UNIVERSITY OF AUCKLAND
Waipapa Taumata Rau
NEW ZEALAND
BUSINESS SCHOOL

Economic Policy Centre

The University of Auckland Business School

Evaluating the Long-Run Effects of Zoning Reform on Urban Development

Ryan Greenaway-McGrevy

March 2023

Economic Policy Centre

WORKING PAPER NO. 013

Evaluating the Long-Run Effects of Zoning Reform on Urban Development *

Ryan Greenaway-McGrevy†

March 2023

Abstract

This paper presents a framework for assessing the long-run effects of zoning reform on urban development, housing construction, and dwelling prices. We begin by developing a version of the monocentric model of the city that features regulatory restrictions on the capital intensity of housing that vary across different residential zones. A key implication of the model is that differences in land price gradients between zones reflect differences in the optimal floorspace that developers build to. This facilitates empirical estimation of the anticipated supply response to zoning reform through the measurement of changes in land price differentials soon after policy announcement. We use the framework to evaluate a recently-implemented upzoning policy in Auckland, New Zealand, finding that changes in land price gradients in upzoned areas compared to non-upzoned areas is consistent with an approximate 18.1% increase in aggregate floorspace. Using plausible estimates of the house price elasticity of demand from the extant literature, this supply increase equates to a reduction in dwelling prices ranging between 23% and 39%, relative to the counterfactual of no upzoning.

Keywords: Upzoning, Land Use Regulations, Redevelopment, Housing Construction, AMM model.

JEL Classification Codes: R14, R31, R52

*We thank Auckland Council for providing the dataset of ratings valuations, and Geoff Cooper for helpful suggestions on an earlier draft of the paper. This research was funded by the Royal Society of New Zealand under Marsden Fund Grant UOA2013.

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

1 Introduction

Zoning reform is increasingly advocated in response to a variety of social problems, including increasingly unaffordable housing (Glaeser and Gyourko, 2003; Freeman and Schuetz, 2017), environmental sustainability (Wegmann, 2020) and inequality (Manville et al., 2020). Advocates argue that stringent land use regulations (LURs) restrain housing supply in the places people want to live, pushing costs up, and families out. Relaxing these restrictions would enable more housing through the redevelopment of low density housing into more intensive forms, such as plexes, rowhouses and apartments.[1]

Many municipal, gubernatorial and national governments are now implementing zoning reforms to enable redevelopment. Between 2019 and 2022, the states of Oregon and Maine, and the cities of Minneapolis and Charlotte, passed laws to abolish single-family zoning, and a similar bill currently sits before the Washington State Senate. In March 2023, Victoria, Canada upzoned approximately 45% of single- and double- family lots to allow between 6 and 12 dwellings, while Arlington, Virginia banned single family zoning.[2] Auckland, New Zealand, upzoned approximately three quarters of residential land in 2016 (Greenaway-McGrevy and Jones, 2023), and the New Zealand government has subsequently passed reforms to enable medium density and transit-oriented, high density housing in its largest cities.

To date, the results of these reforms are mixed. Minneapolis' reforms legalized duplexes and triplexes on land previously restricted to detached single family housing, but plexes continue to constitute only a small fraction of total permits. In contrast, the large-scale upzoning in Auckland stimulated an immediate and significant increase in housing construction (Greenaway-McGrevy and Phillips, 2023), while Houston issues building permits at a much higher per capita rate than other US cities after its reduction in minimum lot sizes in 1998 (Gray, 2022). Meanwhile, supply responses to localized upzonings tend to be slow (Limb and Murray, 2022; Peng, 2023) or altogether lacking (Freemark, 2019), contravening anticipated outcomes, contributing to doubt regarding the efficacy of zoning reform to achieve policy objectives (Rodríguez-Pose and Storper, 2020).

Policy evaluation can be hampered by the long time horizon associated with urban redevelopment. Zoning reform targets outcomes that can be reasonably expected to manifest over several decades, making it difficult assess the efficacy and aggregate impact of the reforms based on immediate, short-term trends in building activity. As time passes, the effects of other confounding policies or events become more salient, further inhibiting policy assessment. Even in cases such as Auckland, where upzoning stimulated an immediate and significant increase in housing construction, it remains unknown whether the reform will meet long-run policy targets. Meanwhile, evaluating effects on house prices is complicated by the fact that upzoning can increase land prices, causing properties that are sufficiently intensive in land to appreciate in value (Greenaway-McGrevy et al., 2021; Greenaway-McGrevy and Phillips, 2023). Reductions in house prices via increased supply require the redevelopment of such land intensive properties into more capital intensive housing

[1] "Plexes" refers to duplexes, triplexes, sixplexes, etc.

[2] Proportion obtained via personal correspondence with City of Victoria staff.

forms.

To assist in the difficulties of policy evaluation, this paper develops a method to evaluate the long-run effects of upzoning based on immediate land market responses to zoning changes. We develop a version of the conventional monocentric model that allows land use regulations (LURs) to vary between different zones of the city, showing that upzoning (i.e. a relaxation of LURs) generates changes in land price gradients between upzoned and non-upzoned areas that reflect increases in the optimal floor area that developers build to. Armed with this result, we can calculate the increase in floorspace capacity enabled by the policy by carefully measuring changes in land price gradients between zones with different zoning regulations. These relative changes in land price gradients can be readily estimated using spatial and time series variation in residential zoning and property prices. Then, using information on the amount of land assigned to each zone, we can back out the aggregate increase in floorspace supply enabled under the policy, and, given a price elasticity of demand for housing, the impact of the supply increase on dwelling prices.

