The Impact of Upzoning on Housing Construction in Auckland[*]

Ryan Greenaway-McGrevy[†]   Peter C. B. Phillips[‡]

May 2022

Abstract

There is a growing debate about whether upzoning is an effective policy response to housing shortages and unaffordable housing. This paper provides empirical evidence to further inform debate by examining the various impacts of recently implemented zoning reforms on housing construction in Auckland, the largest metropolitan area in New Zealand. In 2016, the city upzoned approximately three quarters of its inner suburban land to facilitate construction of more intensive housing. We use a quasi-experimental approach to analyze the short-run impacts of the reform on construction, allowing for potential shifts in construction from non-upzoned to upzoned areas (negative spillovers) that would, if unaccounted for, lead to an overestimation of treatment effects. We find strong evidence that upzoning stimulated construction. Treatment effects remain statistically significant even under implausibly large spillovers that would necessitate a six-fold increase in the trend rate of construction in control areas under the counterfactual of no-upzoning. Our findings support the argument that upzoning can stimulate housing supply and suggest that further work to identify factors that mediate the efficacy of upzoning in achieving wider objectives of the policy would assist policymakers in the design of zoning reforms in the future.

Keywords: Upzoning, Land Use Regulations, Redevelopment, Housing Construction, Housing Affordability.

JEL Classification Codes: R14, R31, R52

[†]University of Auckland. Corresponding author. Postal address: The University of Auckland, Private Bag 92019 Auckland 1142, New Zealand. Email: r.mcgrevy@auckland.ac.nz.

[‡]Yale University, University of Auckland, Singapore Management University, and University of Southampton. Email: peter.phillips@yale.edu.

1

1 Introduction

Housing has become prohibitively expensive in many of the world's major cities, precipitating serious and widespread housing affordability crises (Wetzstein, 2017). A growing coalition of researchers argue that part of the solution is to 'upzone' cities by relaxing land use regulations (LURs) to allow construction of more intensive housing, such as townhouses, terraced housing and apartment buildings (Glaeser and Gyourko, 2003; Mayer and Somerville, 2000; Freeman and Schuetz, 2017; Manville, Monkkonen, and Lens, 2019). Policymakers have begun to listen to these supply-side solutions and, in response, several metropolitan and gubernatorial authorities have pursued upzoning strategies in recent years (National Public Radio, 2019).

These policy reforms are underpinned by the argument that LURs increase house prices by restricting housing supply (Thorson, 1997; Mayer and Somerville, 2000; Quigley and Raphael, 2006; Ihlanfeldt, 2007; Saiz, 2010; Zabel and Dalton, 2011; Kahn, Vaughn, and Zasloff, 2010; Jackson, 2016). Relaxing those regulations through upzoning, it is argued, enables new more intensive development, thereby increasing housing supply and putting downward pressure on prices. However, these arguments are not universally accepted and many commentators remain sceptical of the capacity for these market-led policies to deliver affordable and inclusive housing (Rodríguez-Pose and Storper, 2020; Wetzstein, 2021). Instead it is suggested that government intervention is needed to tackle the problem through means such as state-led construction (Favilukis, Stijn, Nieuwerburgh, Mabille, and Nieuwerburgh, 2019; Wetzstein, 2021), the repurposing of public space (Wetzstein, 2021; Freemark, 2021), and policies that limit demand (Wetzstein, 2021).

Unfortunately our understanding of the manifold impact of upzoning is presently limited by an acute lack of empirical research on the subject (Schill, 2005; Freeman and Schuetz, 2017; Freemark, 2019). Only a handful of studies have offered empirical evidence of the effects of a relaxation of LURs (Atkinson-Palombo, 2010; Freemark, 2019) and these have concentrated on small-scale policy changes involving transit-oriented rezoning, not the sweeping policy reforms currently being implemented in many US cities. The limited empirical work that is available has findings that often contravene the outcomes anticipated by proponents of supply-side regulatory reforms. For example, Freemark (2019) found that transit-oriented upzoning in Chicago failed to encourage construction, calling into question the fundamental premise of the supply-side argument (Rodríguez-Pose and Storper, 2020). More recently, Limb and Murray (2021) conclude that rezoning in Brisbane, Australia, generated no significant increase in housing construction.

The lack of empirical evidence on the effects of large-scale upzoning is largely due to the fact that, until very recently, no city has systematically upzoned large shares of land as a mechanism to promote affordability (Freeman and Schuetz, 2017, p. 229). In 2016, however, the city of Auckland, New Zealand, implemented large-scale zoning reforms under the Auckland Unitary Plan (AUP). Motivated by housing affordability concerns (Auckland Unitary Plan Independent Hearings Panel, 2016), the plan removed single dwelling restrictions on three-quarters of residential land and trebled the number of dwellings that could be built within the residential areas (Greenaway-McGrevy, Pacheco, and Sorensen, 2021), providing researchers with a unique opportunity to study large-scale

upzoning reform of the kind that was hitherto lacking.

The present paper examines impact of upzoning on housing construction in Auckland to provide further evidence on whether zoning reforms can fulfil the fundamental premise of the supply-side policy response, namely, that upzoning increases housing supply. A quasi-experimental framework is adopted that exploits geographic variation in the incidence of upzoning to estimate causal effects through the comparison of outcomes in upzoned areas with outcomes in non-upzoned areas. Our dataset consists of geocoded building consents (or 'permits' in North American parlance) that are matched to planning maps that detail the geographic incidence and intensity of upzoning. Because the empirical design exploits temporal changes in zoning rules via a policy intervention, it has the capacity to mitigate many of the concerns stemming from the endogeneity of regulations that afflict studies which rely only on spatial variation in LURs (Gyourko and Molloy, 2015). The approach is concordant with other work where changes in the geographic variation in zoning has been used in quasi-experimental designs to examine casual impacts (Chakraborty, Knaap, Nguyen, and Shin, 2010; Cunningham, 2006; Kahn, Vaughn, and Zasloff, 2010; Thorson, 1997; Zhou, Mcmillen, and Mcdonald, 2008).

Our empirical strategy pays particular attention to the possibility of negative spillover effects that can lead to upward bias in estimated treatment effects. One potential consequence of upzoning is that it can reallocate construction that would have otherwise occurred in non-upzoned areas to upzoned areas.[1] This violation of the stable unit treatment value assumption (SUTVA) in the Neyman-Rubin causal framework means that a simple comparison of outcomes in treatment and control groups can generate a confounding overstatement of the treatment effect. To address this problem, we adapt the set identification approach suggested by Rambachan and Roth (2022) (hereafter 'RR') for remediating violations of the standard parallel trends (or 'pre-trends') assumption that is required under the difference-in-differences (DID) framework. RR extrapolate pre-treatment trends to generate a set of counterfactual outcomes in the treatment group. In the present paper we repurpose this strategy by using pre-treatment trends in the *control* group to extrapolate a set of counterfactual outcomes that are used to bound the magnitude of the spillover effect. The intuition underpinning both the RR strategy and our strategy is the idea that observed trends immediately prior to the policy intervention are informative of the counterfactual scenario. Adapting the RR method to our application yields a confidence set of treatment effects that is robust to spillover effects and amenable to inference.

The empirical findings using this methodology reveal strong statistical evidence that upzoning increased housing construction. Our preferred model specification shows a statistically significant increase in consents even under counterfactual sets that span approximately six times the extrapolated linear trend in the control group. For example, a linear trend fitted to pre-treatment observations in the control group implies that 1,344 additional dwellings would have been built in non-upzoned areas in 2021 compared to 2015 under the counterfactual of no upzoning. We find that the estimated treatment effect for 2021 remains statistically significant even when we permit

---

[1]Monocentric models of urban development predict a reallocation of construction from the city fringes to the inner suburbs when height restrictions on residential construction are relaxed. See Bertaud and Brueckner (2005).

an additional 7,744 dwellings in the counterfactual set of outcomes. Put differently, counterfactual scenarios that imply an approximate six-fold (since $7744/1344 \simeq 6$) increase in consents over the pre-treatment trend would be needed in order for the estimated treatment effects to become statistically insignificant. There is no policy change concurrent to the upzoning policy that could plausibly generate such a substantive increase in construction.

We also use the extrapolated counterfactual trend in control areas to generate a spillover-robust estimate of the number of additional dwellings enabled through upzoning. To do so, we restrict the counterfactual set to the extrapolated linear trend, so that the set collapses to a point. This approach implies that 26,903 additional dwellings were consented over the five years following upzoning, corresponding to approximately 5.07% of the dwelling stock of the Auckland region.[2]

Our results have implications for ongoing debates about the efficacy of upzoning. In particular, the findings support the view that large-scale upzoning can encourage construction. This is particularly important in the light of recent work by Freemark (2019) and Limb and Murray (2021), who find that zoning changes had minimal impact on housing construction in Chicago and Brisbane, respectively. Further work examining potential mediating factors that enabled increased construction will hopefully help explain why the policy was more effective in Auckland, and assist policymakers in tailoring rezoning and housing policies to facilitate construction elsewhere.

The remainder of the paper is organized as follows. Section two provides the background institutional context and timeline of the key events in the city of Auckland and section three describes the dataset used in our empirical work. Section four presents the empirical DID model. Section five describes and applies our methodology for dealing with potential spillover effects, and presents spillover-robust estimates of treatment effects. Section six concludes.

2   Institutional Background

This section provides some background demographic and administrative features of Auckland city with information concerning relevant policies and processes preceding the relaxation of land use regulations under the Auckland Unitary Plan.

Auckland is the largest city in New Zealand with a population of approximately 1.5 million within the greater metropolitan region (as of 2017). Since 2010, the entire metropolitan area, as well as several towns, populated islands, and a large amount of the rural land beyond the fringes of its outermost suburbs, has been under the jurisdiction of a single local government, the Auckland Council. Centred on a long isthmus of land between two harbours, this jurisdiction extends over 4,894 km^2 of land area.

In March 2013, the Auckland Council announced the 'draft' version of the Auckland Unitary Plan. The draft version of the plan went through several rounds of consultations, reviews and revisions before the final version became operational on 15 November 2016. Each version of the AUP contained new LURs that would potentially change restrictions on the extent of site development,

---

[2]Unfortunately we do not have precise measures of dwellings demolished when properties are redeveloped. This is because demolition permits are only required for buildings that are less than three storeys.

depending on the site location. In most areas these LURs were relaxed in order to enable residential intensification and greater population density, including multi-family housing such as terraced housing and apartments. These proposed changes could be viewed online, so that any interested member of the public could observe the specific LURs proposed for a given parcel of land. This meant that it was relatively simple for developers to observe the new land use regulations and to commence planning prior to the policy becoming operational.

The amount of development permitted on a given site is restricted by the residential planning zone in which the site is located. In this study we focus on four zones, listed in declining levels of permissible site development: Terrace Housing and Apartments (THA); Mixed Housing Urban (MHU); Mixed Housing Suburban (MHS); and Single House (SH). Thus THA permits the most site development, and SH permits the least. Table 1 summarizes the various LURs for each of the four residential zones considered. These regulations include site coverage ratios, minimum lot sizes for new subdivisions, and height restrictions, among others. For example, between five to seven storeys and a maximum site coverage ratio of 50% is permitted in THA, whereas only 2 storeys and a coverage ratio of 35% is permitted in SH.

[Insert Table 1 here]

Our empirical design treats the introduction of the AUP as a quasi-natural experiment. The SH zone is used as the control, while the MHS, MHU and THA zones are used as the treatment groups. Most of our analysis will be based on aggregating these three zones into a single treatment group. However, towards the conclusion of our empirical work we focus on each zone separately.

3 Dataset

Our dataset is based on annual building permits issued for new dwelling units by the Auckland Council from 2010 to 2021.[3] The permits include the number of dwellings consented. Each observation includes the longitude and latitude of the parcel, which have been used to map each permit to its corresponding zone under the Unitary Plan.[4]

We use local board areas as suburbs of Auckland, including territorial authority subdivisions. Figure 1 illustrates the local areas. Local areas are obtained by the intersection of the 2016 Community Board Areas and Territorial Authority subdivisions. In addition, we delineate regional areas for Kumeu-Riverhead-Waiamauku and Titirangi by amalgamating statistical areas in these suburbs.[5]

The four residential zones we use in our sample (SH, MHS, MHU and THA) comprise the vast majority of residential land in Auckland. Figure 2 depicts the geographic distribution of the zones.

---

[3]Permits for extensions to existing dwellings are not included in our analysis.

[4]A small proportion of coordinates are located marginally outside of parcels on roads. These observations are omitted. We anticipate repairing these observations in future iterations of the paper.

[5]We use 2018 Statistical Area 2 units for this exercise.

For clarity we zoom in on the central urban core of Auckland. Evidently MHS covers the most area, closely followed by MHU. SH is predominantly located either very close to the Auckland city central business district (CBD) or at the outskirts of the city. THA covers the least amount of area. Within what we refer to as the suburban core of Auckland (the yellow area in Figure 1), the THA, MHU and MHS zones collectively comprise 77.6% of residential land.

[**Insert Figure 1 here**]

[**Insert Figure 2 here**]

Care is taken to remove construction activity generated by other government policies during our sample period. In particular, consents issued under 'Special Housing Area' (SpHA) authority are omitted from our sample. These designated areas came with incentives for developers to provide affordable housing by offering accelerated consenting processing, and would be considered similar to an 'inclusionary zoning' policy in the United States. Once it came into effect, the Unitary Plan superseded the SpHA development incentives. We exclude permits issued within SpHAs prior to 2017 since this policy intervention preceded upzoning. After filtering out consents issued under SpHAs and zones that fall outside of the four primary zones of interest, the set of local areas included shrinks to 32. (The *Waiheke Island* local area has no SH, MHS, MHU or THA zoning.)

Figure 3 exhibits the aggregate dwellings consented in non-upzoned areas (SH) compared to upzoned areas (THA, MHU and MHS) over the 2010 to 2021 period. We also decompose permits into attached and detached dwellings. There is a clear increase in the number of dwellings consented in upzoned areas after the policy is implemented (from 2017 onward). The number of attached dwellings consented per year in upzoned areas increases from under 1,000 in 2016 to near 10,000 by 2021 – more than a tenfold increase. Over the same period, detached housing increases from just over 2,000 consents per year to approximately 4,500. By 2019, there were more attached dwellings being consented than detached, consistent with the upzoning goal of incentivising more capital intensive structures. In addition, there is an interesting fall in detached dwelling consents between 2019 and 2021 in upzoned areas.

Prior to treatment, consents in subsequently upzoned areas consistently exceeded consents in non-upzoned areas by between 1400 to 1900 consents annually. This difference remains remarkably consistent even as consents in non-upzoned areas more than doubled between 2011 and 2015. These patterns are consistent with modelling permits in levels in the DID framework, since it implies an absolute difference in the level of the two series under the counterfactual (Kahn-Lang and Lang, 2020).

[**Insert Figure 3 here**]

There is a notable decrease in the number of dwellings consented in the non-upzoned areas after 2015. Consents follow a steady upwards trend until 2015. Thereafter there is a discrete shift as consents trend downwards. This break in trend is statistically significant (refer to the Appendix for details). This outcome is consistent with a negative spillover effect, as construction that would otherwise have occurred in non-upzoned areas may have been relocated to upzoned areas as a result of the treatment. The outcome is mainly driven by a decline in detached dwellings, which lends further support to the negative spillover interpretation of the switch in trend from 2015.

Economic theories of urban development suggest that these spillover effects will be more apparent in locations distant from the city centre. The canonical Alonso-Muth-Mills (AMM) spatial equilibrium model of the monocentric city predicts that a relaxation of LURs will shift future construction from the outskirts of the city to the centre (Bertaud and Brueckner, 2005). To explore whether this is the case in Auckland, Figure 4 divides the sample into core and non-core areas of Auckland (see Figure 1 above for the geographic delineation of the core and non-core areas). Consistent with this prediction, the reduction in consents in non-upzoned areas (relative to the pre-treatment trend) from 2015 onwards appears to be driven more by a fall outside of the urban core. Consents in non-upzoned areas outside the core experienced rapid growth prior to 2015, increasing from 401 in 2010 to 1200 in 2015, thereafter declining to a local nadir of 772 consents in 2019. Consents in non-upzoned areas inside the core are comparatively flatter, although there is a mild increase above trend in 2014 and 2015. Consents increased from 434 in 2010 to 755 in 2015, thereafter declining back to 447 consents in 2016. In 2021, 533 consents were granted. Although the data suggest that the negative spillover effects are stronger outside of the urban core, our empirical strategy for addressing spillovers is not based on this prediction of the AMM model.

[**Insert Figure 4 here**]

Figure 5 breaks down the upzoned areas into the three constituent AUP zones. Despite having more restrictive constraints than MHU and THA, it is unsurprising that MHS accounts for most of the increase in consenting activity because it covers the most geographic area.

[**Insert Figure 5 here**]

4 Empirical Model and Results

Construction activity in the different residential zones is assembled by mapping individual construction consent data to residential zones using the longitude and latitude of the building site. Let $c_{i,j,t}$ denote the (log) number of dwellings built in zone j in suburb $i = 1, \ldots, n$ in period $t = -\underline{T}, \ldots, 0, \ldots, \overline{T}$, where $\underline{T}$ denotes the number of time series observations prior to the treatment, and $\overline{T}$ denotes the number of time series observations post-treatment. Treatment occurs in

period $t = 0$. We use $j = 0$ to indicate the control group (i.e, consents in non-upzoned areas) and $j = 1$ to signify the treatment group (consents in upzoned areas). The causal impact of upzoning is then estimated using a multi-period difference-in-differences (DID) specification of the form

$$c_{i,j,t} = \alpha_{i,j} + \sum_{s=-\underline{T},s\neq0}^{\overline{T}} \phi_s \mathbf{1}_{s=t} + \sum_{s=-\underline{T},s\neq0}^{\overline{T}} \beta_s \mathbf{1}_{s=t,j=1} + \varepsilon_{i,j,t} \tag{1}$$

where $\alpha_{i,j}$ are suburb-zone fixed effects, ϕ_s are time s specific fixed effects, and $\mathbf{1}_{s=t}$ are indicators for each time period except the period of implementation, which occurs at $t = 0$. $\mathbf{1}_{s=t,j=1}$ are indicators for each time period (except the period of policy implementation) interacted with the treatment effect indicator. Thus $\{\beta_s\}_{s=1}^{\overline{T}}$ represent the treatment effects over time of upzoning. The empirical estimates of these parameters capture the increase in construction activity in treatment areas relative to the control areas in each period after upzoning is implemented. Following convention, estimates of $\{\beta_s\}_{s=-\underline{T}}^{-1}$ will be used to assess whether any potential confounding pre-treatment trends are evident. The time specific fixed effects ϕ_s measure outcomes in the control group relative to the treatment period $t = 0$. They also capture common variation in consents across different zones and suburbs that is due to macroeconomic or city-wide shocks and policy changes.

