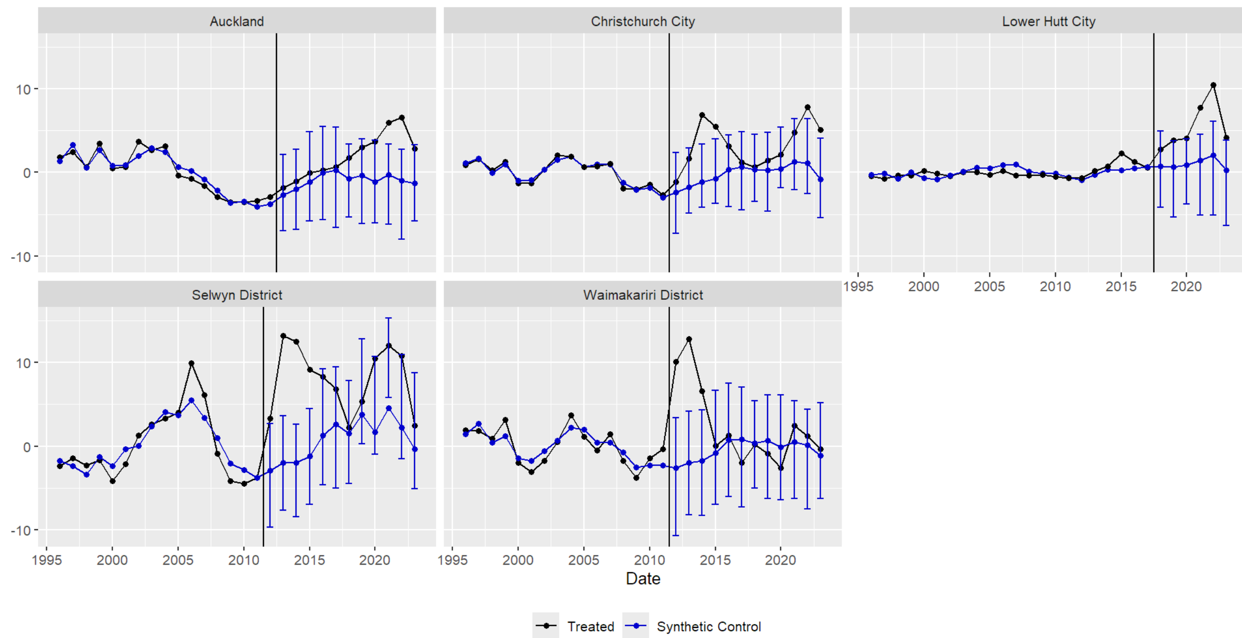

Figure 18: Individual synthetic and actual normalised consents per thousand residents, with prediction intervals