We apply the method to evaluate the effects of the large-scale upzoning policy in Auckland. We find that differences in land price gradients that emerged soon after the reforms are consistent with a 18.1% increase in the dwelling floorspace capacity of the metropolitan area, based on our preferred empirical specification. In percentage terms, this increase is approximately forty percent of the increase in floorspace capacity estimated by Greenaway-McGrevy and Jones (2023) from the relaxation of FARs, and four to five times the increase in the dwellings that Greenaway-McGrevy and Phillips (2023) attribute to the policy over the first five years. Using plausible estimates of the house price elasticity of demand from the extant literature, the increase in housing supply translates to a reduction in dwelling prices of between 23.0 and 39.3%, relative to the counterfactual of no-upzoning.

Our method relies on a monocentric model in the tradition of Alonso (1964), Mills (1967) and Muth (1969), and the restrictions on capital intensity employed by Arnott and MacKinnon (1977) and Bertaud and Brueckner (2005) to model LURs. However, we develop the model further by allowing restrictions on capital intensity to vary across the disk of the city. Specifically, areas of residential land are assigned to different zones, each of which has a unique maximum restriction on the capital intensity of housing construction. This feature of the model brings the monocentric model closer to urban planning in practice, where specific areas of the city are targeted for increase density via different zones. The model is then used to examine the impact of upzoning by comparing static equilibria under different regimes of restrictions. Building on results presented in Duranton and Puga (2015), the model yields a number of identities that relate differences in land price gradients between zones with different floor area ratio (FAR) restrictions to endogenous outcomes of the model, including dwelling densities and floorspace. The identities prove very useful in a policy evaluation exercise, since measured changes in land price gradients can be used to anticipate the impacts of an upzoning policy on future development patterns. In particular, differential changes in land price gradients between upzoned and non-upzoned areas are equivalent to differences in the optimal floorspace that developers build to in each area. This result holds when floorspaces are constrained by exogenous FAR restrictions imposed by a regulator. It also holds under a wide family of household utility and floorspace production functions.

These theoretical findings can be used to empirically evaluate the impact of upzoning in practice using detailed individual-level housing data and geospatial data on residential zoning. Policy impacts on land prices for Auckland are estimated using parcel-level information on house prices that are geocoded, so each parcel can be mapped to planning zones introduced under the zoning reform, allowing us to identify whether the parcel was upzoned. We can then estimate the relative impacts of the policy on land prices using a difference-in-differences (DID) model. Because individual dwellings differ in their land endowments, we use the capital intensity of each dwelling to mediate the effect of upzoning in the empirical model, thereby allowing us to uncover the effect of upzoning on the value of the underlying land. This feature of the method means that it does not rely on data on vacant lots, which are typically sparse in urban areas.

Unfortunately DID is ill-suited to uncovering the impacts of upzoning on the absolute level of dwelling prices. Causal inference requires the Stable Unit Treatment Value Assumption (SUTVA), meaning that the treatment cannot affect outcomes in the control group. However, upzoning in one area of the city affects house prices in all areas under standard economic models of urban development, since spatial equilibrium implies that policy interventions in one area of the city will have impacts on other areas. This "spillover effect" is frequently encountered in DID applications (Imbens and Wooldridge, 2009) and violates the SUTVA condition. Thus, while the DID framework can reliably estimate *relative* changes in house and land prices, we cannot use it to infer the effects of upzoning on the *absolute* level of house prices.

We can nonetheless interpret the changes in land prices differentials through the lens of our theoretical model in order to anticipate development patterns that are consistent with these observed changes, including increases in the supply of housing. Then, using plausible estimates of price elasticity of demand for housing, we can obtain changes in house prices implied by the increase in housing supply from the reforms. Following results presented in Albouy et al. (2018), we consider elasticities between -0.5 and -0.8, resulting in implied price decreases of between 39.3 and 23.0% relative to the counterfactual of no-upzoning.

This paper makes several contributions to the extant literature. First, it contributes to the small but growing literature on upzoning, which has been an understudied topic in urban economics to date (Schill, 2005; Freeman and Schuetz, 2017), by providing theoretically-grounded, empirical framework for assessing its impacts. Second, it extends the canonical monocentric model to allow restrictions to vary across different areas of the city according to different zones. To date, the models have assumed a single restriction, such as a height restriction, apply uniformly across the city (Arnott and MacKinnon, 1977; Bertaud and Brueckner, 2005; Kulish et al., 2012). Third, it contributes to research focused on examining the specific impacts of the Auckland Unitary Plan on housing outcomes. Greenaway-McGrevy et al. (2021) study the immediate impact of the AUP on relative house prices, while Greenaway-McGrevy and Phillips (2023) empirically measure the impact of upzoning on dwelling construction over the first five years after the policy change. Our paper complements this work by producing an assessment of the long-run effects of the reform.

The remainder of the paper is organized as follows. The following section introduces the theoretical model and our approach to modeling LURs that vary according to different zones in the city disk. Section three describes our data and the institutional context and details of the upzoning policy in Auckland that we study. In section four we present our empirical model for measuring the effects of upzoning on land price gradients. In section five we interpret these results through the lens of our theoretical model to arrive at an estimate of the increase in floorspace capacity and the reduction in dwelling prices from the policy. Section six concludes.