Dwelling consents are modelled in levels. As discussed above, levels seem more appropriate given the observed pre-treatment trends in upzoned and non-upzoned areas, which differ over time by a near constant amount until the AUP is implemented. An additional benefit of modelling outcomes in levels is that it allows us to define counterfactual sets in terms of model parameters. By definition, spillovers are measured in levels. For example, consider a spillover that generates $\epsilon \in \mathbb{R}^+$ fewer consents in non-upzoned areas – and ϵ more consents in upzoned areas – in the first treatment period. The corrected treatment effect would be $\hat{\beta}_1 - \epsilon$ and the corrected period effect would be $\hat{\phi}_1 + \epsilon$. This direct mapping is lost or transformed if consents are instead modelled in logs or another non-linear transformation.

4.1 Selection of the Treatment Date

Throughout this study 2015 is used as the final year prior to treatment (i.e. the treatment date). There are three main reasons for this selection.

First, as discussed in Section 2, the AUP became operational in November 2016, meaning that the final month and a half of 2016 fall under the new regulations. Using 2015 as the treatment date ensures that any differential increase in consents between upzoned and non-upzoned areas over this final month and a half of 2016 are correctly identified as a treatment effect.

Second, this date yields a more conservative estimate of the treatment effect if (a) negative spillovers from non-upzoned to upzoned areas are present, and (b) developers begin to respond to new regulations prior to those regulations coming into effect. Policy interventions can begin to manifest prior to the policy change if agents are notified of the policy change in advance. This is likely in the case of the Unitary Plan because the first version of the plan, with clear notification of the intent for dwelling intensification, was released in 2013, more than three years before the final version became operational (see Section 2 above). Negative spillovers can manifest as a decrease

8

in dwelling consents in non-upzoned areas prior to 2016 if developers delay and shift planned construction to upzoned areas. The observed trend in consents in non-upzoned areas depicted in Figure 3 is consistent with a negative spillover effect beginning prior to the policy change. In fact, consents trend upwards until reaching a peak in 2015, a year before the policy change in late 2016, before trending downwards to a nadir in 2019.

Third, the most appropriate alternative treatment date, 2016 (the year the policy is enacted), results in larger estimated treatment effects because consents in non-upzoned areas are lower in 2016 than 2015. This is also the case for the set identified treatment effects, since the extrapolated linear trend that provides the basis for the counterfactual sets is flatter when the treatment date does not coincide with the peak in the control group. The result is a smaller spillover effect in the counterfactual set, and thus a larger treatment effect in the identified set, compared to that obtained with 2015 as the treatment date.

4.2   Results

Figure 6 shows the estimates of the coefficients alongside 95% confidence intervals (standard errors are clustered by suburb). Recall that we set the treatment to occur in 2016 since only the final month and a half of 2016 falls under the AUP. The top panel of the Figure displays results for all dwellings, while the middle and bottom panels display results for detached and attached dwellings, respectively.

Evidently there is no apparent trend in the estimated treatment effects prior to the treatment date in any of the three samples. We proceed under the assumption that there is no confounding variable generating a difference in consents between treatment and control areas prior to policy implementation. In other words, the parallel trends assumption appears to hold. After upzoning the estimated treatment effects are increasing over time. In 2021, six years after the policy was introduced, some 333.2 additional dwellings are built, on average, in upzoned areas compared to non-upzoned areas in each of the 32 local areas. (This would correspond to an additional 10,662 = (343.6 × 32 local areas) dwellings across the city.) These are mostly attached dwellings. The estimated treatment effect for 2021 is 270 for attached dwellings and 73.6 for detached dwellings.

[Insert Figure 6 here]

5   Spillover Effects

As discussed earlier, it is possible that the implementation of upzoning reallocated construction to upzoned areas that would have otherwise occurred in non-upzoned areas. This negative spillover effect would lead to an overstatement of the treatment effect, since some of the construction in upzoned areas would have occurred in non-upzoned areas under the counterfactual scenario.

The time series plots of dwelling consents in Figure 3 appear to be consistent with a spillover effect. There is a mild upward trend in dwellings built in non-upzoned areas until 2015, one year prior to policy implementation. Thereafter, there is a mild decrease until 2019. Meanwhile, consents in upzoned areas trend upwards from 2016 onwards. This pattern is consistent with a negative spillover effect that shifted construction from non-upzoned to upzoned areas as a result of the policy. The initial upward trend in the control group, and subsequent downturn from 2015 onwards, is most evident in the detached dwellings sample, which is also consistent with a negative spillover.

These patterns are also evident in the estimated period effects (i.e., $\left\{\hat{\phi}_s\right\}$), which are depicted in Figure 7. The coefficients measure outcomes (average number of dwelling consents) in the control group relative to the treatment period (2015). In the all dwellings sample, the coefficients trend upwards until 2015; thereafter they trend down until 2019. These coefficients form the basis for specifying the set of counterfactual scenarios.

[Insert Figure 7 here]

Extant techniques for addressing spatial spillover effects require the spillovers to be localised, in the sense that the magnitude of the spillover between geographically distant areas is assumed to be negligible (Clarke, 2017; Butts, 2021; Huber and Steinmayr, 2021). Such an assumption is unlikely to be tenable in our setting. Spatial models of urban development predict that upzoning will shift construction from the city fringes towards the city centre, and, as discussed above, the patterns in consenting activity illustrated in Figure 4 are consistent with this prediction. We therefore develop an approach to accommodating spillover effects that does not rely on spillovers dissipating with distance.

5.1 Set Identification of Treatment Effects under Spillovers

To account for spillover effects we adopt a confidence set identification approach by repurposing methods recently proposed in Rambachan and Roth (2022). First, we specify a set of plausible treatment effects based on observed pre-treatment trends in the control group. We then test whether the estimated treatment effects are significantly different from the set of counterfactuals.

Pre-treatment trends are frequently used to infer information about the counterfactual in DID frameworks. For example, RR propose extrapolating pre-treatment trends in estimated treatment effects to make inferences about counterfactual outcomes. We propose extrapolating trends in outcomes in the control group (i.e., the estimated period fixed effects) to learn about the counterfactual.

Figure 8 superimposes a linear trend on the outcomes in the control group. The linear trend is only fitted to the pre-treatment sample (2010–2015), and is fitted to pass through the coefficients in the first period (2010) and the treatment date (2015). Extrapolating a linear trend into the treatment period yields 42 additional dwellings permitted in the control group in the final period

of observation (2021) relative to the treatment date (2015). However, we actually observe 6.1 fewer consents (since $\hat{\phi}_{\overline{T}} = -6.1$) in the control group. Using the extrapolated trend as the counterfactual implies that the treatment effect in the final period is overstated by 48.1 (= 42 + 6.1) dwellings, or approximately 1,542 dwellings across $n = 32$ local areas.

[Insert Figure 8 here]

It is desirable to permit some margin for error when using pre-treatment trends to extrapolate a counterfactual scenario. This permits local deviations to potential nonlinearities in the trends. The approach suggested by Rambachan and Roth (2022) is to adopt a set of counterfactuals around the extrapolated trend. For example, in Figure 8, the space in between the dot-dashed pink lines permits a set of counterfactual scenarios in the control group, such that there is a margin of error of 50 dwellings in the final period of the analysis (2021). This means that the control group counterfactual for 2021 lies anywhere between 92 (= 42 + 50) and -8 (= 42 − 50) dwellings built relative to the treatment period (2015).

The technicalities of the set identification approach are now formalized.

5.1.1 Set Identification of Treatment Effects

Let $\beta = \left(\beta_{pre}', \beta_{post}' \right)'$, where $\beta_{pre} = \left(\beta_{-\underline{T}}, \ldots, \beta_{-1} \right)'$ and $\beta_{post} = \left(\beta_1, \ldots, \beta_{\overline{T}} \right)'$. Following RR, we partition $\beta_{post} = \delta_{post} + \tau_{post}$, where τ_{post} is the true treatment effect, and δ_{post} is the difference between the treatment and control groups under the counterfactual scenario. The conventional assumptions required for DID analysis ensure that $\delta_{post} = 0$ In the absence of spillover effects, $\delta_{post} = 0$. [<= inexact] The quantities δ_{post} and τ_{post} are unobserved and unidentified.

Example 1. Suppose that $\overline{T} = 1$ (one post-treatment period) and that upzoning only re-allocated construction from control to treatment areas, so that no additional construction was generated. Then $\beta_1 = \delta_1$ and $\tau_1 = 0$.

RR use β_{pre} to generate a set of possible counterfactual outcomes when the parallel trends assumption ($\beta_{pre} = 0$) does not hold. For example, in the three period DID case, they discuss setting $\delta_1 = -\beta_{-1} \pm M$ for some $M \in \mathbb{R}^+$. The intuition is that observed pre-treatment trends in the treatment group relative to the control group are informative of post-treatment trends under the counterfactual. The case where $M = 0$ imposes a linear extrapolation, which is highly restrictive. Permitting general $M \in \mathbb{R}^+$ allows for nonlinear patterns within a set of counterfactual scenarios.

In our application we want to account for negative spillovers that cause the estimated treatment effects to overstate the true treatment effects. To do so, we need to place bounds on the size of the spillover effect. In this regard, pre-trends in the treatment effect β_{pre} are uninformative.[6] Instead, we propose using pre-treatment trends in control group outcomes to bound the counterfactual

---

[6]However, our approach to modelling counterfactual scenarios based on pre-treatment trends in the control group could easily be extended to incorporate trends in *both* the treatment and the control group. Under such a scenario, β_{pre} is informative and would be used to bound δ_{post}.

outcomes in the treatment group relative to the control group. We therefore similarly define and partition $\phi = \left(\phi'_{pre}, \phi'_{post}\right)'$, where $\phi_{pre} = \left(\phi_{-\underline{T}}, \ldots, \phi_{-1}\right)'$ and $\phi_{post} = \left(\phi_1, \ldots, \phi_{\overline{T}}\right)'$.

Example 2. Suppose that $\overline{T} = \underline{T} = 1$ (so that we have three periods) and that $\phi_{-1} < 0$ (outcomes in the control group trend upwards prior to the treatment) and $\phi_1 < 0$ (outcomes in the control group trend downwards after the treatment). If control group outcomes remained on trend under the counterfactual, negative spillovers account for the observed downward deviation from trend after the policy is implemented. Then $\delta_1 = -\phi_{-1} - \phi_1$ and $\tau_1 = \beta_1 + \phi_{-1} + \phi_1$. In practice, we permit deviations from the linear trend assumption, such that $\delta_1 = -\phi_{-1} - \phi_1 \pm M$ for some $M \in \mathbb{R}^+$.

Once the set of counterfactuals is articulated, we can rely on the inferential architecture supplied by Rambachan and Roth (2022). In our application we select counterfactual sets that are convex and centrosymmetric, which ensures that fixed length confidence intervals (FLCI) are consistent. We provide an overview of the method for completeness. The remainder of this section is based on section 3.1 of RR.

Let $\theta = l'\tau_{post}$ be a linear combination of the treatment parameters of interest, where $l \in \mathbb{R}^{\overline{T}}$. For example, if we are interested in the treatment effect in the final period, $l = (0, \ldots, 0, 1)'$. Next, let $\hat{\lambda}_n$ be a relevant $m-$subvector of the estimate $\hat{\Lambda}_n = \left(\hat{\phi}'_n, \hat{\beta}'_n\right)'$, where $\hat{\lambda}_n \sim \mathcal{N}\left(\lambda, \Sigma_n\right)$. That is, there exists a full column rank $\left(\underline{T} + \overline{T}\right) \times m$ selection matrix $\mathbf{J}$ such that $\hat{\lambda}_n = \mathbf{J}'\hat{\Lambda}_n$. The choice of $\hat{\lambda}_n$ depends on both the parameter of interest θ and the counterfactual set – a specific example is given below. We similarly define $\lambda = \mathbf{J}'\Lambda$, where $\Lambda = (\phi', \beta')'$ can be decomposed as follows

$$\Lambda = \begin{bmatrix} \phi_{pre} \\ \phi_{post} \\ \beta_{pre} \\ \beta_{post} \end{bmatrix} = \begin{bmatrix} \phi_{pre} \\ \phi_{post} \\ \beta_{pre} \\ \delta_{post} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \tau_{post} \end{bmatrix} =: \delta + \tau.$$

The decomposition above accords with that given for β in (2) of RR. We simply extend their framework to include ϕ in the parameter space of interest.

We consider FLCIs based on affine estimators of θ of the general form

$$\mathcal{C}_{\alpha,n}\left(a, \upsilon, \chi\right) = \left(a + \upsilon'\hat{\lambda}_n\right) \pm \chi,$$

where α and χ are scalars, $\upsilon \in \mathbb{R}^m$, and $\alpha \in (0, 0.5]$ denotes a significance level. We choose a and υ to minimize

$$\chi_n\left(a, \upsilon; \alpha\right) = \sigma_{\upsilon,n} \cdot cv_\alpha\left(\bar{b}\left(a, \upsilon\right) / \sigma_{\upsilon,n}\right),$$

where $\sigma_{\upsilon,n} = \sqrt{\upsilon'\Sigma_n \upsilon}$, and $cv_\alpha\left(\cdot\right)$ denotes the $1 - \alpha$ quantile of the folded normal distribution with unit variance, $|\mathcal{N}\left(\cdot, 1\right)|$. The quantity $\bar{b}\left(a, \upsilon\right)$ denotes the worst-case bias of the affine estimator for a given a and υ, namely

$$\bar{b}\left(a, \upsilon\right) := \sup_{\delta \in \Delta, \tau_{post} \in \mathbb{R}^{\overline{T}}} \left|a + \upsilon'\left(\mathbf{J}'\delta + \mathbf{J}'\tau\right) - l'\tau_{post}\right|, \tag{2}$$

where Δ denotes the set of permissible values of δ articulated under the counterfactual.

To showcase the method, we take $\theta = \tau_{\overline{T}}$ (so that interest is focused on the final treatment effect) and define the set of counterfactuals based on the observed trend in the treatment group between $t = \underline{T}$ and $t = 0$ as

$$\Delta_{\bar{T}} = \left\{ \delta_{\overline{T}} : \delta_{\overline{T}} \in \left(-\hat{\phi}_{-\underline{T}} \cdot (\overline{T}/\underline{T}) - \hat{\phi}_{-\underline{T}} - M, -\hat{\phi}_{-\underline{T}} \cdot (\overline{T}/\underline{T}) - \hat{\phi}_{-\underline{T}} + M \right) \right\}, \tag{3}$$

where $-\hat{\phi}_{-\underline{T}} \times (\overline{T}/\underline{T})$ is the counterfactual trend in the final period $\bar{T}$. Thus $\lambda = \left(\phi_{-\underline{T}}, \phi_{\underline{T}}, \beta_{\overline{T}} \right)$ is the subvector of parameters of interest. Using the estimates of $\hat{\phi}_{-\underline{T}} = -35.0$ and $\hat{\phi}_{\underline{T}} = -6.1$ from our empirical application means $\Delta_{\bar{T}}$, as defined in (3), permits $\delta_{\overline{T}} \in \left(35 \times \frac{6}{5} + 6 - 50, 35 \times \frac{6}{5} + 6 + 50 \right) = (-2, 98)$. Figure 8 depicts the corresponding set of counterfactual values of $\phi_{\bar{T}}$ when $M = 50$, leading to the interval $(42 - 50, 42 + 50) = (-8, 92)$. The affine estimator is then defined on $\hat{\lambda}_n = \left(\hat{\phi}_{-\underline{T}}, \hat{\phi}_{\underline{T}}, \hat{\beta}_{\overline{T}} \right)$, and (2) can be more succinctly expressed as

$$\bar{b}(a, \upsilon) := \sup_{\delta_{\overline{T}} \in \Delta_{\bar{T}}, \tau_{\overline{T}} \in \mathbb{R}} \left| a + \upsilon' \left(\begin{bmatrix} \hat{\phi}_{-\underline{T}} \\ \phi_{\underline{T}} \\ \delta_{\overline{T}} \end{bmatrix} + \begin{bmatrix} 0 \\ \phi_{\underline{T}} \\ \tau_{\overline{T}} \end{bmatrix} \right) - \tau_{\overline{T}} \right|.$$

Figure 9 superimposes the confidence region for the set in the final period for $\alpha = 0.05$ and $M = 50$. Notably, the confidence interval sits above zero, meaning the we can reject the null of no treatment effect at the 95% level. It is also centred at a point below $\hat{\beta}_{\overline{T}}$, which is consistent with a negative spillover effect. The span of the confidence set is $(164.1, 426.9)$, indicating the identified set is significant at the 95% level.

[**Insert Figure 9 here**]

The figure also includes confidence sets for 2016, 2017, 2018, 2019 and 2020, for which the counterfactual set is defined in Figure 8. Specifically, for each $t = 1, \ldots, \bar{T}$ the set is given by

$$\Delta_t = \left\{ \delta_t : \delta_t \in \left(-\hat{\phi}_{-\underline{T}} \cdot \frac{t}{\underline{T}} - \hat{\phi}_{\underline{T}} - M \frac{t}{\overline{T}}, -\hat{\phi}_{-\underline{T}} \cdot \frac{t}{\underline{T}} - \hat{\phi}_{\underline{T}} + M \frac{t}{\overline{T}} \right) \right\},$$

where again $M = 50$. The 95% confidence sets for 2018 and 2019 also lie above zero, indicating statistically significant treatment effects, allowing for a counterfactual set defined with $M = 50$.