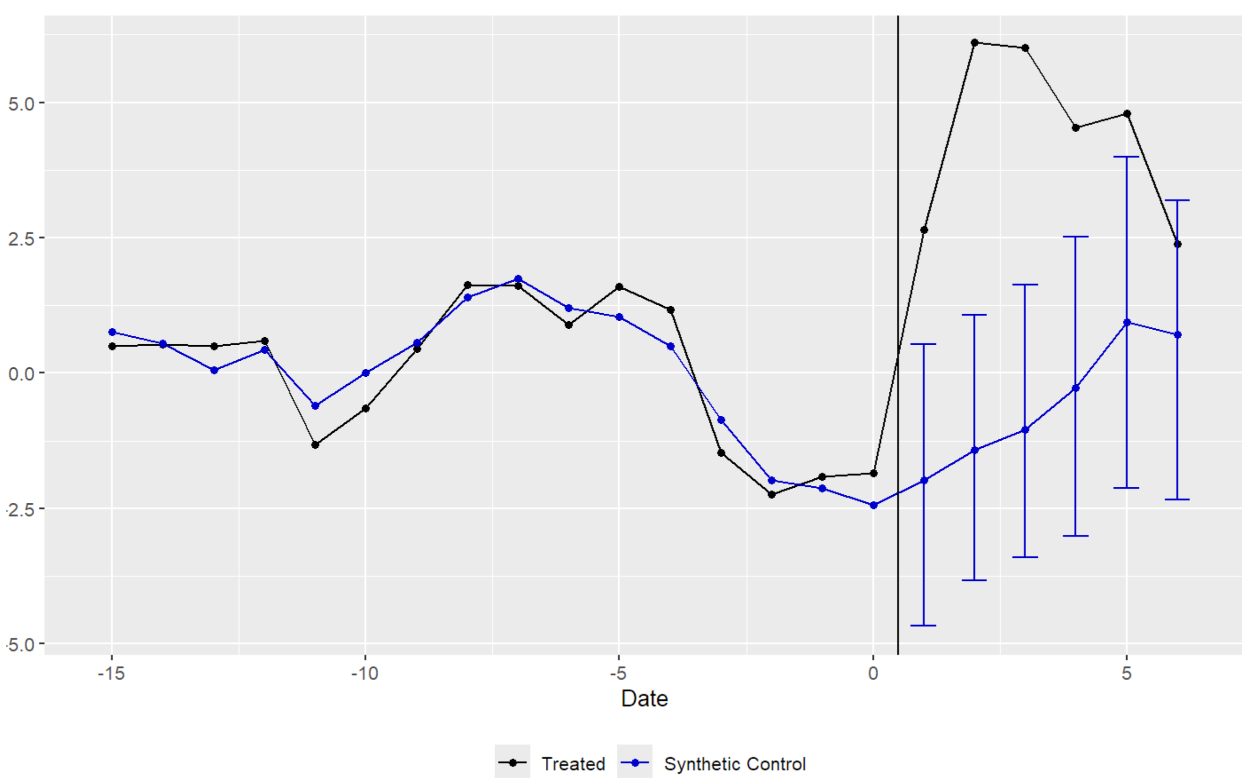

Figure 19: Average synthetic and actual normalised consents per thousand residents, with prediction intervals

F Error Formulae

The following formulae correspond to the in-sample and out-of-sample errors under the Cattaneo et al. (2023) framework.

$$\text{InErr}(\tau_{ik}) = -\mathbf{Y}'_{\mathcal{N}(T_i+k)}\left(\hat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}\right),$$

$$\text{InErr}(\tau_{i}) = -\frac{1}{T-T_i+1}\sum_{k=0}^{T-T_i}\mathbf{Y}'_{\mathcal{N}(T_i+k)}\left(\hat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}\right),$$

$$\text{InErr}(\text{TQ}_k) = \frac{-1}{Q}\sum_{i:T_i\in Q}\mathbf{Y}'_{\mathcal{N}(T_i+k)}\left(\hat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}\right),$$

$$\text{InErr}(\tau.) = \frac{-1}{LJ_1}\sum_{k=1}^{L}\sum_{i\in\mathcal{E}}\mathbf{Y}'_{\mathcal{N}(T_i+k)}\left(\hat{\mathbf{w}}^{[i]} - \mathbf{w}_0^{[i]}\right),$$

$$\text{OutErr}(\tau_{ik}) = u_{i(T_i+k)},$$

$$\text{OutErr}(\tau_{i\cdot}) = \frac{1}{T-T_i+1}\sum_{k=0}^{T-T_i}u_{i(T_i+k)},$$

$$\text{OutErr}(\text{TQ}_k) = \frac{1}{Q}\sum_{i:T_i\in Q}u_{i(T_i+k)},$$

$$\text{OutErr}(\tau_{\cdot\cdot}) = \frac{1}{LJ_1}\sum_{k=1}^{L}\sum_{i\in\mathcal{E}}u_{i(T_i+k)},$$

where $\mathcal{N} = \{i : T_i = \infty\}$ represents the group of "never-reformed" territorial authorities,

$\mathcal{E} = \{i : T_i < \infty\}$ represents the group of treated units (with number of elements $J_1 = |\mathcal{E}|$), L is the number of years after policy adoption, and $Q \subseteq \{1, 2, \ldots, T\}$ is a set of adoption times.

References

Abadie, A. (2021), 'Using synthetic controls: Feasibility, data requirements, and methodological aspects', *Journal of Economic Literature* **59**(2), 391–425.

Abadie, A., Diamond, A. & Hainmueller, J. (2010), 'Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program', *Journal of the American Statistical Association* **105**(490), 493–505.

Abadie, A., Diamond, A. & Hainmueller, J. (2015), 'Comparative politics and the synthetic control method', *American Journal of Political Science* **59**(2), 495–510.

Abadie, A. & Gardeazabal, J. (2003), 'The economic costs of conflict: A case study of the Basque country', *American Economic Review* **93**(1), 113–132.

Athey, S. & Imbens, G. W. (2022), 'Design-based analysis in difference-in-differences settings with staggered adoption', *Journal of Econometrics* **226**(1), 62–79.

Auckland Council (2013), The Proposed Auckland Unitary Plan, Technical report, Auckland Council. Accessed: 2024-10-28.

Ben-Michael, E., Feller, A. & Rothstein, J. (2021), 'The augmented synthetic control method', *Journal of the American Statistical Association* **116**(536), 1789–1803.

Brunsdon, N. (2019), 'Christchurch goes up and out', *Build Magazine* **172**, 73–74. Accessed: 2024-10-28.