2 Theoretical Model

The monocentric model of Alonso (1964), Mills (1967) and Muth (1969) provides the theoretical foundation for our work. Because we are interested in how regulations affect the intensity of housing development, we employ a version of the model that admits substitution between land and capital in the production of floorspace. Regulations restrict developers to a maximum capital-to-land ratio. Under constant returns to scale in housing floorspace production, this is equivalent to a restriction on the Floor to Area Ratio (FAR). We therefore refer to it as a FAR restriction. It is equivalent to the height restrictions employed by Arnott and MacKinnon (1977), Bertaud and Brueckner (2005) and Kulish et al. (2012). We provide a brief description of the standard model before describing our own additions to the set-up.

The city lies on a flat plane and is comprised of a central business district (CBD) surrounded by suburbs that house workers. The land around the CBD suitable for housing development spans θ radians. Workers reside in the suburbs and commute to the CBD to earn wages. Their preferences over housing floorspace H and a consumption numeraire C are described by a utility function $U(H, C)$ that is increasing in both arguments and strictly quasi-concave. Households living at distance $x \in [0, \infty)$ from the CBD incur a commuting cost tx to earn the wage W.[3] Under these assumptions, both the rent of housing floorspace $P(x)$ and floorspace $H(x)$ are decreasing in x and convex (Duranton and Puga, 2015).

We impose standard assumptions on the production of housing floorspace. Developers produce floorspace $H(x)$ using capital K and land L. $H(x)$ is increasing in K and L, exhibits constant returns to scale, and is strictly quasi-concave. The rental price of capital is assumed constant and is set to unity. The rental price of land is denoted $R(x)$. Developers are perfectly competitive. We define $h(x) = H(x)/L(x)$ as the FAR ratio that developers build to at x. Under constant returns to scale, $h(x)$ can be expressed as a function of the capital to land ratio $k(x) = K(x)/L(x)$.

In the Appendix we present a specification of the model based on Cobb-Douglas utility and production. This specification is unnecessary for the exposition presented here.

2.1 FAR Restrictions and Zones

Arnott and MacKinnon (1977), Bertaud and Brueckner (2005) and Kulish et al. (2012) model LURs as constraints on height. For example, $h(x) < \hat{h}$ for some $\hat{h} > 0$ that is selected by a policymaker.

[3] The AMM structure can accommodate dispersed employment across the city under the assumption wages decrease linearly at rate t as the distance between the place of work and the CBD increases (Glaeser, 2008).

Under constant returns to scale production this is equivalent to a constraint on the capital intensity of the property, $k(x)$. In the framework of Arnott and MacKinnon (1977), Bertaud and Brueckner (2005) and Kulish et al. (2012), the FAR restriction $\hat{h}$ applies uniformly across the city. Under this approach, the model is solved via the population constraint (c.f. (9) in Bertaud and Brueckner, 2005):

$$\int_0^{\hat{x}} \frac{\hat{h}}{H(x)}\theta x dx + \int_{\hat{x}}^{\bar{x}} \frac{h(x)}{H(x)}\theta x dx = N$$

and the condition that $\hat{x}$ satisfies $h(\hat{x}) = \hat{h}$, such that $\hat{x}$ denotes the distance at which $\hat{h}$ becomes binding. $\bar{x}$ denotes the radius of the city, while $N$is the number of dwellings in the city.

In practice, however, regulations that affect housing construction typically varies between different zones of a city. To bring the model closer to urban planning in practice, we extend the general framework by permitting different zones across the city disk. Each zone has a different restriction on $h(x)$. Let $\omega_j(x) \in (0,1)$ denote a continuous function in x that describes the proportion of land at distance x that is assigned to residential zone j. Let $\hat{h}_j$ denote the FAR restriction that applies in zone j, such that $h(x) \leq \hat{h_j}$. The population condition becomes

$$\theta \sum_{j=1}^{m_z} \left(\int_0^{\hat{x}_j} \frac{\hat{h}_j}{H(x)} \omega_j(x) x dx + \int_{\hat{x}_j}^{\bar{x}} \frac{h(x)}{H(x)} \omega_j(x) x dx \right) = N \tag{1}$$

where $\hat{x}_j$ satisfies $h(\hat{x}_j) = \hat{h}_j$, such that $\hat{x}_j$ denotes the distance at which $\hat{h}_j$ becomes binding.[4] The remaining conditions for solving the model are the same as in the standard AMM model. Household utility is equal at all locations and land rents at the boundary are equal to agricultural rents $\bar{R}$, i.e.

$$R(\bar{x}) = \bar{R}$$

A simple example of the framework is when $\omega_j(x) = \omega_j \in (0,1)$ for all $j = 1, \ldots, m$. In this case the city disk spanning θ radians is decomposed into circular sectors, with each sector corresponding to a zone. Such a model may be appropriate for transit-oriented zoning, whereby residential areas close to rapid transit and highway corridors are zoned for greater density,[5] or when planners take other features or locations into account when setting land use policy. Figure 1 below provides an example where $\omega_j(x) = \omega_j \in (0,1)$ for all $j = 1, \ldots, m$. However, in practice, planners often locate high density closer to downtown. In such cases we might expect $\omega_j(x)$ to be monotonically decreasing in x for zones that permit high levels of capital intensity in housing, and increasing for zones that permit low levels of capital intensity.