5.2   Set Identification Results

We now consider set identification under various counterfactual set lengths. Figure 10 exhibits the set-identified treatment effects for $M = 100, 150$ and 200. In all cases, the identified set lies above zero for 2021. The largest counterfactual set, M = 200, spans 7,744 ($= 32 \times (42 + 200)$ to $-5,056$ ($= 32 \times (42 - 200)$). The upper bound of this set is approximately six times the extrapolated linear trend (since $7744/1344 = (200 + 42)/42 \simeq 6$). This serves to pin down the magnitude of

13

the increase in consents that would be necessary under a counterfactual scenario to render the treatment effects statistically insignificant. Specifically, we must allow for counterfactual scenarios that imply a six-fold increase in consents over the pre-treatment trend in order for the estimated treatment effects to become statistically insignificant. Such a scenario appears highly improbable. There is no concurrent policy change in the narrative record that could plausibly generate such a substantive increase in construction.

To help illustrate the substantive growth in consents that would be required under the counterfactual to render all of the treatment effects insignificant, Figure 15 in the Appendix presents the counterfactual set when $M = 200$. The counterfactual set can even accommodate limited forms of exponential growth in consents over the six year post-treatment period, including a year-on-year growth rate of 25.8%.[7]

[**Insert Figure 10 here**]

5.3 Treatments Effects for Attached and Detached Housing

This section applies the set identification approach to the subsamples of attached and detached housing.

5.3.1 Attached Housing

The empirical model is first fitted to dwelling consents for attached housing. Figure 15 in the Appendix exhibits the estimated period effects and extrapolated linear trend between the first period and the treatment period. Evidently, there is only a slight upward trend in attached dwelling consents prior to treatment, with the extrapolated trend yielding only an additional 4.7 dwellings in 2021 compared to 2015. This implies an additional 146 (= 4.7 × 32) dwellings in total across the 32 local areas under the extrapolated counterfactual scenario.

Figure 11 exhibits the estimated treatment effects and set-identified confidence intervals under various counterfactual sets ($M = 125, 150, 175$). The 95% confidence interval for the identified set for the $M = 175$ case sits above zero for 2021. Thus, treatment effects are statistically significant at the 95% level even when the counterfactual set spans an additional 5,747 (= (175 + 4.7) × 32) dwellings under the counterfactual scenario. This is about forty times the 146 additional consents in the control area implied by the extrapolated linear trend. The conclusion seems unequivocal that upzoning has had a substantive impact on the construction of attached dwellings.

[**Insert Figure 11 here**]

---

[7]In 2015 there were 1,955 consents in non-upzoned areas. The upper bound permits up to 7,744 consents in 2021, corresponding to an annual growth rate of 25.8%.

5.3.2 Detached Housing

Next we fit the model to the subsample of detached dwelling consents. Figure 15 in the Appendix exhibits the estimated period effects and extrapolated linear trend between the first period and the treatment period. Evidently, there is a marked increase in detached dwelling consents prior to treatment. The extrapolated trend then yields 37.3 additional dwellings in non-upzoned areas in 2021 compared to 2015. This would entail an additional 1,194 ($= 37.3 \times 32$) dwellings in total under the extrapolated counterfactual scenario.

Figure 12 exhibits the estimated treatment effects and set-identified confidence intervals under various counterfactual sets ($M = 10, 30, 50$). For $M = 30$, the 95% confidence set sits above zero for every year except 2016 and 2021. In contrast, the confidence set for the $M = 50$ case sits above zero only for 2018 and 2019. The counterfactual extrapolated trend is 24.9 in 2018, while the counterfactual set spans $(24.9 - 50 \times (4/6), 24.9 + 50 \times (4/6)) = (-8.4, 58.2)$. Thus, the treatment effect is statistically significant at the 95% level even when the counterfactual set spans approximately two times the additional consents in the control area implied by the extrapolated linear trend (since $2.33 \simeq 58.2/24.9$). We conclude that upzoning had a substantive impact on the construction of detached dwellings, albeit a smaller impact than that realized for attached dwellings.

[Insert Figure 12 here]

5.4 Treatment Effects by Zone

We now analyze the impact of upzoning in each of the three treatment zones separately. Figure 16 in the Appendix exhibits the treatment effects for THA, MHU and MHS. Treatment effects in all three zones are generally positive and statistically significant in the post-treatment period. The estimated treatment effects for 2021 are 91.0, 129.9 and 124.3 for the THA, MHU and MHS zones, respectively. There is some evidence of a downward pre-treatment trend for THA. The treatment effects in this zone may therefore be under-estimated.

We also consider set identification of confidence intervals. Under the set identification approach there is no obvious method to allocate spillovers from the control group to the treatment groups. For example, all of the spillover could be allocated to a single zone, such as MHS. This would be consistent with *all* of the construction in the control zone spilling over into the MHS zone, and none into the MHU and THA zones.

In the absence of an obvious alternative, we elect to allocate the spillover to each of the three zones according to baseline levels of construction in each zone prior to upzoning. Over 2010-2014, 14.08% of construction activity in these three zones was in THA, 29.22% was in MHU, and 56.7% in MHS. Let w_j denote the weights in zone j. For each $t = 1, \ldots, \bar{T}$ the set of counterfactuals is given by

$$\Delta_t = \left\{ \delta_t : \delta_t \in w_j \left(-\hat{\phi}_{-\underline{T}} \cdot \tfrac{t}{\underline{T}} - \hat{\phi}_{\underline{T}} - M \tfrac{t}{T}, -2\hat{\phi}_{-\underline{T}} \cdot \tfrac{t}{\underline{T}} - \hat{\phi}_{\underline{T}} + M \tfrac{t}{T} \right) \right\}.$$

15

Figure 16 exhibits the results with $M = 200$, which is the largest set considered earlier. For THA and MHU the treatment effects are statistically significant for 2018, 2019, 2020 and 2021. Meanwhile, the estimated effects for MHU are statistically insignificant at the 5% level (two-sided). This is due to the fact that the treatment effects are small relative to the proportion of construction activity occurring in this zone prior to the policy. For example, the treatment effect for MHU was 129.9, which is comparable to the 124.3 treatment effect in MHS. But MHS is allocated 57% of the spillover whereas MHU receives only 29% of it.

5.5 Robustness Checks

This section reports robustness checks of our baseline results. The first check is to include a set of controls in the model. The second is to use a much finer resolution of geographic area as the unit of analysis, which affords the use of a larger set of control variables.

Our set of controls includes information on local household incomes and transportation accessibility. These controls can account for changes in macroeconomic conditions that might manifest as spatial variation in construction across different zones. Household income proxies differential access to credit that might manifest as spatial variation in housing demand within Auckland. Distances to the CBD and transport network nodes can account for how changes in congestion affect spatial variation for housing demand in the city.

5.5.1 Including Control Variables

In all model specifications there is little evidence to suggest that the parallel trends assumption does not hold. This means that the significant treatment effects are highly unlikely to be driven by a confounding variable that generates the observed long-run differences in upzoned and non-upzoned areas after the treatment. However, parallel trends do not rule out the possibility of other concurrent events or policy changes being responsible for generating the observed increase in consents in upzoned areas when compared to non-upzoned areas.

To explore this possibility, we augment the empirical regression with a set of local controls in order to examine the extent to which the treatment effects could be driven by a confounding policy or variable. To preview the findings, our results show that including demographic and geographic variables in the regression has next to no impact on the estimated treatment effects or the period fixed effects.

In the presence of controls, the empirical model becomes

$$
\begin{aligned}
c_{i,j,t} \;=\; & \alpha_{i,j} + \sum\nolimits_{s=-\underline{T},s\neq 0}^{\overline{T}} \phi_s \mathbf{1}_{s=t} + \sum\nolimits_{s=-\underline{T},s\neq 0}^{\overline{T}} \beta_s \mathbf{1}_{s=t,j=1} + \\
& \sum\nolimits_{s=-\underline{T},s\neq 0}^{\overline{T}} \mathbf{1}_{s=t}\chi'_s X_i + \sum\nolimits_{s=-\underline{T},s\neq 0}^{\overline{T}} \mathbf{1}_{s=t,j=1}\xi'_s X_i + \varepsilon_{i,j,t},
\end{aligned}
$$

where X_i is a vector of observable characteristics for location i. The set of controls includes: average household income, haversine distance from the centroid of the local area to downtown, and dwelling

density. All variables are logged. Additional details on the data can be found in the Appendix.[8]

Figure 13 presents results for all dwellings, attached dwellings and detached dwellings. For each case, we also include the set identified treatment effects under the largest set of control group counterfactual trends that resulted in a statistically significant set-identified treatment effect. We note that the point estimates and confidence sets are little changed from those depicted in Figures 10, 11 and 12.

[Insert Figure 13 here]

5.5.2 Using Smaller Geographic Units

Our baseline model is based on thirty-two Local Areas. For an additional robustness check we consider the impact of using higher resolution geographic units as the basis of analysis. In particular, we use Area Units (AUs) rather that Local Areas as geographic units. AUs are significantly smaller than Local Areas[9]: there are 366 Area Units (instead of 32 Local Areas) in our sample.

Area Units open up a richer range of potential control variables. In addition to census data on dwelling density and median household income of the area unit, we can compute Manhattan distances to the CBD using GIS information on roads and highways and the centroid of the Area Unit. We also compute Manhattan distances to the nearest rapid transit stop (rail, separated-grade busway, or ferry) and the nearest motorway onramp. We can therefore account for the possibility that changes in congestion, or public transport or other costs of commuting act as confounding variables. All control variables are logged and standardized.

Figure 18 in the Appendix exhibits the results. The y-axis spans a smaller range than previously used for local areas due to fewer consents per AU, on average. The patterns in the estimated treatment effects are otherwise strikingly similar to those documented above. These findings lead us to conclude that using a higher resolution geographic unit as the basis of measurement, and a broader set of control variables, has little impact on our findings and conclusions.

5.6 How Much Construction has Upzoning Enabled?

To estimate the impact of upzoning on the total housing stock, confidence sets were constructed with the counterfactual set restricted to the counterfactual trend, i.e. $M = 0$. This implies that the treatment effect is estimated by subtracting $-\hat{\phi}_{-\underline{T}} \cdot \frac{t}{\underline{T}} - \hat{\phi}_t$ from $\hat{\beta}_t$ for each $t = 1, ..., \bar{T}$. In other words, the difference between the extrapolated linear trend and observed consents in the control group is subtracted from the estimated treatment effect. This difference represents the additional consents that would have occurred in the control group had the policy not been implemented. Figure 14 depicts the spillover-robust treatment effects.

---

[8]Other potential controls, such as (log) population density, were omitted due to collinearity concerns. For example, population density is highly correlated with dwelling density.

[9]AUs are non-administrative geographic areas defined by Statistics New Zealand. Within residential urban areas, AUs are typically a collection of city blocks or suburbs and contain 3000–5000 persons.

The spillover-robust treatment effects for each year between 2016 and 2021 are –25.3, 49.2, 131.6, 174.4, 215.3 and 295.5, respectively, yielding a cumulative total of 840.7. This implies 26,903 (= 815.2 × 32) additional dwellings have been consented as a result of the upzoning policy. This figure implies that consents issued per year have approximately doubled as a result of the policy, which averaged 4,148 dwellings per year between 2010 and 2015.[10] The additional 26,903 consents corresponds to 5.07% in the city's extant housing stock. Statistics New Zealand estimates that there were 530,300 dwellings in Auckland by the end of 2016.[11] Because completions of consented dwellings range between 95% and 99% (outside of recessionary periods), the cumulative completed construction enabled by the policy is likely to be between 4.82% (= 5.07% × 0.95) and 5.02% (= 15.07% × 0.99) of the dwelling stock of Auckland. It is important to note, however, that consents in upzoned areas are still trending upwards as of 2021, so the full impact of the policy will likely not be known for several more years.

[**Insert Figure 14 here**]

Residential Zones. The analysis was repeated for estimated treatment effects by residential zone. We used the same weights as used in Section 5.4 above to allocate the spillovers in each of the three zones.

Figure 17 in the Appendix shows the results. For THA, the spillover-robust treatment effects for each year between 2016 and 2021 are 4.4, 6.3, 48.1, 59.1, 60.9 and 87, respectively, yielding a cumulative total of 268.9 consents per local area, or 8,605 (= 269.9 × 32) additional dwellings in total. For MHU, the corresponding estimates are 0.3, 30.5, 58.2, 90.7, 100.3 and 119.9, yielding a cumulative total of 399.9 consents per local area, or 12,797 (= 399.9 × 32) across all local areas. Finally, for MHS the corresponding estimates are –9.5, 29.6, 67.5, 71.0, 89.0 and 100.8, yielding a cumulative total of 347.3 consents per local area, or 11,115 (= 347.3 × 32) in total. These figures imply that 26.4% of the overall increase in consents occurred in THA, 39.4% in MHU, and 34.2% in MHS if we allocate the counterfactual change in consents equally to each of the three zones.[12]

Area Units. Applying the same approach to the sample based on Area Unit geographies with control variables (see Section 5.5.2 above) implies an additional 32,930 consents. This is substantially more than the estimate of approximately 26,903 in the sample without controls. The

---

[10]This number excludes houses built in the special housing areas. As stated above, these are excluded from the analysis.

[11]Source: https://www.stats.govt.nz/experimental/experimental-dwelling-estimates

[12]The sum of the treatment effects across the three zones is not equal to the treatment effect when the treatment effect is obtained when the three zones are aggregated. Let β_s^k, $k \in \{\text{ALL,THA,MHU,MHS}\}$ denote the parameter when k is the treatment group ('ALL' denotes the case where all three zones are in the treatment group). Then $\sum_{k \in \{\text{THA,MHU,MHS}\}} \left(\hat{\beta}_s^k + \hat{\phi}_s \right) = \hat{\beta}_s^{\text{ALL}} + \hat{\phi}_s$ for each year s, where $\hat{\phi}_s$ is the same as when (1) is run using each k as the treatment group. To allocate $\hat{\beta}^{\text{ALL}}$ to each of the three zones, we use $\hat{\beta}_s^k + \frac{2}{3}\hat{\phi}_s$ since $\hat{\beta}_s^{\text{ALL}} = \sum_{k \in \{\text{THA,MHU,MHS}\}} \left(\hat{\beta}_s^k + \frac{2}{3}\hat{\phi}_s \right)$. This is equivalent to assuming that the counterfactual change in consents in each zone is equivalent to one-third of the counterfactual implied by $\hat{\phi}_s$.

difference is due to the inclusion of controls, since the Area Unit and Local Area sample yield the exact same accumulated totals by construction (i.e. if we omit controls from the Area Unit sample regression, the estimated model implies 26,903 additional consents – the same amount as implied by the Local Area sample).

6   Discussion

The empirical findings show strong evidence to support the conclusion that upzoning raised dwelling construction in the city of Auckland. Our fitted model shows that 26,903 additional dwellings have been consented as a result of the policy over the first five years subsequent to the policy. This figure is equivalent to 5.07% of the existing housing stock prior to policy implementation. The data reveal that much of this increase is in the form of the more capital intensive, attached (or multifamily) structures in the suburban core of the city.

These findings are a positive sign for proponents of upzoning as a solution to housing affordability. Increased supply of housing is a necessary (albeit not sufficient) first step towards achieving housing affordability through a market-led supply response. In future work we will examine the impact of the policy on house prices.

We conclude by noting that the impact of upzoning on housing construction and housing markets will continue to be felt over coming years. Consents for attached dwellings are still trending upwards and consents for detached dwellings remain significantly above their pre-upzoning average. In future work, and as new data become available, the impact of the policy on the housing stock will be updated and new research will seek to determine the particular characteristics of parcels that predict the uptake of redevelopment. Such findings should be useful in assisting the design and refinement of future upzoning policies.

7 Figures and Tables to appear in the Main Text

Table 1: Summary of Land Use Regulation by Residential Zone under the Unitary Plan

Regulation	Terraced Housing Apartments Zone	Mixed Use Urban Zone	Mixed Use Suburban Zone	Single House Zone
Max. height	16m (five storeys)	11 to 12m (three storeys)	8 to 9m (two storeys)	8 to 9m (two storeys)
Height in relation to boundary	$3m + 45°$ recession plane	$3m + 45°$ recession plane	$2.5m + 45°$ recession plane	$2.5m + 45°$ recession plane
Setback	0m	1m	1m	1m
Site Coverage	50%	45%	40%	35%
Impervious Area	70%	60%	60%	60%
Min. dwelling size (1 bedroom)	$45m^2$	$45m^2$	$45m^2$	n/a
Max. dwellings per site	does not apply	3	3	1
Min. Lot Size	$1200m^2$ (subdivision)	$300m^2$ (subdivision)	$400m^2$ (subdivision)	$600m^2$ (subdivision)

Notes: Tabulated restrictions are 'as of right' and can be exceeded through resource consent notification. Height in relation to boundary and Setbacks apply to side and rear boundaries. Less restrictive height in relation to boundary rules than those tabulated apply to side and rear boundaries within 20m of site frontage. Maximum dwellings per site are the number permitted as of right in the Unitary Plan. Minimum lot sizes do not apply to extant residential parcels.