Brunskill, D. (2023), 'Christopher Luxon's National Party has scrapped a bi-partisan housing agreement and will instead allow cities to choose where and how to build 30-years of growth', *Interest.co.nz* . Accessed: 2024-10-30.

Büchler, S. & Lutz, E. C. (2021), The local effects of relaxing land use regulation on housing supply and rents, Technical Report 21/18, MIT Center for Real Estate.

Canterbury Earthquake Recovery Authority (CERA) (2015), Land Use Recovery Plan monitoring report, Technical report, Canterbury Earthquake Recovery Authority.

Cattaneo, M. D., Feng, Y., Palomba, F. & Titiunik, R. (2022), 'scpi: Uncertainty quantification for synthetic control methods', *arXiv.org* .

Cattaneo, M. D., Feng, Y., Palomba, F. & Titiunik, R. (2023), 'Uncertainty quantification in synthetic controls with staggered treatment adoption', *arXiv.org* .

Cattaneo, M. D., Feng, Y. & Titiunik, R. (2021), 'Prediction intervals for synthetic control methods', *Journal of the American Statistical Association* **116**(536), 1865–1880.

Cheung, K. S., Monkkonen, P. & Yiu, C. Y. (2024), 'The heterogeneous impacts of widespread upzoning: Lessons from Auckland, New Zealand', *Urban Studies* **61**(5), 943–967.

Christchurch City Council (2015), 'Christchurch District Plan'. Accessed: 2024-10-30.

Christchurch City Council (2024), 'Proposed Plan Change 14', Available: `https://ccc.govt.nz/the-council/plans-strategies-policies-and-bylaws/plans/christchurch-district-plan/changes-to-the-district-plan/proposed-changes-to-the-district-plan/pc14/`. Accessed: 2024-10-30.

Colbert, J., Sila-Nowicka, K. & Yao, J. (2022), 'Driving forces of population change following the Canterbury earthquake sequence, New Zealand: A multiscale geographically weighted regression approach', *Population, Space and Place* **28**(8).

Dong, H. (2021), 'Exploring the impacts of zoning and upzoning on housing development:

A quasi-experimental analysis at the parcel level', *Journal of Planning Education and Research* **44**.

Doudchenko, N. & Imbens, G. W. (2016), Balancing, regression, difference-in-differences and synthetic control methods: A synthesis, Working Paper 22791, National Bureau of Economic Research.

Duguid, J. & Chan, C. (2014), Auckland's First Unitary Plan: Creativity, Innovation and Opportunity, Technical report, Auckland Council. Accessed: 2024-10-28.

Environment Canterbury Regional Council (2013), 'Canterbury Regional Policy Statement 2013', Available: https://ccc.govt.nz/the-council/plans-strategies-policies-and-bylaws/plans/christchurch-district-plan/changes-to-the-district-plan/proposed-changes-to-the-district-plan/pc14/. Revised December 2013.

Ferman, B. & Pinto, C. (2021), 'Synthetic controls with imperfect pretreatment fit', *Quantitative Economics* **12**(4), 1197–1221.

Freedman, S. M., Hollingsworth, A., Simon, K. I., Wing, C. & Yozwiak, M. (2023), Designing difference in difference studies with staggered treatment adoption: Key concepts and practical guidelines, Working Paper 31842, National Bureau of Economic Research.

Gabbe, C. J. (2019), 'Changing residential land use regulations to address high housing prices: Evidence from Los Angeles', *Journal of the American Planning Association* **85**(2), 152–168.

Greater Wellington Regional Council (2020), 2019/20 Annual Monitoring Report on the Regional Land Transport Plan, Monitoring Report 2020-63, Greater Wellington Regional Council.

Greenaway-McGrevy, R. (2023), 'Can zoning reform increase housing construction? Evidence from Auckland', *Economic Policy Centre* **17**.

Greenaway-McGrevy, R. & Jones, J. A. (2023), 'Can zoning reform change urban development patterns? Evidence from Auckland'.

Hutt City Council (2017), 'Proposed District Plan Change 43: Residential and Suburban Mixed Use Volume 2'. Accessed: 2024-10-27.

Hutt City Council (2022*a*), 'Implementing Government Requirements for Housing Intensification', Available: https://www.huttcity.govt.nz/council/district-plan/district-plan-changes/completed/implementing-government-requirements-for-housing-intensification. Accessed: 2024-10-29.

Hutt City Council (2022*b*), 'Regional Housing and Business Development Capacity Assessment – Housing Update, May 2022'. Accessed: 2024-10-27.