[4]Note that condition this does not preclude $\hat{x}_j > \bar{x}$, which indicates that the FAR restriction is binding out to the radius $\bar{x}$ in zone j.

[5]Transit corridors that allow faster commute times do however warp the topology of the city disk. See Baum-Snow (2007) or Greenaway-McGrevy and Jones (2024).

2.1.1 Discretization into Annulus Sectors

The model can also be solved via numerical integration when the functions $\{\omega_j(x)\}_{j=1}^{m_z}$ are discretized via a step function over x. Discretization maybe desirable when the zoning topography of the model is intended to match zoning topography in practice.

Discretization involves decomposing the disk into a grid of annulus sectors and assigning a zone to each sector. We assign m_s circular sectors, each of angle θ/m_s. Each annulus has an annulus radius of length l. Let (a, s) index the annulus and sector, such that $s = 1, \ldots, m_s$ indexes the sectors and $a = 1, 2, \ldots$ indexes the countably infinite annuli. We define $\omega_{j,a} \in [0, 1]$ to be the proportion of sectors within the ath annulus that are assigned to zone j.

Figure 1 exhibits four examples of the annulus sectors framework for a case where there are three different zones. Each annulus sector within the disk is assigned a zone and signified by a different color (magenta, cyan or gray). In the example in the top left panel of Figure 1, we have the case where the proportion of land assigned to each zone is constant for all distances x. In the example in the bottom right panel, the cyan zone is more prevalent close to the CBD, while the gray zone becomes more prevalent towards the outskirts of the city. Because distance to the CBD determines outcomes in the monocentric AMM model, the ordering of the zones within the ath annulus is inconsequential in the disks presented in Figure 1. Under the discretization into annulus sectors, the population condition (1) becomes more complicated. We use x_a^* to denote a discrete measure of distance to the CBD. Specifically, x_a^* is the distance of the outer edge of the ath annulus to the CBD, such that $x_a^* \in \{x_0^*, x_1^*, x_2^*, x_3^*, \ldots\} = \{0, l, 2l, 3l, \ldots\}$. Then (1) becomes

$$\theta \sum_{j=1}^{m_z} \left(\begin{array}{c} \sum_{a=1}^{a_j} \int_{x_{a-1}^*}^{x_a^*} \omega_{j,a} \frac{\hat{h}_j}{H(x)} x dx + \int_{x_{a_j}^*}^{\hat{x}_j} \omega_{j,a+1} \frac{\hat{h}_j}{H(x)} x dx + \omega_{j,a_j+1} \int_{\hat{x}_j}^{x_{a+1}^*} \frac{h(x)}{H(x)} x dx + \\ \sum_{a=a_j+2}^{\bar{a}} \omega_{j,a} \int_{x_{a-1}^*}^{x_a^*} \frac{h(x)}{H(x)} x dx + \omega_{j,\bar{a}+1} \int_{x_{\bar{a}}^*}^{\bar{x}} \frac{h(x)}{H(x)} x dx \end{array} \right) = N$$

where $a_j = \lfloor \frac{\hat{x}_j}{l} \rfloor$ and $\bar{a} = \lfloor \frac{\bar{x}}{l} \rfloor$, and where $\lfloor \cdot \rfloor$ denotes the largest integer less than or equal to the argument. Thus the FAR restriction for zone j becomes non-binding in the $(a_j + 1)$th annulus, and the radius of the city lies within the $(\bar{a} + 1)$th annulus.

2.1.2 Numerical Simulation

We present an example of urban development outcomes when spatial equilibrium holds under the annulus sectors disk depicted on the bottom left panel of Figure 1. For the purposes of this instructive exercise, we impose a FAR restriction of 3 on the magenta annulus sectors. The cyan areas have a FAR restriction of 6, while the gray areas have a FAR restriction of 1. Together with the spatial distribution of zones depicted in the bottom left panel of Figure 1, this set of FAR restrictions constitute what we will refer to as a "regulatory regime", for instructive purposes. To complete the model, we follow Bertaud and Brueckner (2005) and employ Cobb-Douglas utility and production functions. The parameters of the model are described in the notes to Figure 2. Further details on the model are contained in the Appendix. Noting that there is a direct mapping from rents to prices under a suitable capitalization rate, we hereafter refer to $P(x)$ and $R(x)$ as "house