Figure 1: Auckland Region and Local Areas

20 km
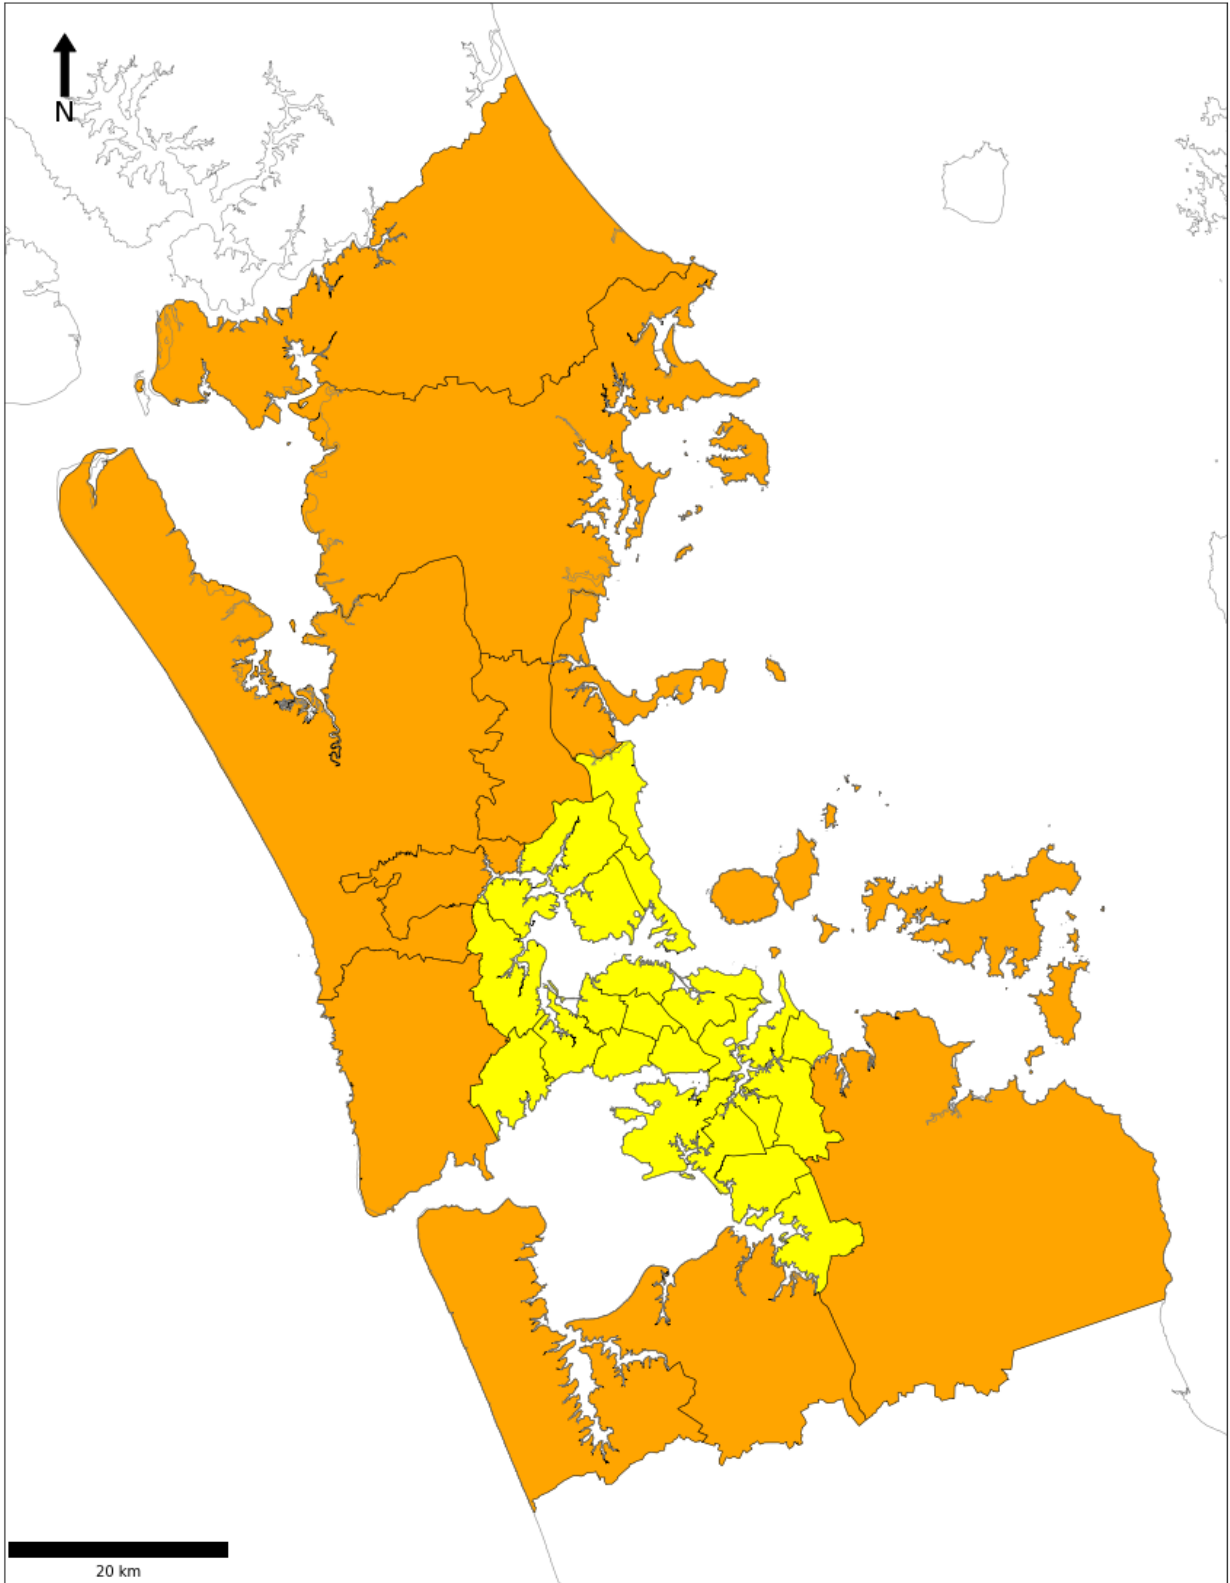
Notes: Local Areas of Auckland (shaded regions). Core urban area in yellow.

Figure 2: Upzoned areas in Auckland

- Residential - Mixed Housing Suburban Zone
- Residential - Mixed Housing Urban Zone
- Residential - Terrace Housing and Apartment Building Zone

- Residential - Single House Zone
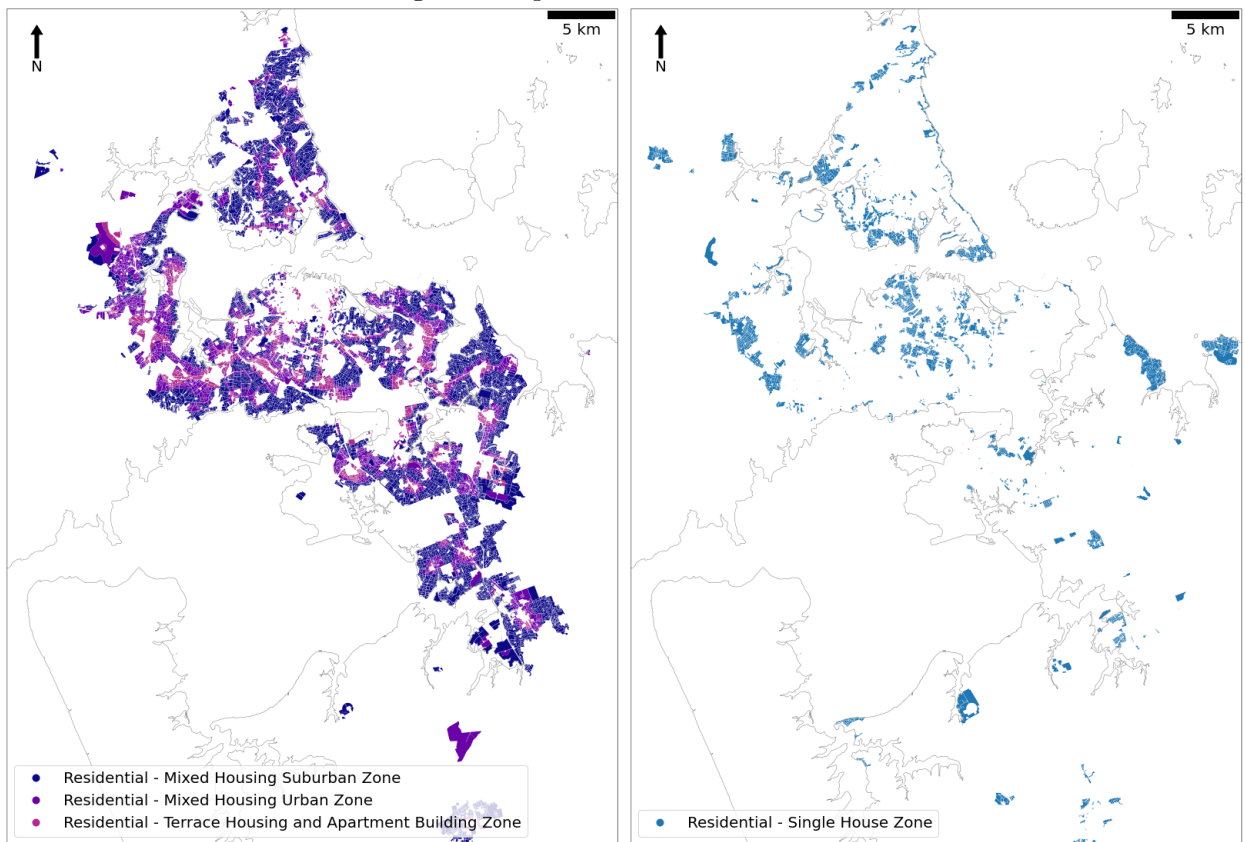
Figure 3: Dwelling consents, 2010-2021

All Dwellings

Attached Dwellings

Detached Dwellings
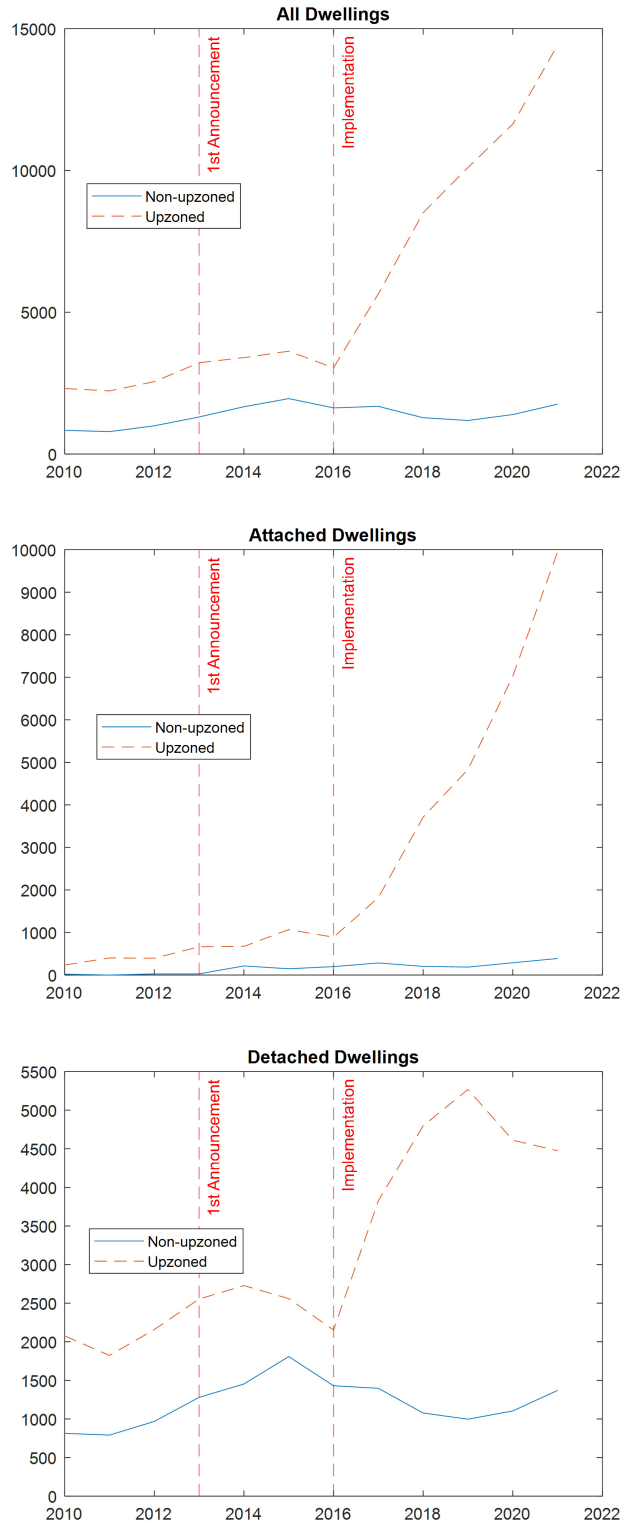
Note: Dwelling consents issued in upzoned and non-upzoned areas.

Figure 4: Dwelling consents inside and outside urban core, 2010-2021
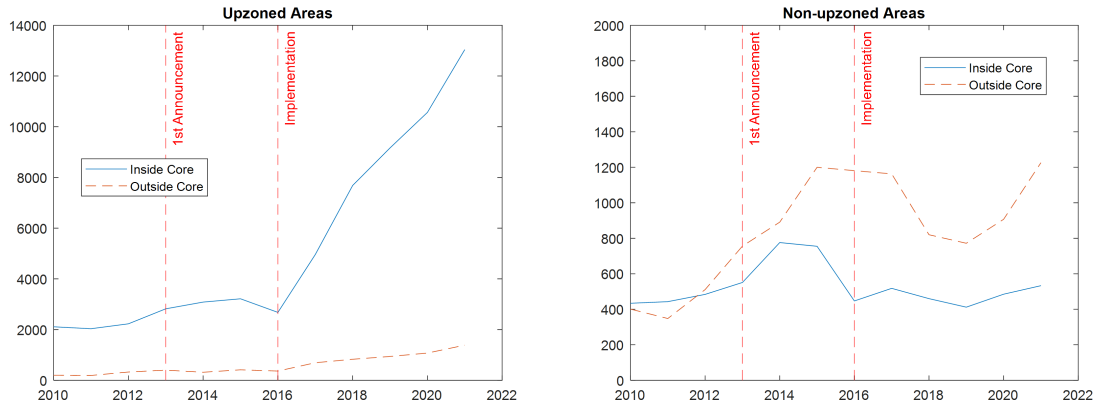
Note: Dwelling consents issued inside and outside urban core by upzoned (left) and non-upzoned areas (right).

Figure 5: Dwelling consents by residential zone, 2010-2021
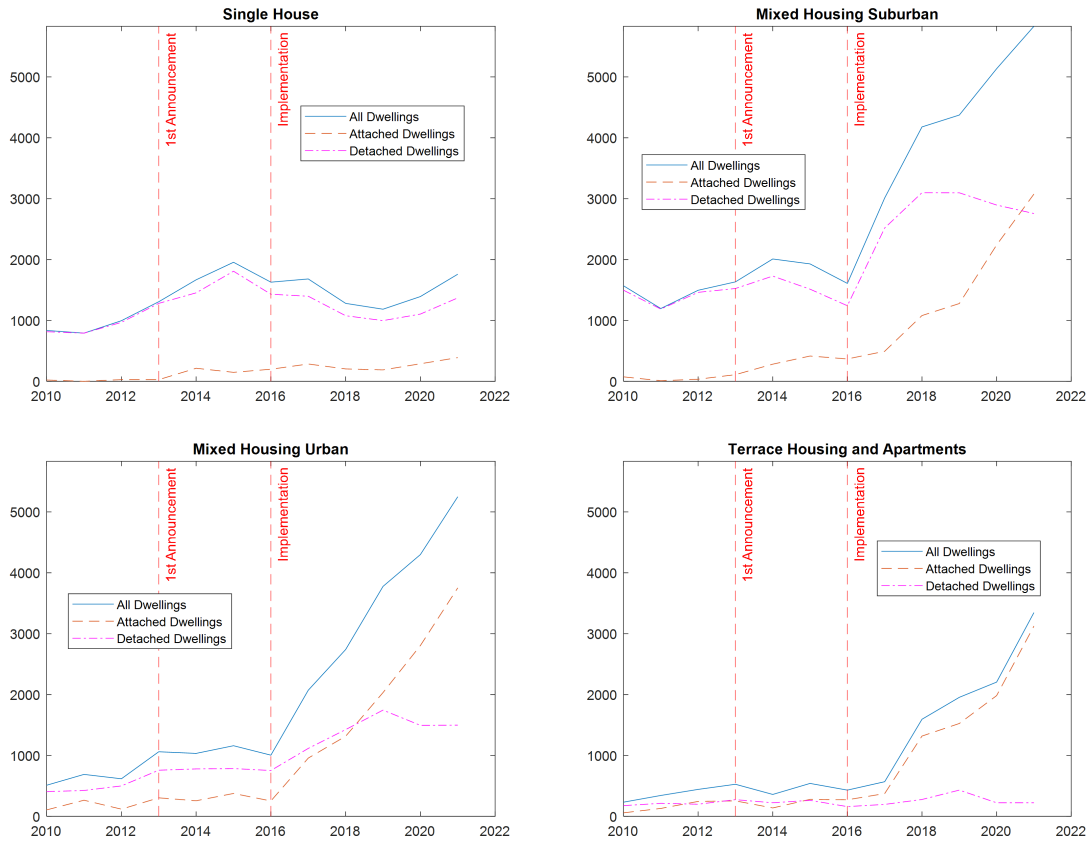
Note: Dwelling consents issued in the four AUP residential zones considered.