LaCour-Little, M. & Staer, A. (2016), 'Earthquakes and price discovery in the housing market: Evidence from New Zealand', *International Real Estate Review* **19**(4), 493–513.

Maltman, M. & Greenaway-McGrevy, R. (2024), 'Going it alone: The impact of upzoning on housing construction in Lower Hutt', *Economic Policy Centre* **18**.

Manaaki Whenua Landcare Research (2021), 'Urban Land Cover Indicators', Available: https://www.stats.govt.nz/indicators/urban-land-cover/. Accessed: 2024-10-27.

McClelland, R. D. & Gault, S. (2017), The Synthetic Control Method as a Tool to Understand State Policy, Technical report, Urban Institute.

Ministry for the Environment and Ministry of Business, Innovation and Employment (2016), National Policy Statement on Urban Development Capacity, National policy statement, New Zealand Government, Wellington, New Zealand. Released under the Resource Management Act 1991.

Ministry of Business, Innovation and Employment (2024*a*), 'Building and construction sector trends – annual report 2023', Available: https://www.mbie.govt.nz/.

Ministry of Business, Innovation and Employment (2024*b*), 'Building system reforms - granny flats', Available: https://www.building.govt.nz/getting-started/building-system-reforms/granny-flats.

Ministry of Housing and Urban Development (2024), 'Urban Development', Available: https://www.hud.govt.nz/urban-development/national-policy-statement-on-urban-development-nps-ud/. Accessed: 2024-10-25.

MRCagney (2022), 'Development Insights: Lower Hutt Housing and the Removal of Parking Minimums', Available: https://www.mrcagney.com/about/blog/development-insights-lower-hutt-housing-and-the-removal-of-parking-minimums/. Accessed: 2024-10-27.

New Zealand Infrastructure Commission (2022), The decline of housing supply in New Zealand: Why it happened and how to reverse it, Technical report, New Zealand Infrastructure Commission / Te Waihanga, Wellington.

New Zealand Productivity Commission (2015), Using land for housing, Technical report, New Zealand Productivity Commission. Accessed: 2024-10-24.

Panel, I. H. (2015), 'Christchurch Replacement District Plan: Decision 10 – Residential (Part)', Available: https://resources.ccc.govt.nz/files/policiesreportsstrategies/dpr_ihpdecision10_residential_stage_one.PDF, note=Accessed: 2024-10-30.

Pickett, R. E. M., Hill, J. & Cowan, S. K. (2022), 'The Myths of Synthetic Control: Recommendations for Practice'.

Robinson, S. (2013), 'Christchurch housing 'worst it's ever been'', *NZ Herald* .

Scoop (2016), ‘Wellington house prices rising faster than Auckland, up 21 per cent in 12 months’, Available: https://wellington.scoop.co.nz/?p=93127. Accessed: 2024-10-28.

Stats NZ (2020), ‘Regional council 2020 clipped (generalised)’, Available: https://datafinder.stats.govt.nz/layer/104253-regional-council-2020-clipped-generalised/. Accessed: 2024-10-27.

Stats NZ (2023), ‘Territorial authority 2023 (generalised)’, Available: https://datafinder.stats.govt.nz/layer/111194-territorial-authority-2023-generalised/. Stats NZ Geographic Data Service.

Stats NZ (2024*a*), ‘Aotearoa Data Explorer - Subnational Population Estimates’. Accessed: 2024-10-27.

Stats NZ (2024*b*), ‘Infoshare - building consents issued’, Available: https://infoshare.stats.govt.nz/ViewTable.aspx?pxID=a8d7d702-20bf-4f65-93e3-0eb3def4298d. Accessed: 2024-11-09.

Survey, U. S. G. (2011), ‘Christchurch earthquake - region info’, Available: https://earthquake.usgs.gov/earthquakes/eventpage/usp000huvq/region-info. Accessed: 2024-10-28.

Theelen, M. (2014), ‘Strategic Directions Hearing: Evidence in Chief’, Available: https://chchplan.ihp.govt.nz/wp-content/uploads/2015/03/Strategic-Directions-Hearing-Evidence-in-Chief-Michael-Theelen-Pr.pdf. Accessed: 2024-10-28.

Wellington Regional Leadership Committee (2023), Housing and Business Development Capacity Assessment: Wairarapa-Wellington-Horowhenua, Capacity assessment report, Wellington Regional Growth Framework, Wellington, New Zealand.

West, E. (2024), 'Up-zoning new zealand: The localisation of a globally mobile policy idea', *Urban and Spatial Economics Research Hub* **3**.

Wetzstein, S. (2017), 'The global urban housing affordability crisis', *Urban Studies* **54**(14), 3159–3177.