Figure 1: Examples of Annulus Sector Zones in the Alonso-Muth-Mills model

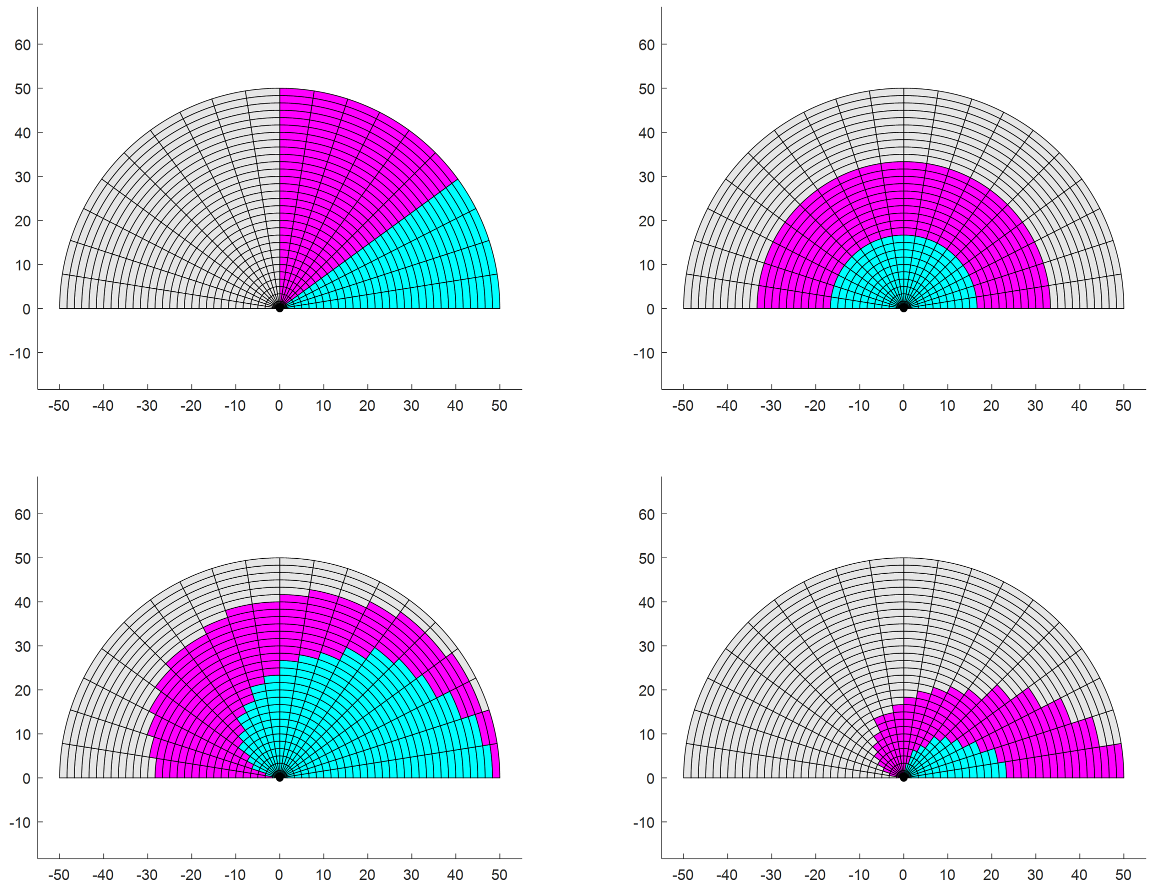

Notes: Disk decomposed into annulus sectors, $\theta = \pi$. Colors represent different zones. In the example on the left, the proportion of sectors in each annuli assigned to each zone is constant. In the example on the right, the proportion of annuli assigned to each zone is dependent on the distance of the annuli to the CBD. Housing demand is independent of the angular displacement of any annulus sector from any ray from the origin. Therefore the ordering of zones within a given annulus has no impact on model outcomes.

prices" and "land prices" (Kulish et al., 2012).

Figure 2 depicts outcome variables of interest under spatial equilibrium. It depicts dwelling prices $P(x)$ and floorspaces $H(x)$, which depend only on x and $\bar{x}$, as well as land prices $R(x)$, FARs $h(x)$, and densities $1/L(x)$, all of which vary according to zone. To illustrate how liberal zoning lowers land input costs to housing production, Figure 2 also depict total land costs, $R(x)L(x)$, and the land cost per unit of floor space, $R(x)L(x)H^{-1}(x) = R(x)h^{-1}(x)$ (i.e., the land cost to floor area ratio). We compare the set of urban development outcomes to a different regulatory regime in which the tightest FAR restriction ($= 1$) holds in all annulus sectors of the city.[6] The latter outcomes are depicted by the dashed blue lines in Figure 2. We hold population constant, meaning that the comparison is made under closed city assumptions. We describe this regulatory regime as "more restrictive", and compare its equilibrium outcomes to those of the "less restrictive" regulatory regime that has the three different different zones with FAR restrictions no less than those of the more restrictive regime.

First, we examine outcomes within the less restrictive regulatory regime. Dwelling prices $P(x)$ and floorspace $H(x)$ gradients depend on distance from the CBD but not the underlying FAR restriction of the zone in which the dwelling is located – i.e., floorspace and house prices at a given distance x are the same in each zone. This follows from the equilibrium condition that utility is equal at all locations, and that transport costs are equal at distance x, which ensures that income after commuting costs is the same at all distances x. This does not, however, mean that the geographic distribution of zones and FAR restrictions have no impact dwelling price and floorspace gradients. Rather, FAR restrictions impact house prices and floorspace gradients through their impact on city radius $\bar{x}$ in equilibrium. Dwelling prices and floorspaces are determined by the radius of the city, $\bar{x}$, wages W, and travel costs t. Holding all else constant, tighter FAR restrictions increase $\bar{x}$, resulting in higher house prices and lower floor areas at each distance $x \in [0, \bar{x})$ from the CBD.