Figure 6: Estimated Treatment Effects, 2010-2021
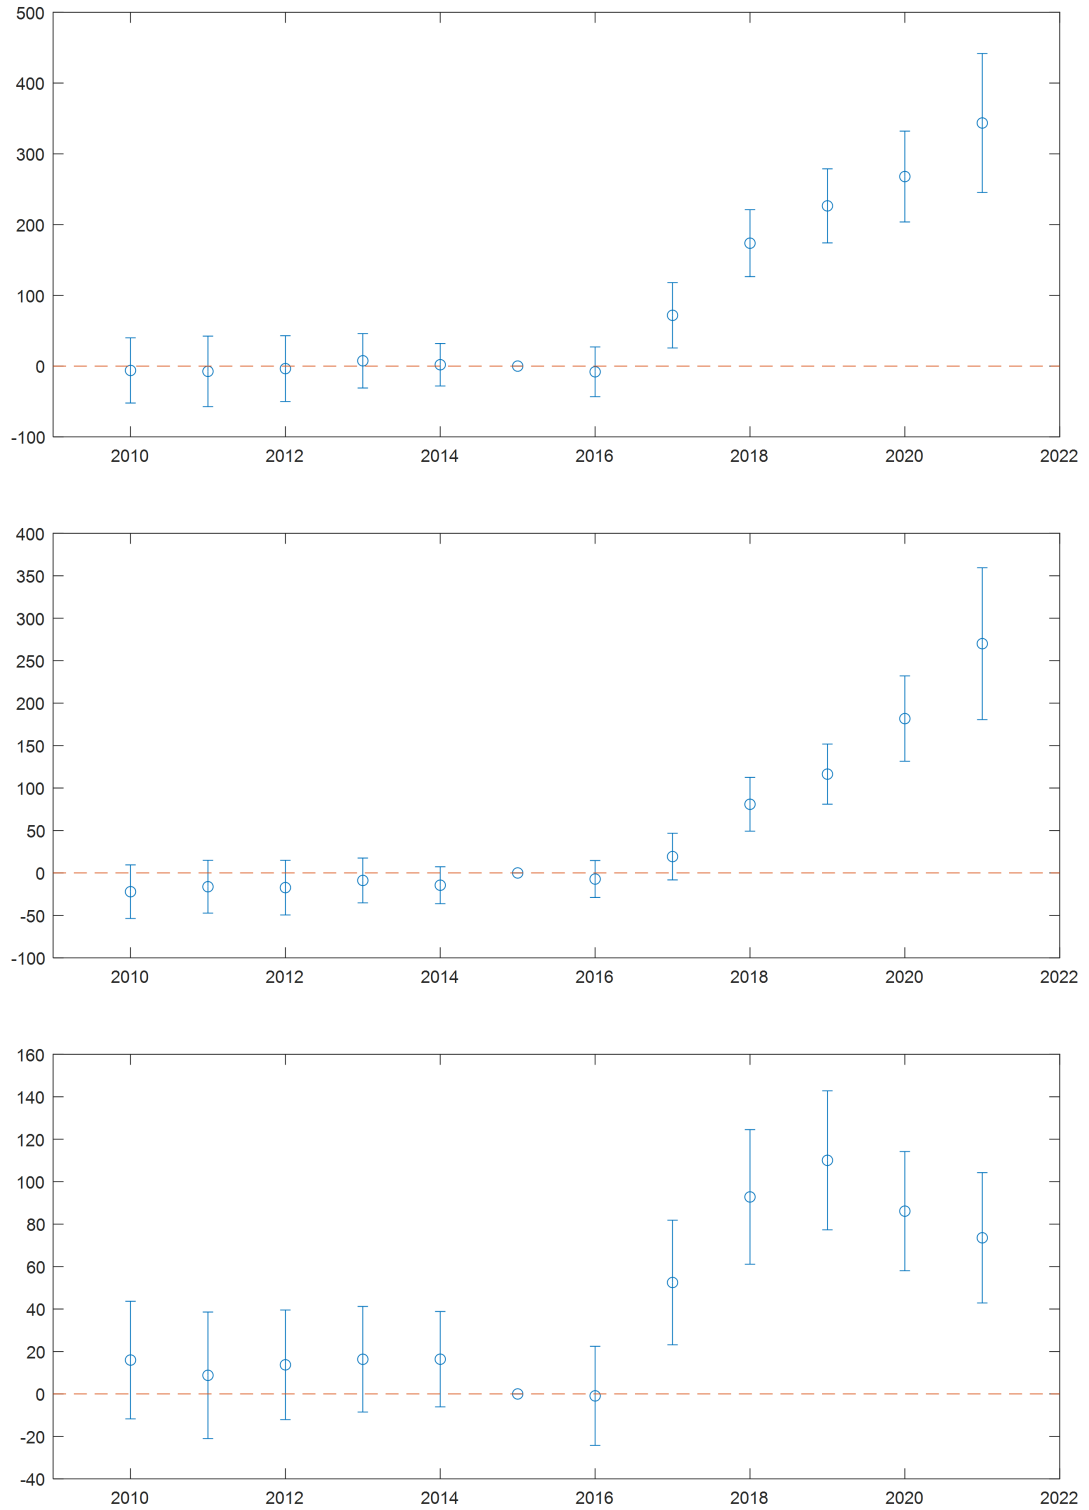
Notes: Estimated multi-period treatment effects for all dwellings (top), attached (middle) and detached dwellings (bottom). OLS estimates of $\{\beta_s\}$ (circles) and 95% confidence intervals (lines).

Figure 7: Estimated Period Effects, 2010-2021
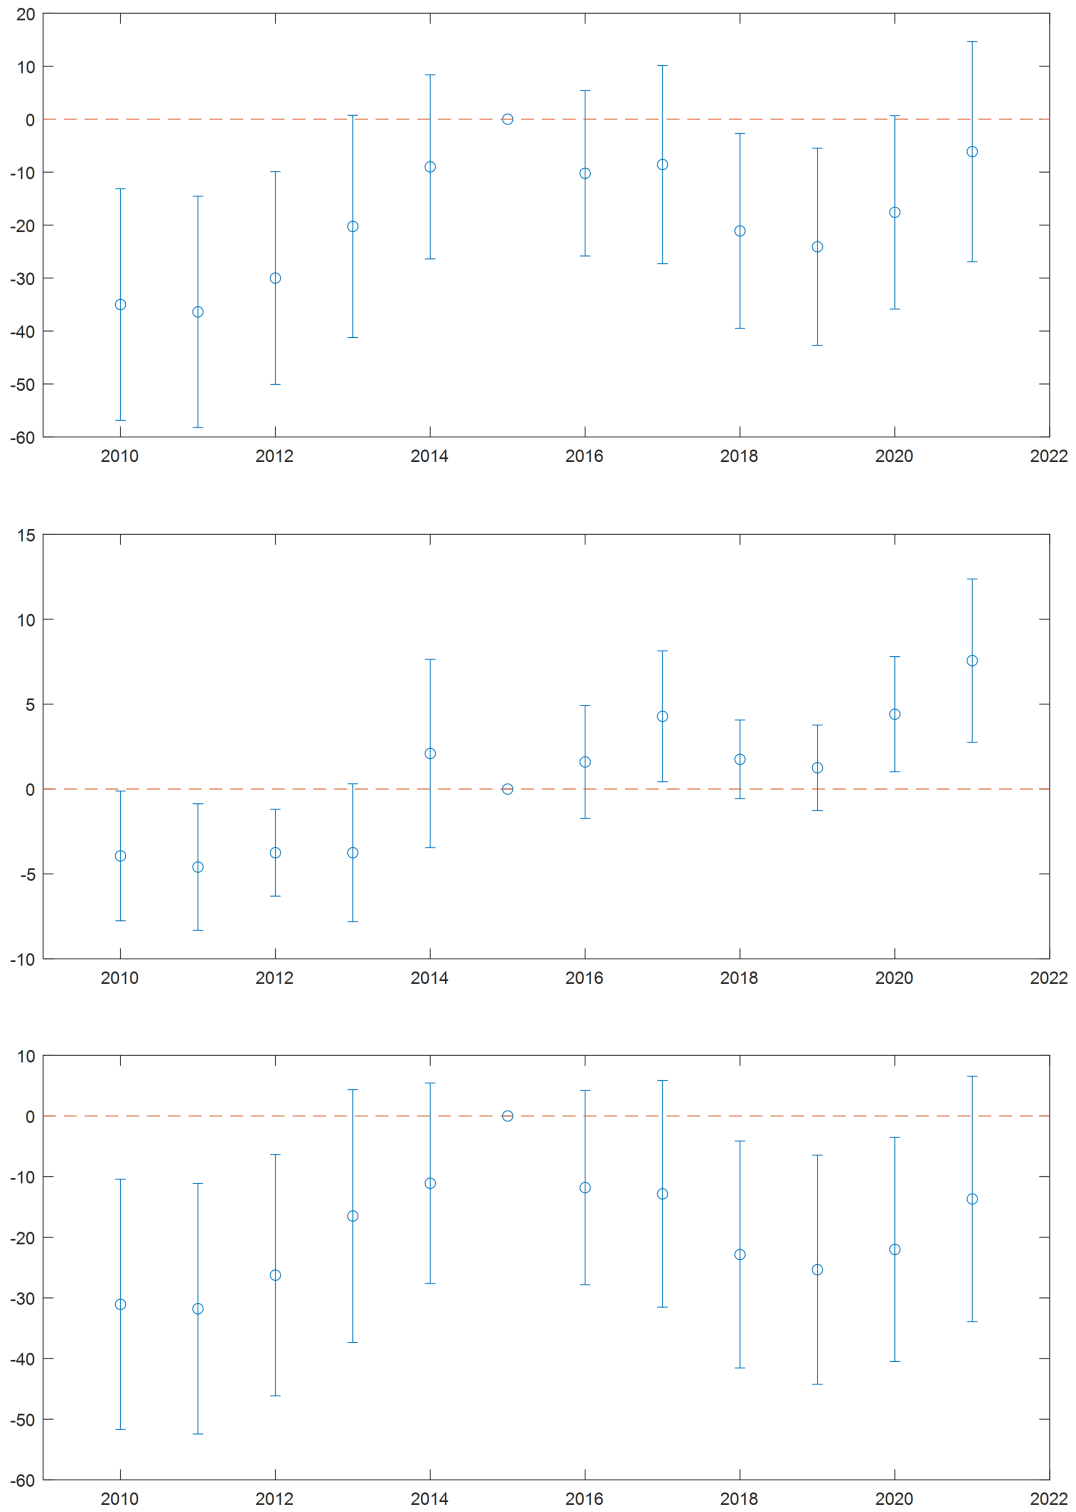
Notes: Estimated period effects for all dwellings (top), attached (middle) and detached dwellings (bottom). The estimated period effects are the average dwelling consents (across local areas) in non-upzoned areas relative to the treatment year 2016. OLS estimates of $\{\phi_s\}$ (circles) and 95% confidence intervals (lines).

Figure 8: Estimated Period Effects and Counterfactual Sets, 2010-2021
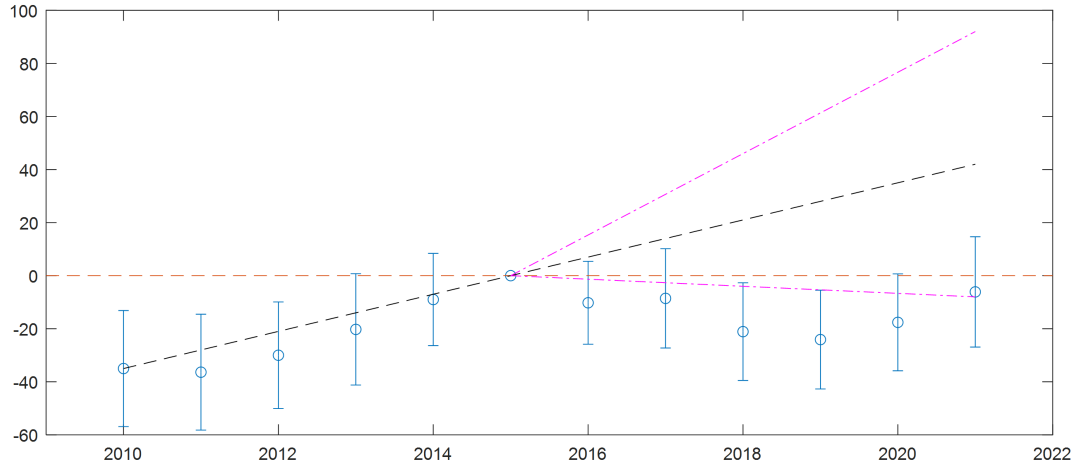
Notes: Outcomes in control group. OLS estimates of $\{\phi_s\}$ (blue circles) alongside 95% confidence intervals (blue lines). Extrapolated pre-treatment trend given by the (black) dashed line. Set of counterfactual outcomes for 2016 to 2021 given by the space between the (pink) dot-dash lines. The counterfactual set is specified as ± 50 dwellings in the final period (equivalent to ± 1600 dwellings over the 32 local areas).

Figure 9: Estimated Treatment Effects and Confidence Sets, 2010-2021
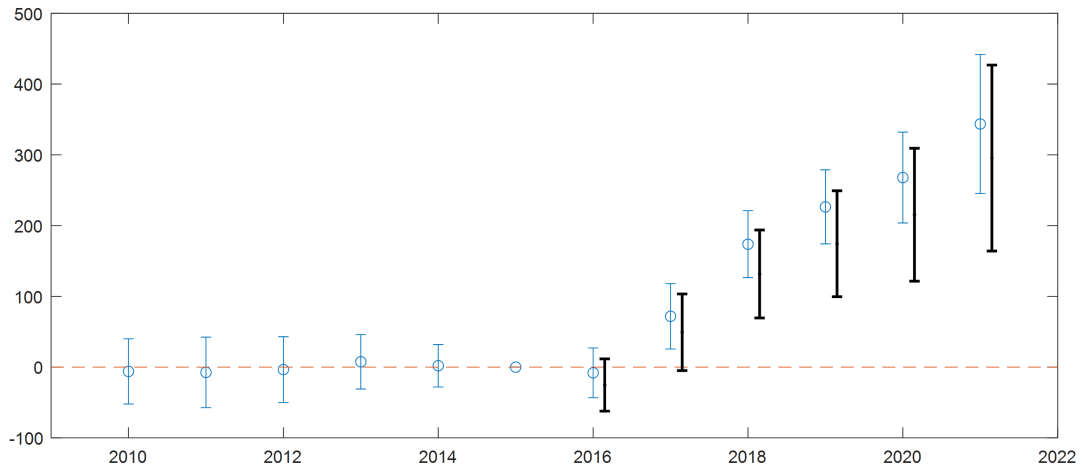
Notes: Multi-period treatment effects and 95% confidence intervals (blue circles and thin lines), and confidence sets for post-treatment interventions, 2016 to 2021 (thick black lines). The counterfactual set is specified as $M = 50$ dwellings in the final period.

Figure 10: Estimated Treatment Effects and Confidence Sets, 2010-2021.
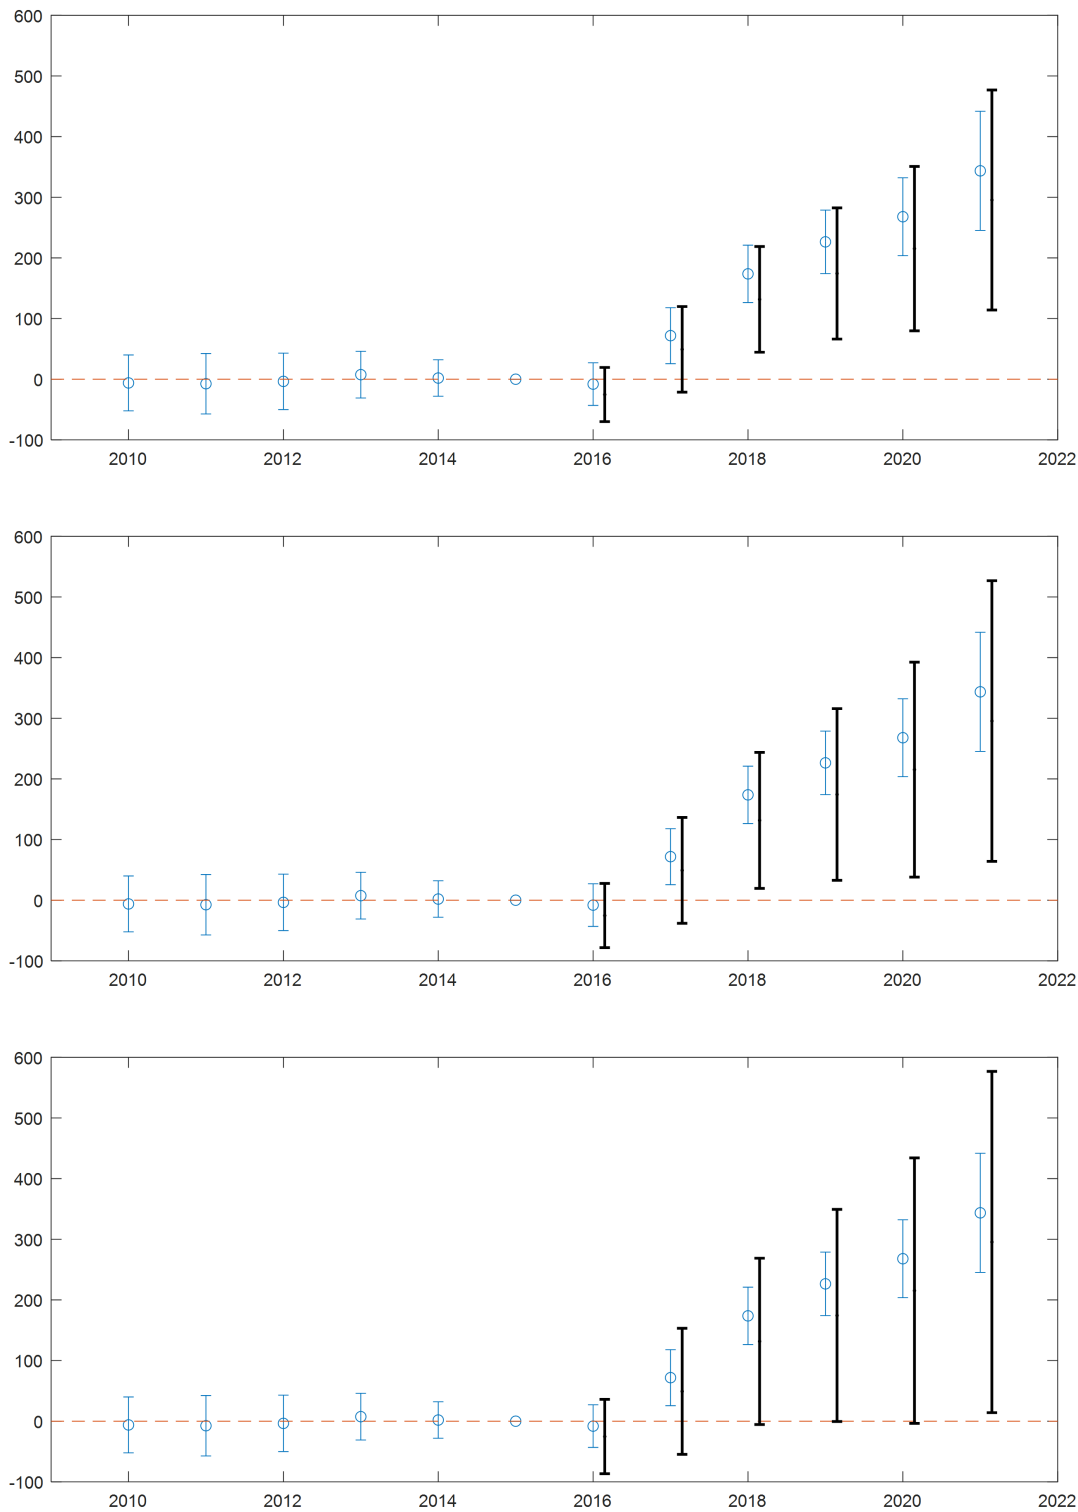
Notes: Multi-period treatment effects and 95% confidence intervals (blue circles and thin lines), and confidence sets for post-treatment interventions, 2016 to 2021 (thick black lines). Treatment occurs in 2015. Top panel $M = 100$; middle panel $M = 150$; bottom panel $M = 200$.
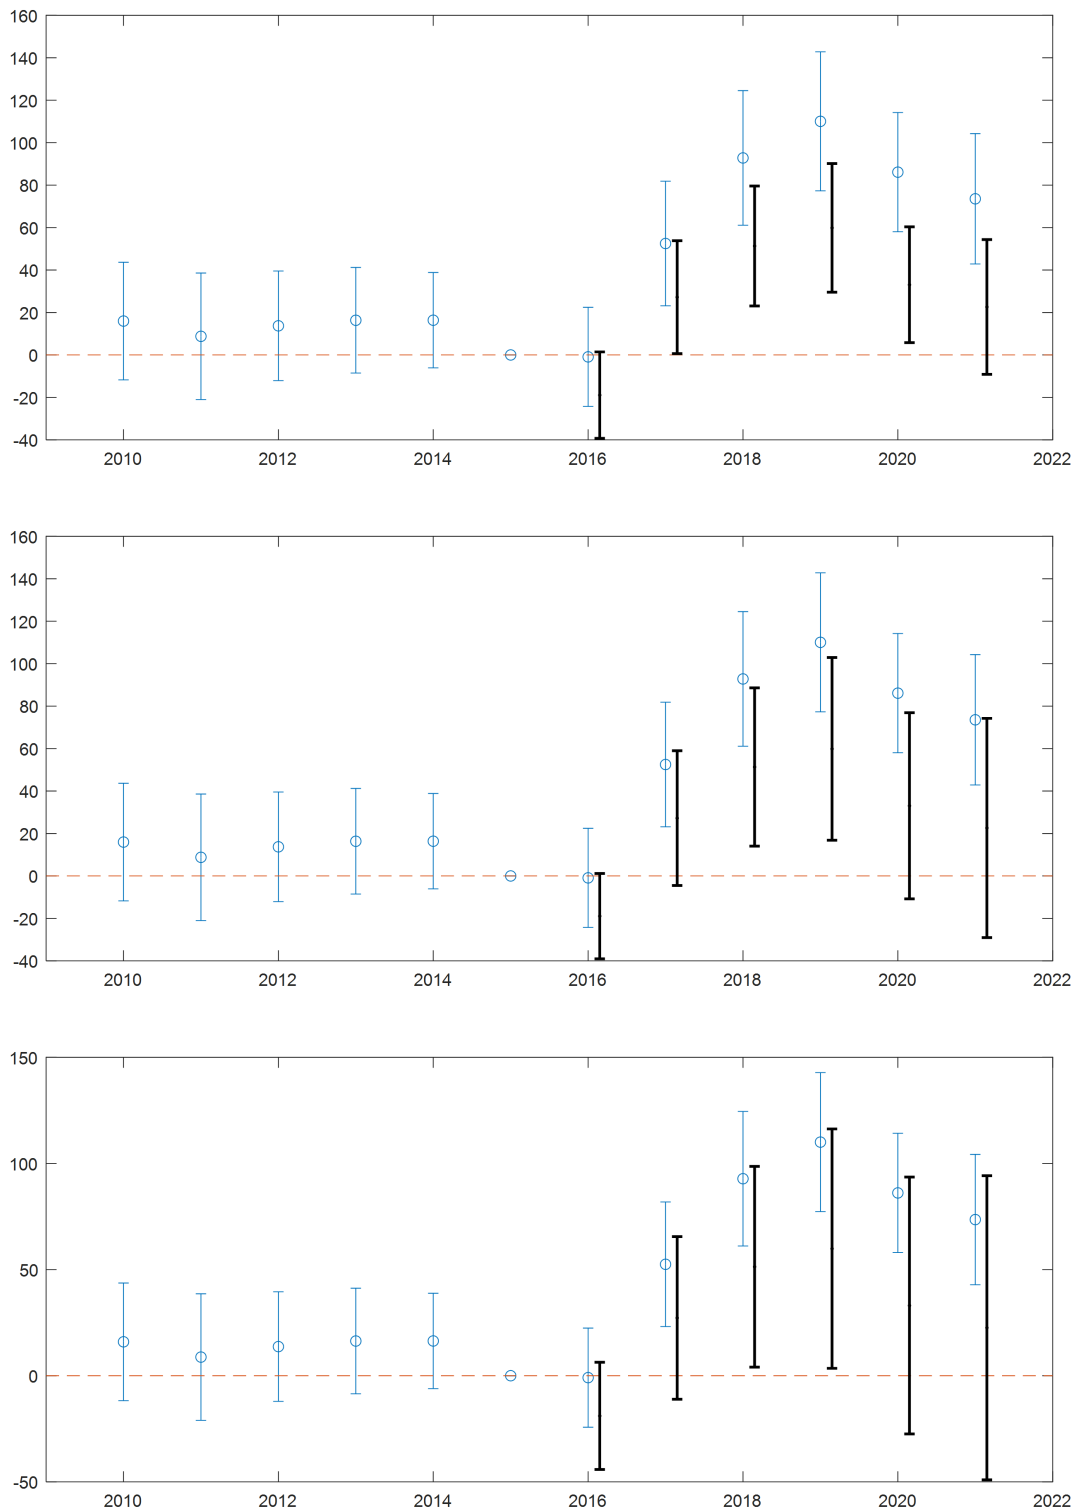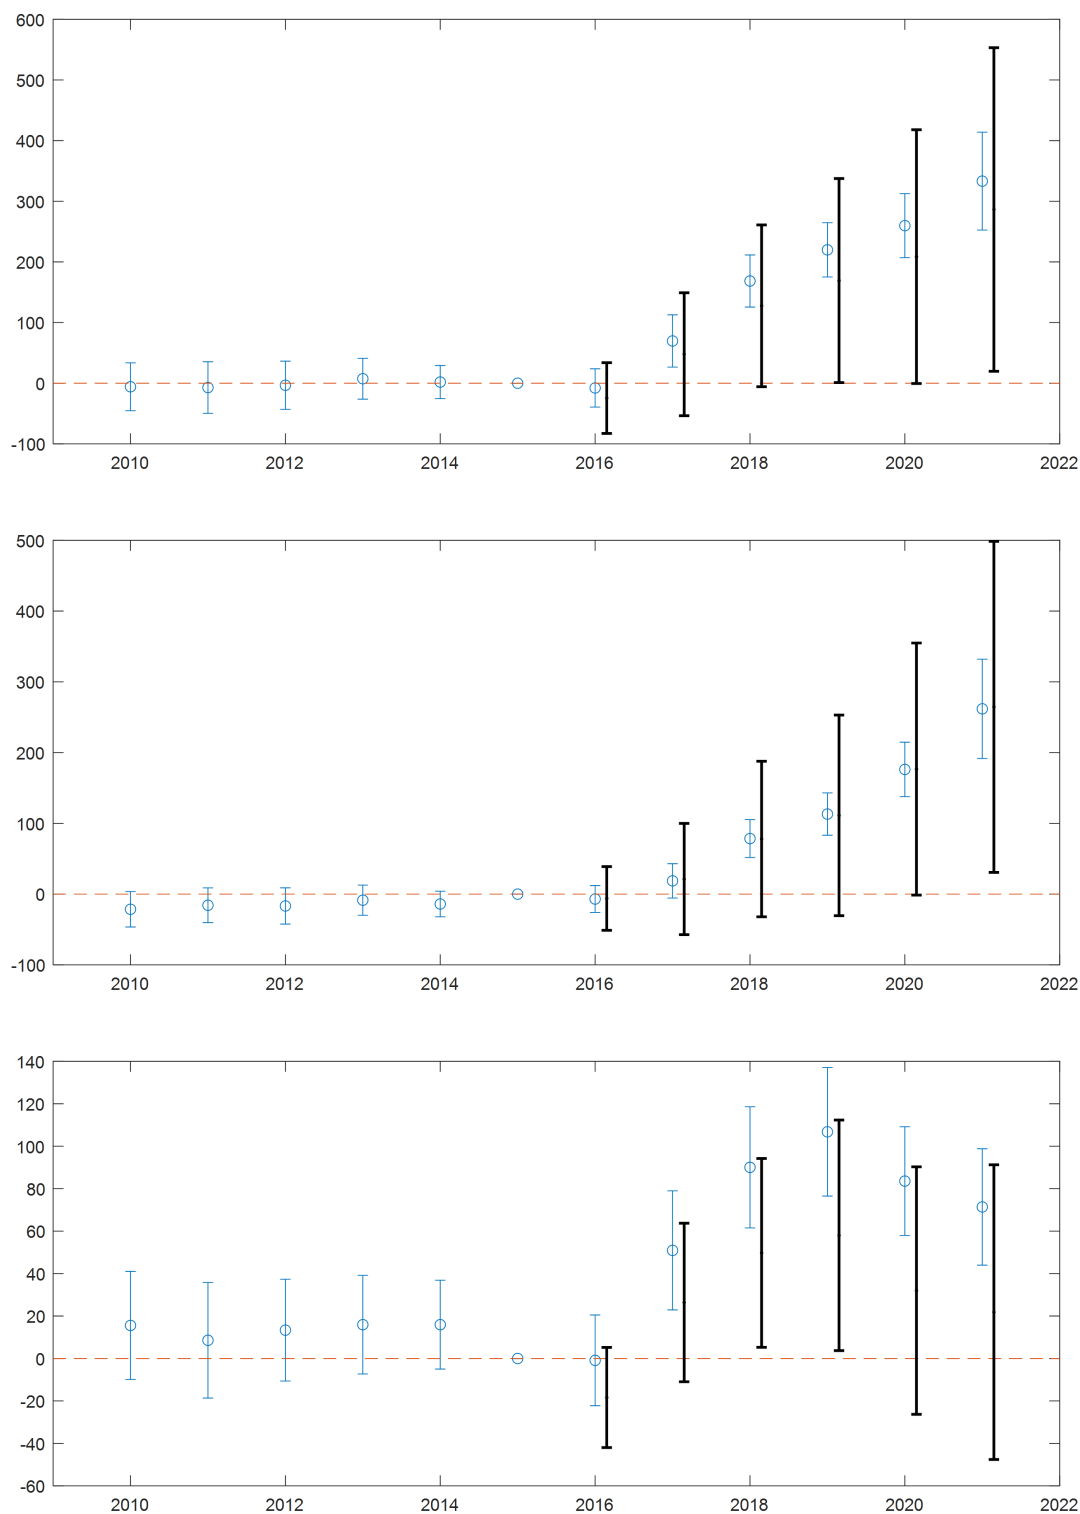
Figure 11: Estimated Treatment Effects and Confidence Sets, Attached Housing, 2010-2021
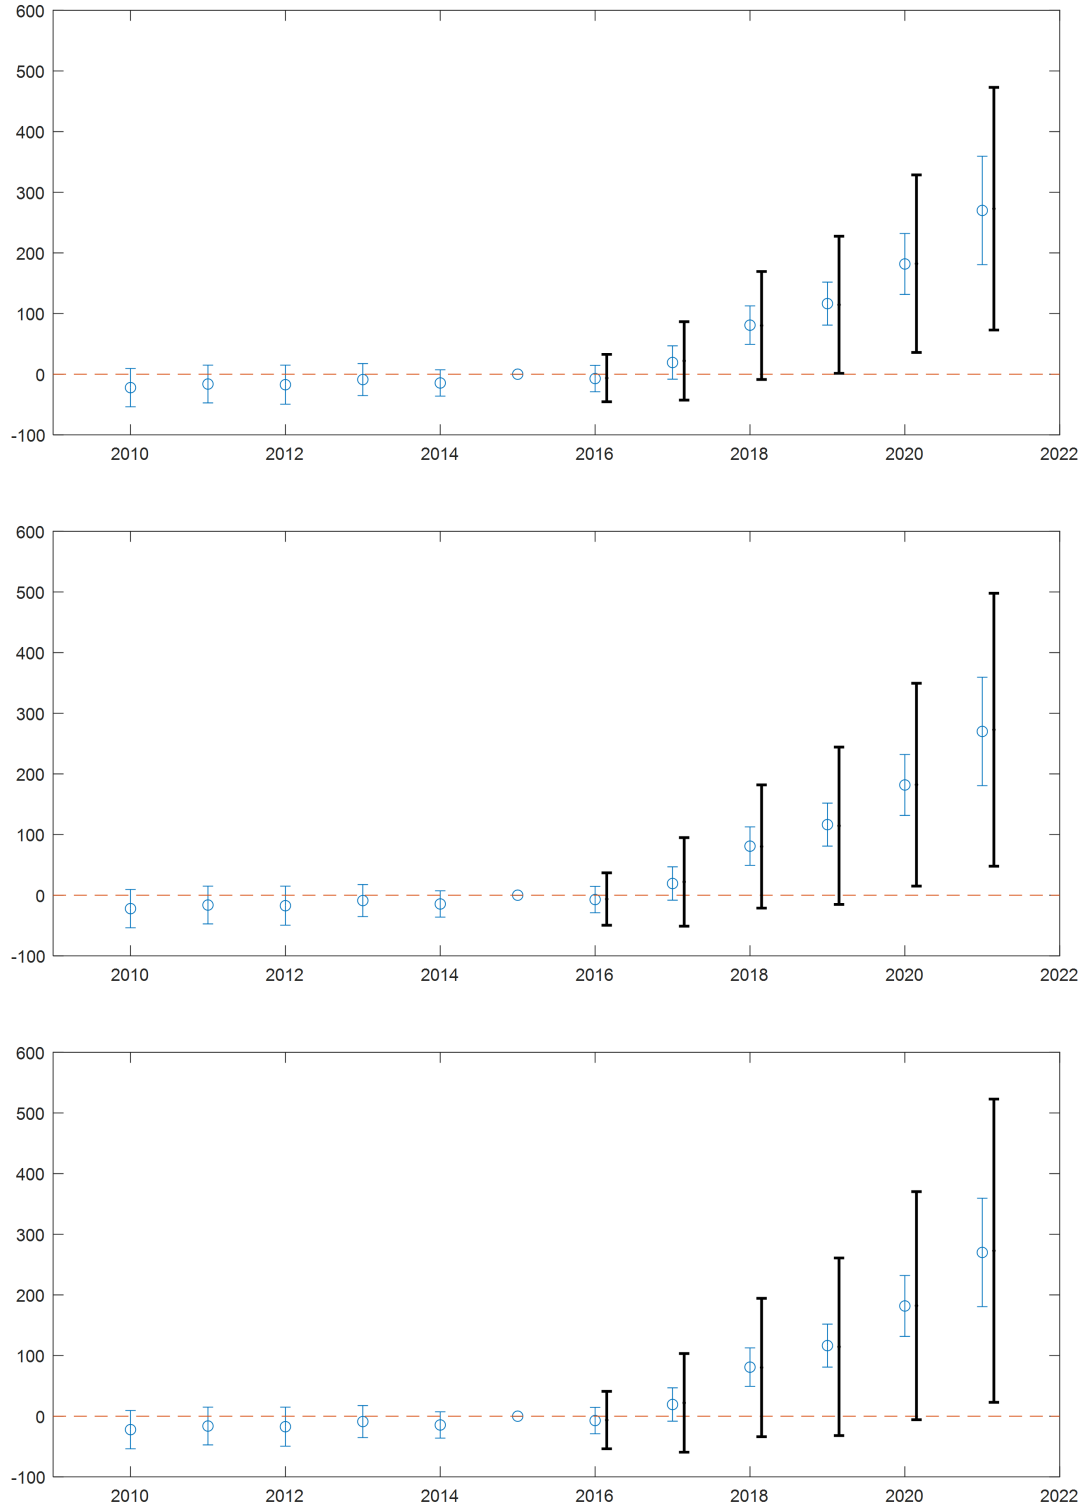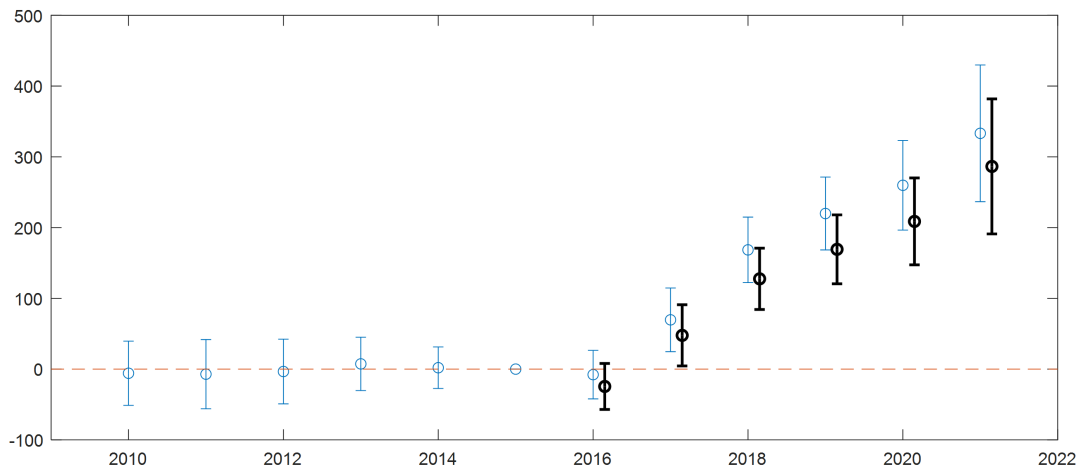

Notes: Multi-period treatment effects and 95% confidence intervals (blue circles and thin lines), and confidence sets for post-treatment interventions, 2016 to 2021 (thick black lines). Treatment occurs in 2015. Top panel $M = 125$; middle panel $M = 150$; bottom panel $M = 175$.