The remaining variables – including land prices, FARs, and dwelling densities – do differ with zone. Land prices are lower in zones with tighter restrictions on capital intensity at distances where the restriction is binding (i.e. for $x < \hat{x}_j$). Zones with more restrictive limits on FARs reduce land prices because they restrict the marginal product of land. In locations where the restriction is non-binding, so that developers build to a FAR that is less than the restriction imposed, land prices are equivalent between the zones. FARs are clearly constrained by the upper bound imposed in each zone out to a certain distance from the CBD, after which the optimal FAR falls. We refer to this point as the distance at which the constraint binds. The distance at which the constraint binds is decreasing in the upper bound on the FAR. Thus we have $\hat{x}_{j+1} < \hat{x}$, $j = 1, \ldots, m_z$, where recall that zones are ordered from least to most permissive in terms of the FAR restriction. Dwelling densities are higher in zones with more permissive FARs.

Next, we compare equilibrium outcomes in the less restrictive regulatory regime to those of the more restrictive regime. The more restrictive regime has higher house prices and smaller floorspaces compared to the less restrictive regime. It is also more expansive, covering a radius of 53 km from

[6] This regulatory regime is analogous to the simulations of height limits studied by Arnott and MacKinnon (1977) and Bertaud and Brueckner (2005).

Figure 2: Zoning Reform in the AMM model with Annulus Sectors

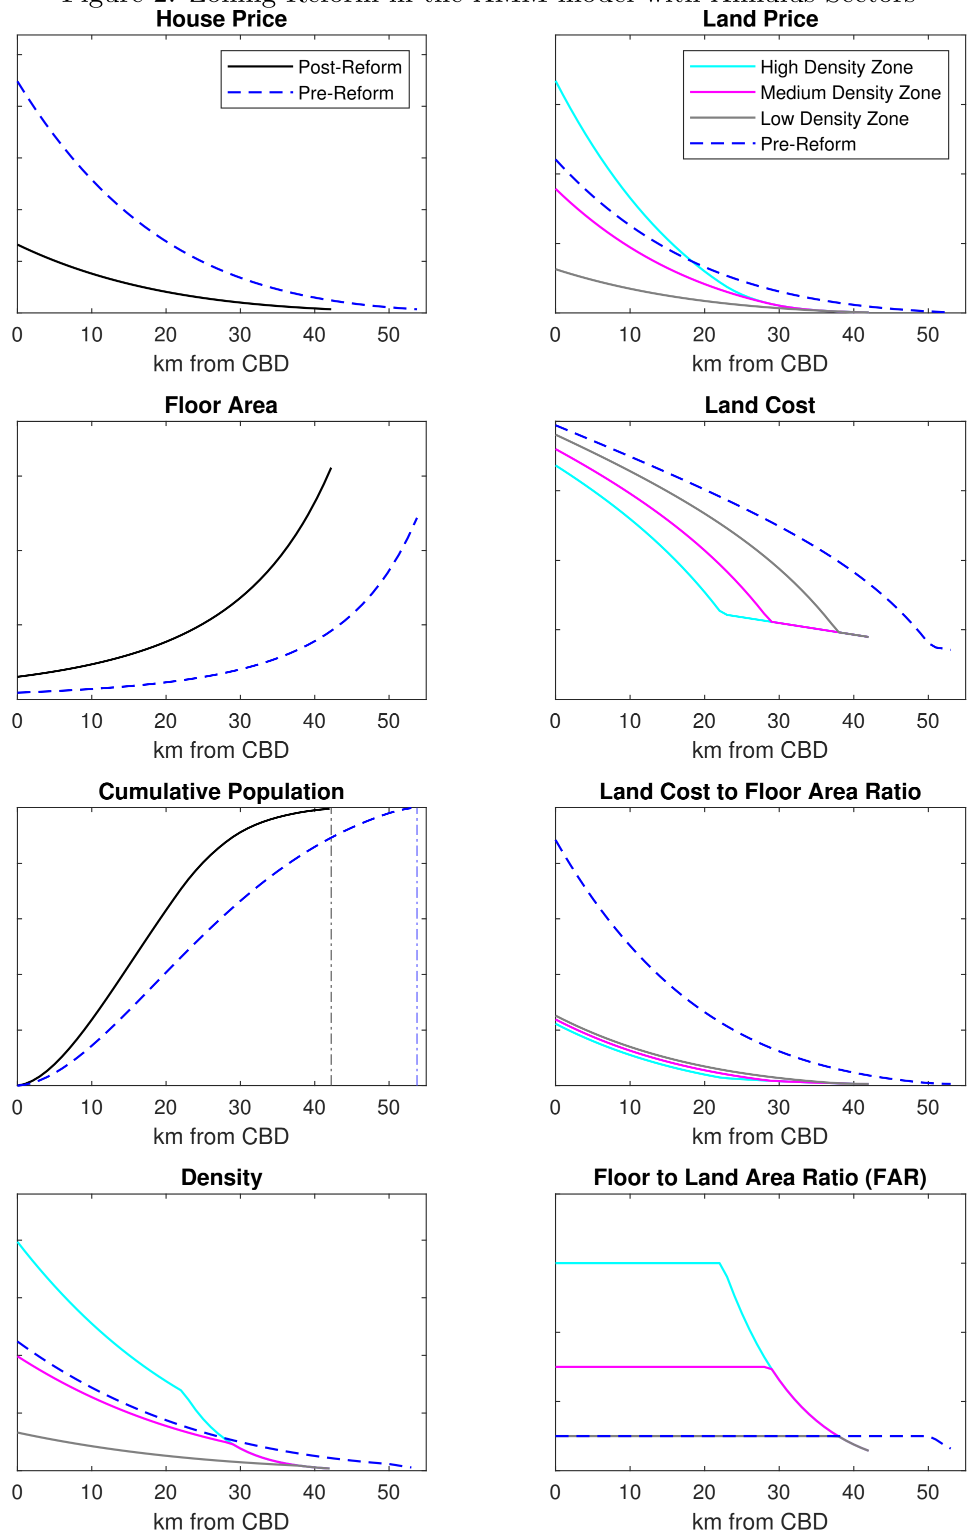

Simulated equilibrium outcomes using the annulus sectors depicted in the bottom left panel of Figure 1. High density FAR restriction = 6; Medium density FAR restriction = 3; Low density FAR restriction = 1. The dashed blue line are equilibrium outcomes when the entire city is zoned to low density (FAR restriction = 1). Population is constant under both regulatory regimes. Utility is Cobb Douglas with a housing share of 0.2. Floorspace production is Cobb-Douglas with a capital share of 0.6 and TFP of 5e^{-4}. These parameters are adopted from Kulish et al. (2012). $W = \$100,000$, $t = \$1{,}046$, $\bar{R} = \$80,000$, $N = 1,000,000$ and $\theta = \pi$. The vertical dashed lines in the panel labeled 'Cumulative Population' denote the radius of the city (i.e. the distance from the CBD at which the cumulative population is N).

the CBD, whereas the radius of the high density city is 47.8 km. These results reflect those already established in the literature for uniform restrictions on capital intensity. For example, tighter height restrictions generate larger radius, higher house prices and smaller floorspaces compared to cities with looser height restrictions (Arnott and MacKinnon, 1977; Bertaud and Brueckner, 2005; Kulish et al., 2012).

Within a given radius of the CBD, land prices are higher in the least restrictive zone (FAR = 6) compared to land prices in the more restrictive regulatory regime. Meanwhile land prices in the zone with the most restrictive FAR constraint (= 1) are lower under the less restrictive zoning regime. These observations reflect two effects that work in opposite directions. First, the marginal product of land is increasing in the FAR restriction (i.e. as a FAR restriction is relaxed, the marginal product of the affected parcels increases). At a given distance x, differences in the price of land between zones within the less restrictive zoning regime reflect differences in marginal product. Second, as FARs restrictions are relaxed, city radius decreases (when population is held constant) as households are able to located closer to the CBD. This reduces house prices and consequently land prices. Differences in land prices in the zone with the same FAR restriction (= 1) under both regimes reflect only the latter effect since the marginal product of land is the same.