Figure 12: Estimated Treatment Effects and Confidence Sets, Detached Housing, 2010-2021

Notes: Multi-period treatment effects and 95% confidence intervals (blue circles and thin lines), and confidence sets for post-treatment interventions, 2016 to 2021 (thick black lines). Treatment occurs in 2015. Top panel $M = 10$; middle panel $M = 30$; bottom panel $M = 50$.

Figure 13: Estimated Treatment Effects and Confidence Sets, 2010-2021. Regression includes Control Variables

Notes: Multi-period treatment effects and 95% confidence intervals (blue circles and thin lines), and confidence sets for post-treatment interventions, 2016 to 2021 (thick black lines). Treatment occurs in 2015. Top panel is all dwellings with $M = 200$; middle panel is attached dwellings with $M = 175$; bottom panel is detached dwellings with $M = 50$.

Figure 14: Estimated Treatment Effects under Counterfactual Trend for All Dwellings, 2010-2021

Notes: Multi-period treatment effects and 95% confidence intervals (blue circles and thin lines) and treatment effects under counterfactual trend in the control group, 2016–2021 (thick black circles and lines). All dwellings. Estimates under the counterfactual trend are given by restricting the span of the counterfactual set to zero, i.e. $M = 0$. Thick black circles denote spillover-robust point estimates of treatment effects.

8   Appendix

8.1   Background and Detailed Timeline of the Auckland Unitary Plan

Auckland city is the largest metropolitan area in New Zealand with a population of some 1.5 million occupying a narrow isthmus of land in the North Island of the country. Prior to 2010, the greater Auckland metropolitan region comprised one regional council and seven city and district councils. The seven district councils used different land use zones and regulations. On 1 November 2010, Auckland Council (AC) was formed when the eight previous governing bodies in the region were amalgamated. Special legislation was also passed by the central government requiring AC to develop a consistent set of planning rules for the whole region under the Local Government Act 2010. This set of planning rules is embodied in the Auckland Unitary Plan (AUP).

Key dates in the development and implementation of the AUP are as follows:

• 15 March 2013: AC releases the draft AUP. The next 11 weeks comprised a period of public consultation, during which AC held 249 public meetings and received 21,000 items of written feedback.

• 30 September 2013: AC released the Proposed AUP (PAUP) and notified the public that the PAUP was open for submissions. More than 13,000 submissions (from the public, government, and community groups) were made, with over 1.4 million separate points of submission.

• April 2014 to May 2016: an Independent Hearings Panel (IHP) was appointed by the central government, which subsequently held 249 days of hearings across 60 topics and received more than 10,000 items of evidence.

• 22 July 2016: the IHP set out recommended changes to the PAUP. One of the primary recommendations was the abolition of minimum lot sizes for existing parcels. The AC considered and voted on the IHP recommendations over the next 20 working days. On 27 July the public could access and view the IHP's recommendations.

• 19 August 2016: AC released the 'decisions version' of the AUP, including the new zoning maps. Several of the IHP's recommendations were voted down, including a IHP recommendation to abolish minimum floor sizes on apartments. However, the abolition of minimum lot sizes for existing parcels was maintained. This was followed by a 20-day period for the public to lodge appeals on the 'decisions version' in the Environment Court. Appeals to the High Court were only permitted if based on points of law.

• 8 November 2016: A public notice was placed in the media notifying that the AUP would become operational on 15 November 2016.

• 15 November 2016: AUP becomes operational. There were two elements of the AUP that were not fully operational at this time: (i) any parts that remain subject to the Environment Court and High Court under the Local Government Act 2010; and (ii) the regional coastal plan of the PAUP that required Minister of Conservation approval.

All versions of the AUP ('draft', 'proposed', 'decisions' and 'final') could be viewed online.

8.2 Structural Break in Trend in Consents

For $j = 0$ (i.e., control group), we estimate

$$c_{i,j,t} = \alpha_{i,j} + \beta_j 1_{t \geq 1} + \delta_j t + \gamma_j 1_{t \geq 1} t + \varepsilon_{i,j,t}$$

where recall that $c_{i,j,t}$ is the number of consents in zone j in local area i in year t, $1_{t \geq 1}$ denotes an indicator set to one for time periods after the treatment date (2015). OLS estimates are $\hat{\beta}_0 = 35.45$ (t-statistic $= 1.85$), $\hat{\delta}_0 = 7.40$ (t-stat $= 3.67$) and $\hat{\gamma}_0 = -7.67$ (t-stat $= -2.67$), indicating a statistically significant reduction in the upwards linear trend after 2015. t-statistics are based on panel data Newey-West standard errors with the temporal bandwidth set to two periods.

8.3 Control Variables

We use two different sets of controls. Socioeconomic and demographic information is based on 2013 Census data. Local Areas are comprised of 2013 Area Units (AUs), which is one of the geographic resolutions at which 2013 census information is reported. We use the 2013 Census as it is the final census before the AUP is implemented in 2016.

In the local area dataset, average household income is based on average household income by 2013 AU. The AU averages are aggregated up to Local Area averages using the reported observation counts for each AU. Population density is the ratio of 2013 Census population to total land area of the local area. Dwelling density is the ratio of 2013 Census dwelling count to total land area of the local area.

Geographic information is based on GIS information for each Local Area. In the Local Area dataset, distance to downtown is the Haversine distance from the centroid of the Local Area to the skytower in downtown Auckland. In the Area Unit datasets, Manhattan distances are minimum distances using on GIS maps of roads and highways.

Each variable is standardized prior to being included in the regression.

## 8.4 Additional Figures

Figure 15: Estimated Periods Effects and Post-Treatment Counterfactual Sets, 2010-2021.
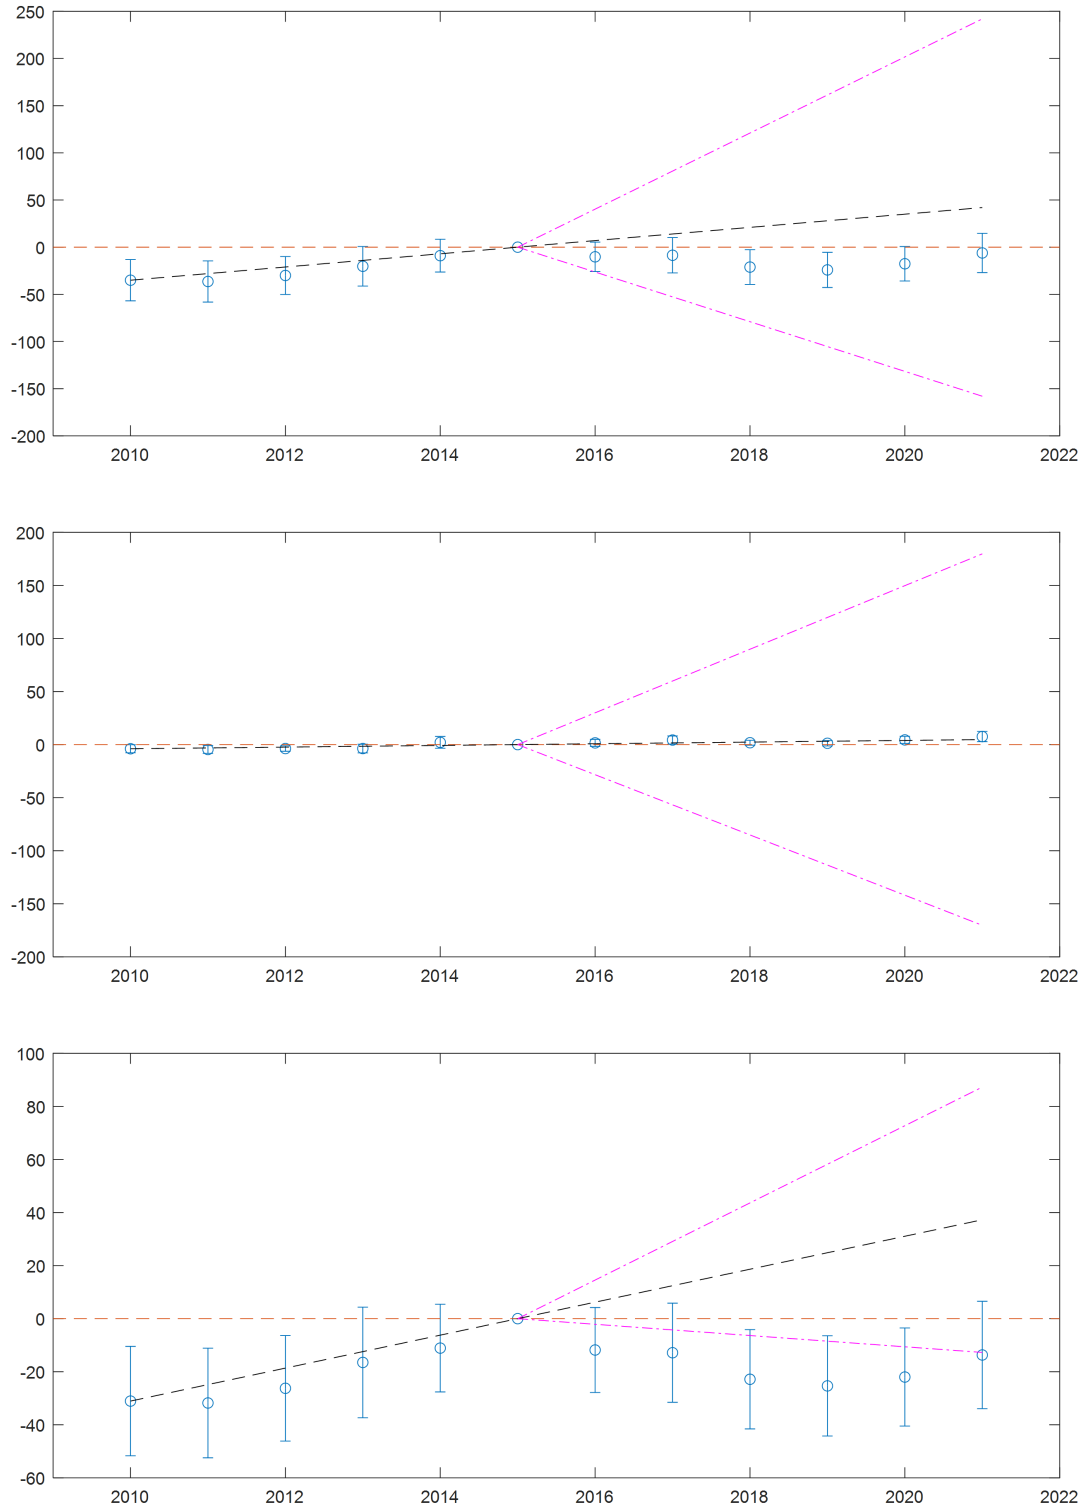
Notes: Outcomes in control group. OLS estimates of $\{\phi_s\}$ (blue circles) alongside 95% confidence intervals (blue lines). Extrapolated pre-treatment trend over 2016 to 2021 given by the (black) dashed line. Set of counterfactual outcomes for 2016 to 2021 given by the space between the (pink) dot-dash lines. Top panel is all dwellings with counterfactual set specified as ± 200 dwellings in the final period; middle panel is attached dwellings with counterfactual set specified as ± 175 dwellings in the final period; bottom panel is detached dwellings with counterfactual set specified as ± 50 dwellings in the final period
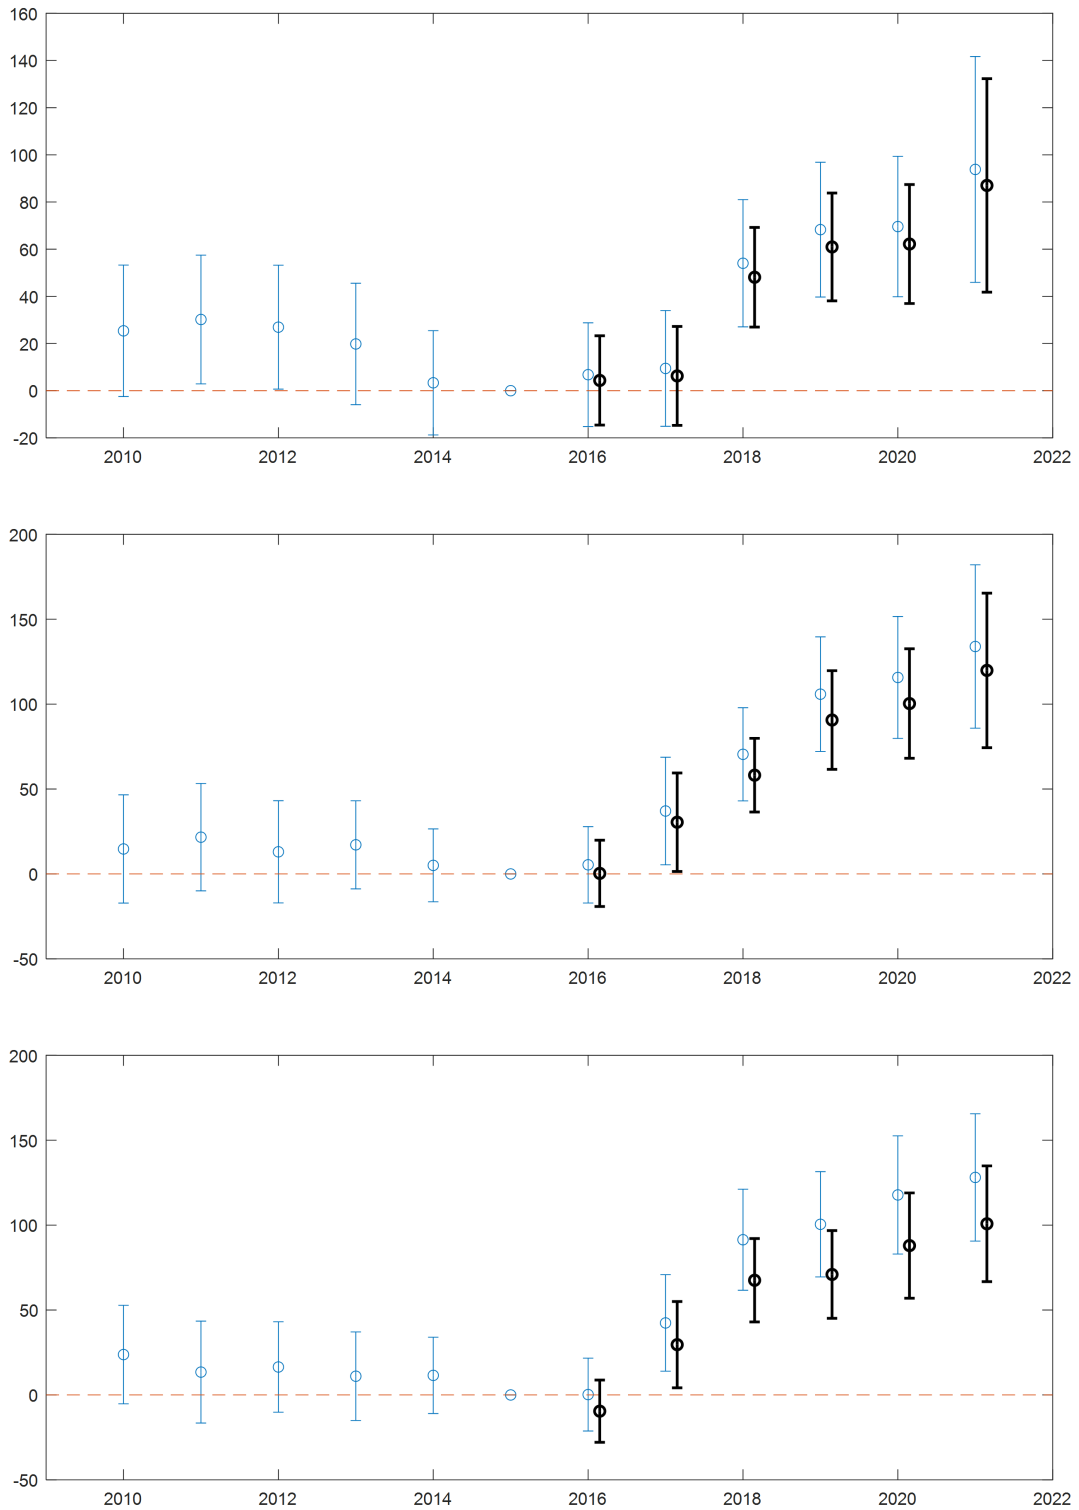
36

Figure 16: Estimated Treatment Effects by Residential Zone, 2010-2021
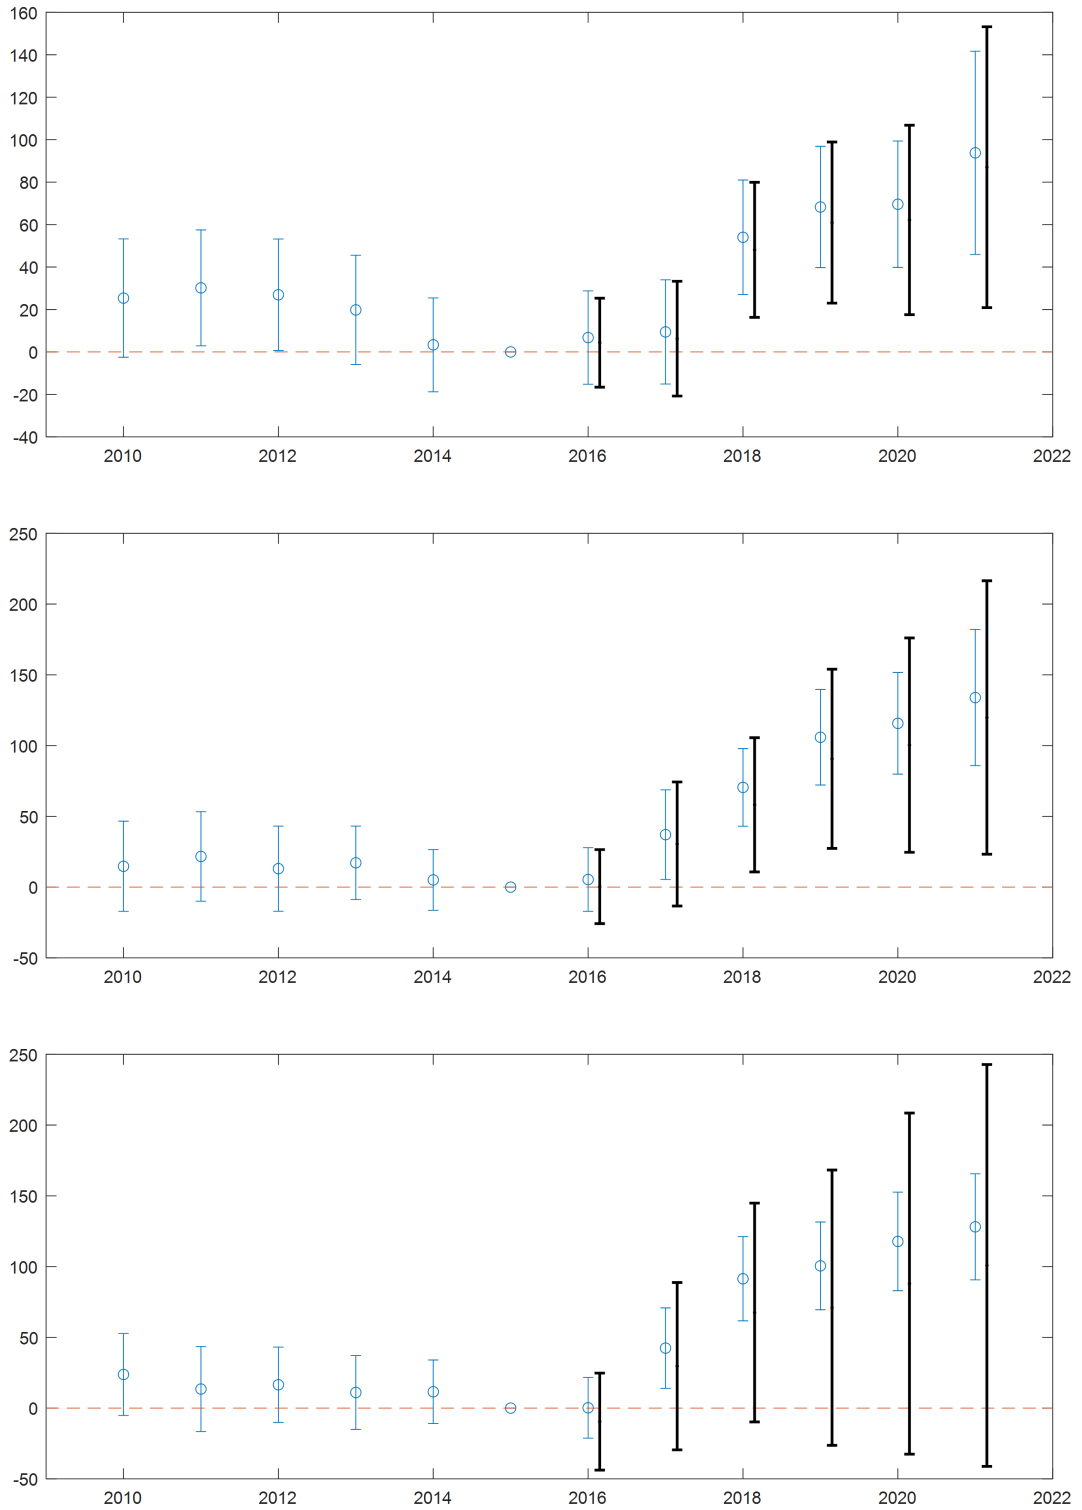
Notes: Multi-period treatment effects (blue) and set-identified confidence intervals for post-treatment interventions (black). THA (top), MHU (middle) and MHS (bottom). $M = 200$.

Figure 17: Estimated Treatment Effects under Counterfactual Trend by Residential Zone, 2010-2021

Notes: Multi-period treatment effects and 95% confidence intervals (blue circles and thin lines) and treatment effects under counterfactual trend in the control group, 2016–2021 (thick black circles and lines). THA (top), MHU (middle) and MHS (bottom). Estimates under the counterfactual trend are given by restricting the span of the counterfactual set to zero, i.e. $M = 0$. Thick black circles denote spillover-robust point estimates of treatment effects.

Figure 18: Estimated Treatment Effects based on Area Units, 2010-2021. Regression includes Control Variables
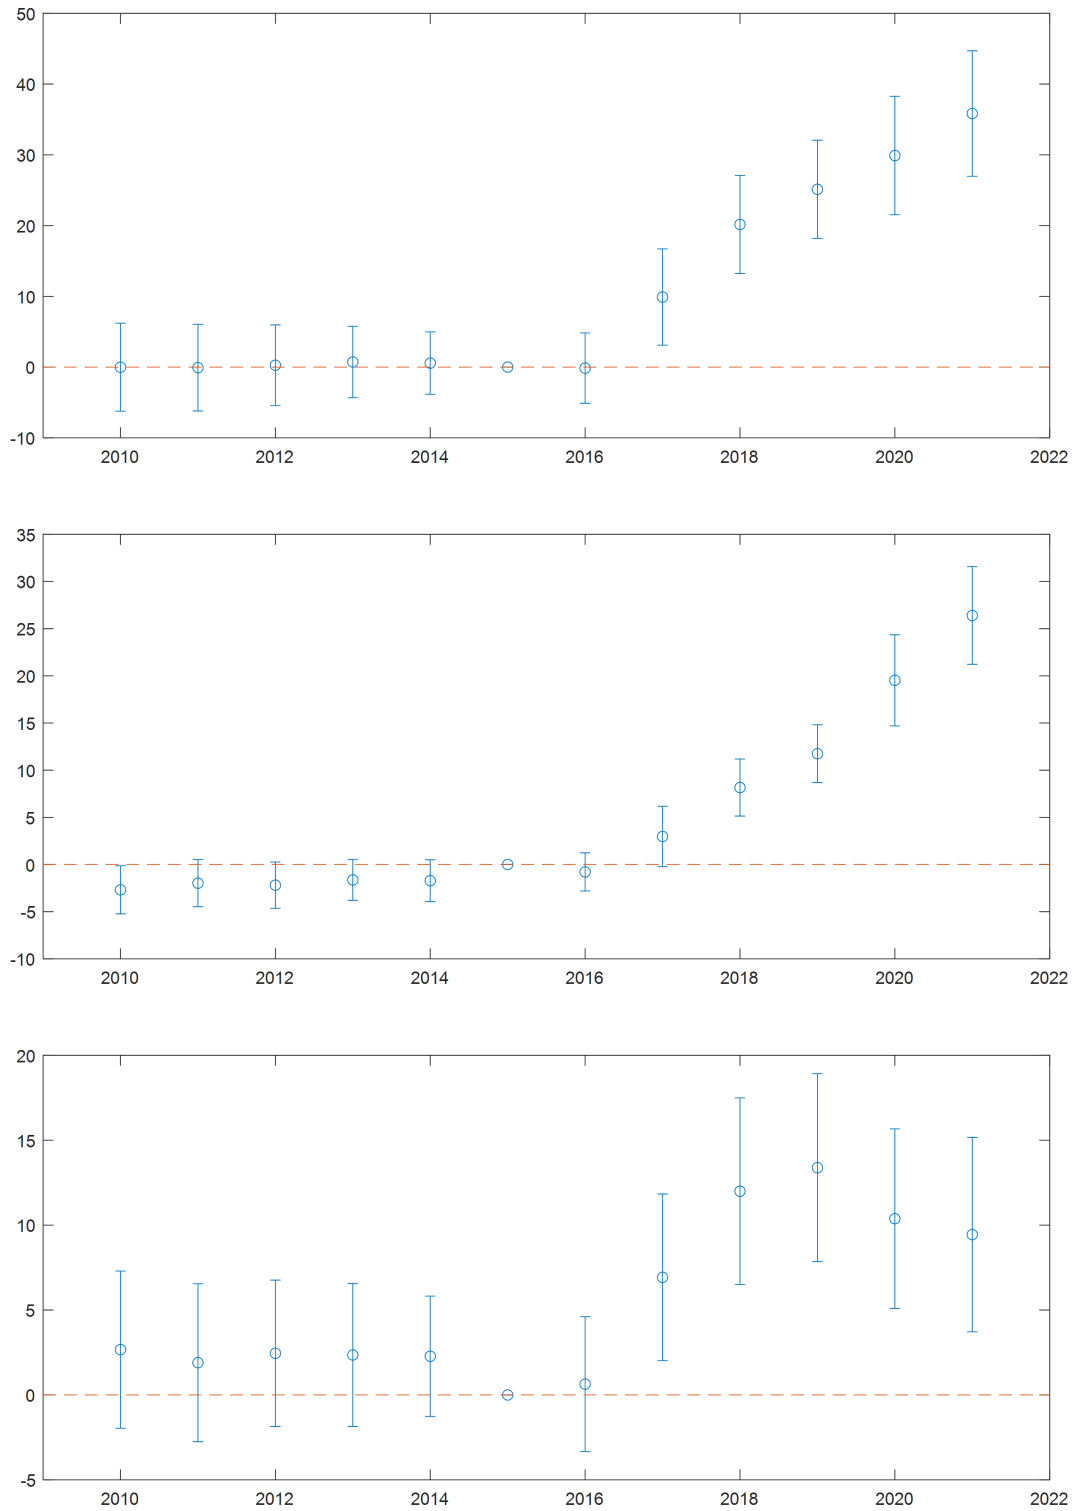
Notes: Multi-period treatment effects and 95% confidence intervals (blue circles and thin lines). Top panel is all dwellings, middle is attached dwellings and bottom is detached dwellings.

References

ATKINSON-PALOMBO, C. (2010): "Comparing the capitalisation benefits of light-rail transit and overlay zoning for single-family houses and condos by neighbourhood type in metropolitan phoenix, Arizona," *Urban Studies*, 47(11), 2409–2426, http://dx.doi.org/10.1177/0042098009357963 2

AUCKLAND UNITARY PLAN INDEPENDENT HEARINGS PANEL (2016): "Report to Auckland Council Overview of recommendations on the proposed Auckland Unitary Plan," https://www.aucklandcouncil.govt.nz/plans-projects-policies-reports-bylaws/our-plans-strategies/unitary-plan/history-unitary-plan/ihp-designations-reports-recommendations/Documents/ihpoverviewofrecommendations.pdf 2

BERTAUD, A., AND J. K. BRUECKNER (2005): "Analyzing building-height restrictions: Predicted impacts and welfare costs," *Regional Science and Urban Economics*, 35(2), 109–125, http://dx.doi.org/10.1016/j.regsciurbeco.2004.02.004 3, 7

BUTTS, K. (2021): "Difference-in-Differences Estimation with Spatial Spillovers," 2105.03737, http://arxiv.org/abs/2105.03737 10

CHAKRABORTY, A., G.-J. KNAAP, D. NGUYEN, AND J. H. SHIN (2010): "The Effects of High-density Zoning on Multifamily Housing Construction in the Suburbs of Six US Metropolitan Areas," *Urban Studies*, 47(2), 437–451 3

CLARKE, D. (2017): "Munich Personal RePEc Archive Estimating Difference-in-Differences in the Presence of Spillovers Estimating Difference-in-Differences in the Presence of," *Munich Personal RePEc Archive*, 52(81604), https://mpra.ub.uni-muenchen.de/81604/ 10

CUNNINGHAM, C. R. (2006): "House price uncertainty, timing of development, and vacant land prices: Evidence for real options in Seattle," *Journal of Urban Economics*, 59(1), 1–31, http://dx.doi.org/10.1016/j.jue.2005.08.003 3

FAVILUKIS, J., P. M. STIJN, V. NIEUWERBURGH, P. MABILLE, AND S. V. NIEUWERBURGH (2019): "Affordable Housing and City Welfare," *SSRN*, http://dx.doi.org/10.3386/W25906 2

FREEMAN, L., AND J. SCHUETZ (2017): "Producing Affordable Housing in Rising Markets: What Works?," *Cityscape: A Journal of Policy Development and Research â¢*, 19 2

FREEMARK, Y. (2019): "Upzoning Chicago: Impacts of a Zoning Reform on Property Values and Housing Construction:," *Urban Affairs Review*, 56(3), 758–789, http://dx.doi.org/10.1177/1078087418824672 2, 4

———— (2021): "Doubling housing production in the Paris region: a multi-policy, multi-jurisdictional response," *International Journal of Housing Policy*, 21(2), 291–305, http://dx.doi.org/10.1080/19491247.2019.1682233 2

GLAESER, E., AND J. GYOURKO (2003): "The Impact of Building Restrictions on Housing Affordability," *Economic Policy Review*, 9(2), 21–39 2

GREENAWAY-MCGREVY, R., G. PACHECO, AND K. SORENSEN (2021): "The effect of upzoning on house prices and redevelopment premiums in Auckland, New Zealand:," *Urban Studies*, 58(5), 959–976, http://dx.doi.org/10.1177/0042098020940602 2

GYOURKO, J., AND R. MOLLOY (2015): "Regulation and Housing Supply," *Handbook of Regional and Urban Economics*, 5, 1289–1337, http://dx.doi.org/10.1016/B978-0-444-59531-7.00019-3 3

HUBER, M., AND A. STEINMAYR (2021): "A Framework for Separating Individual-Level Treatment Effects From Spillover Effects," *Journal of Business and Economic Statistics*, 39(2), 422–436, http://dx.doi.org/10.1080/07350015.2019.1668795 10

IHLANFELDT, K. R. (2007): "The effect of land use regulation on housing and land prices," *Journal of Urban Economics*, 61(3), 420–435, http://dx.doi.org/10.1016/j.jue.2006.09.003 2

JACKSON, K. (2016): "Do land use regulations stifle residential development? Evidence from California cities," *Journal of Urban Economics*, 91, 45–56, http://dx.doi.org/10.1016/j.jue.2015.11.004 2

KAHN, M. E., R. VAUGHN, AND J. ZASLOFF (2010): "The housing market effects of discrete land use regulations: Evidence from the California coastal boundary zone," *Journal of Housing Economics*, 19(4), 269–279, http://dx.doi.org/10.1016/j.jhe.2010.09.001 2, 3

KAHN-LANG, A., AND K. LANG (2020): "The Promise and Pitfalls of Differences-in-Differences: Reflections on 16 and Pregnant and Other Applications," *Journal of Business and Economic Statistics*, 38(3), 613–620, http://dx.doi.org/10.1080/07350015.2018.1546591 6

LIMB, M., AND C. K. MURRAY (2021): "We zoned for density and got higher house prices: Supply and price effects of upzoning over 20 years," Discussion paper, https://osf.io/zkt7v/ 2, 4

MANVILLE, M., P. MONKKONEN, AND M. LENS (2019): "It's Time to End Single-Family Zoning," *Journal of the American Planning Association*, 86(1), 106–112, http://dx.doi.org/10.1080/01944363.2019.1651216 2

MAYER, C. J., AND C. T. SOMERVILLE (2000): "Land use regulation and new construction," *Regional Science and Urban Economics*, 30(6), 639–662, http://dx.doi.org/10.1016/S0166-0462(00)00055-7 2

National Public Radio (2019): "Across America: How U.S. Cities Are Tackling The Affordable Housing Crisis : NPR," https://www.npr.org/2019/08/28/755113175/1a-across-america-how-u-s-cities-are-tackling-the-affordable-housing-crisis 2

Quigley, J. M., and S. Raphael (2006): "American Economic Association Regulation and the High Cost of Housing in California," *American Economic Review*, 95(2), 323–328 2

Rambachan, A., and J. Roth (2022): "A More Credible Approach to Parallel Trends," *Working Paper*, https://scholar.harvard.edu/files/jroth/files/honestparalleltrends_main.pdf 3, 10, 11, 12

Rodríguez-Pose, A., and M. Storper (2020): "Housing, urban growth and inequalities: The limits to deregulation and upzoning in reducing economic and spatial inequality," *Urban Studies*, 57(2), 223–248, http://dx.doi.org/10.1177/0042098019859458 2

Saiz, A. (2010): "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 125(3), 1253–1297, http://qje.oxfordjournals.org/ 2

Schill, M. H. (2005): "Regulations and Housing Development: What We Know," *Cityscape: A Journal of Policy Development and Research â¢*, 8(1), 5–19 2

Thorson, J. A. (1997): "The Effect of Zoning on Housing Construction," *Journal of Housing Economics*, 6, 81–91 2, 3

Wetzstein, S. (2017): "The global urban housing affordability crisis," *Urban Studies*, 54(14), 3159–3177, http://dx.doi.org/10.1177/0042098017711649 2

——— (2021): "Assessing post-GFC housing affordability interventions: a qualitative exploration across five international cities," *International Journal of Housing Policy*, 21(1), 70–102, http://dx.doi.org/10.1080/19491247.2019.1662639 2

Zabel, J., and M. Dalton (2011): "The impact of minimum lot size regulations on house prices in Eastern Massachusetts," *Regional Science and Urban Economics*, 41(6), 571–583, http://dx.doi.org/10.1016/j.regsciurbeco.2011.06.002 2

Zhou, J., D. P. Mcmillen, and J. F. Mcdonald (2008): "Land Values and the 1957 Comprehensive Amendment to the Chicago Zoning Ordinance," *Urban Studies*, 45(8), 1647–1661 3