[7] The medium restriction zone (FAR = 3) has lower land prices than under the more restrictive regime, indicating that the productivity effect is comparatively smaller in this example.

Land costs (as opposed to prices) are lower in all three zones under the less restrictive regime, and, among the three zones, are lowest in the high density zone. Since floor areas are the same in all three zones (at each distance x), this also implies that land price per unit of floor area is lowest in the high density zone. Thus, while zoning has an ambiguous effect on the absolute level of land prices, it has an unambiguous effect on land costs. This indicates that land costs (i.e., the cost of land necessary to produce a dwelling or a given amount of floorspace) should be used to evaluate the effect of upzoning on housing costs – not land prices.

2.2 Useful Results

The AMM model with zoned FAR restrictions yields several notable identities that allow us to quantify the magnitude of differences in outcome variables resultant from a change in regulatory regime. As discussed above, a regulatory regime refers to a specific spatial distribution of zones and associated FARs. Outcome variables in each regime are given by the spatial equilibrium conditions.

We describe each identity through the propositions stated below.

Proposition 1. *Differences in land price gradients between different zones within the same regulatory regime reflect differences in the optimal FAR that developers build to in each zone.*

To see why this is the case, under standard conditions, the ratio of the land price gradient to

[7] A common alternative assumption in comparative statics is to hold utility constant (i.e. the open city assumption). Under this assumption city radius is unchanged, meaning that that land price differentials between different zoning regimes only reflect differences in marginal productivity of land.

house price gradient is equal to the FAR (see, e.g., Duranton and Puga, 2015). That is,

$$\frac{\frac{dR(x)}{dx}}{\frac{dP(x)}{dx}} = h(x) \tag{2}$$

Condition (2) follows from the previously stated assumptions on the production function and competition in housing markets. It therefore extends to cases in which FARs are restricted.

Condition (2) also holds in cases where FAR restrictions vary across different zones of the city. Because dwelling prices $P(x)$ not not vary by zone, (2) implies that for any two different zones $i, j \in \{1, \ldots, m_z\}$ we have

$$\frac{\frac{dR_j(x)}{dx}}{\frac{dR_i(x)}{dx}} = \frac{h_j(x)}{h_i(x)} \tag{3}$$

Thus the ratio of land price gradients between different zones is equal to the ratio of FARs in each zone. Relative differences in land price gradients in different zones tell us about the relative capital intensity that it is optimal for developers to build to. This insight informs our empirical strategy described in more detail in Section 5 below. Because these parameters are given by *relative* differentials in land prices, rather than absolute levels, they can be identified using a difference-in-differences empirical strategy, which identifies relative differences between treated (upzoned) and control (non-upzoned) groups.

Proposition 2. *Differences in land price gradients between different regimes and between different zones reflect differences in dwelling densities.*

To see why this is the case, combining (2) with the Alonso-Muth condition (see (4) in Duranton and Puga 2015) yields

$$\frac{dR(x)}{dx} = -t \cdot \frac{h(x)}{H(x)} \tag{4}$$

where recall that $\frac{h(x)}{H(x)} = \frac{1}{L(x)}$is dwelling density. The condition holds across different spatial equilibria and thus across different zoning regimes. Let $R_{j,A}(x)$ denote land prices in zone j under zoning regime A and let $R_{i,B}(x)$ denote land prices in zone i under regime B. Because travel costs t are constant, and floorspace $H(x)$ is the same across different zones within the same regulatory regime, the above condition implies that

$$\frac{\frac{dR_{j,A}(x)}{dx}}{\frac{dR_{i,B}(x)}{dx}} = \frac{h_{j,A}(x)}{h_{i,B}(x)} \cdot \frac{H_B(x)}{H_A(x)} \tag{5}$$

which means that differences in land price gradients across different regimes and between different zones reflect differences in dwelling densities $h(x)/H(x)$.

Proposition 3. *Differences in land price gradients between different regulatory regimes in zones with identical FAR restrictions reflect differences in floorspace at distances where the FAR* *restrictions are binding.*

This result is an extension of (5). The ratio $\frac{h_{j,A}(x)}{h_{i,B}(x)}$ is fixed (and independent of x) for all distances x at which the FAR restrictions that apply in i and j are binding. That is, $\frac{h_{j,A}(x)}{h_{i,B}(x)} = \frac{\hat{h}_{j,A}}{\hat{h}_{i,B}}$ for all $x \leq \min(\hat{x}_{j,A}, \hat{x}_{i,B})$. That is,

$$\frac{\frac{dR_{j,A}(x)}{dx}}{\frac{dR_{i,B}(x)}{dx}} = \frac{\hat{h}_{j,A}}{\hat{h}_{i,B}} \cdot \frac{H_B(x)}{H_A(x)}, \quad x \in [0, \min(\hat{x}_{j,A}, \hat{x}_{i,B})] \tag{6}$$

This condition also holds when zone j and i are subject to the same FAR restriction, so that $j = i$ and $\hat{h}_{j,A} = \hat{h}_{i,B}$. Then we have

$$\frac{\frac{dR_{j,A}(x)}{dx}}{\frac{dR_{i,B}(x)}{dx}} = \frac{H_B(x)}{H_A(x)}, \quad x \in [0, \min(\hat{x}_{j,A}, \hat{x}_{i,B})] \tag{7}$$

Thus differences in land price gradients between different regimes in zones that have the same FAR restriction reflect differences in floorspace between the regimes.

The differences in floorspaces under each regime can be mapped to differences in house prices given a specific utility function. For example, under Cobb-Douglas utility, the consumption share of housing in the consumption basket is constant, meaning that

$$\frac{P_A(x)}{P_B(x)} = \frac{H_B(x)}{H_A(x)}$$

and thus (7) implies that

$$\frac{\frac{dR_{j,A}(x)}{dx}}{\frac{dR_{i,B}(x)}{dx}} = \frac{P_A(x)}{P_B(x)}, \quad x \in [0, \min(\hat{x}_{j,A}, \hat{x}_{i,B})]$$

Thus differences in land price gradients between zones that have the same FAR restrictions (i.e. $\hat{h}_j = \hat{h}_i$) will reflect differences in dwelling prices, provided that LURs are binding in those locations. Moreover, as shown in the Appendix, under Cobb Douglas utility, $P(x) = B(W - tx)^{\frac{1}{\alpha}}$ for a constant B that is independent of x, which means that the ratio $\frac{P_A(x)}{P_B(x)}$ is constant for all x. Thus, accurate measures of the change in land price gradients in non-upzoned areas after the implementation of an upzoning policy would tell us about the impacts of the policy on house prices. However, it is difficult to find an accurate counterfactual for changes over time, and thus we do not utilize this result in our empirical analysis. Moreover, we do not want our results to hinge on any particular utility function.

3 Data and Institutional Background

In this section we describe the data employed in the empirics of the paper. In order to contextualize the data, we first describe the institutional background underpinning the upzoning policy intervention in Auckland that motivates the empirical analysis.

3.1 Institutional Background

Auckland is the largest city in New Zealand with a population of approximately 1.57 million within the greater metropolitan region (as of 2018 census). 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. Centered on a long isthmus of land between two harbors, 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 LURs that would potentially change restrictions on the extent of site development, 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 incorporate these into their expectations prior to the policy becoming operational. See Greenaway-McGrevy and Jones (2023) for further details on the institutional background.

The amount of development permitted on a given site is restricted by the residential planning zone in which the site is located. For the purposes of estimating changes in land price differentials, 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. These four residential zones (SH, MHS, MHU and THA) comprise the vast majority of residential land within the urban core of Auckland. Table 6 in the Appendix summarizes the various LURs for each of the residential zones. These regulations include site coverage ratios, minimum lot sizes for new subdivisions, and height restrictions, among others. For example, five 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.

Figure 4 depicts the geographic distribution of four zones across the city. For clarity we zoom in on the central urban area of Auckland. Evidently MHS covers the most area, closely followed by MHU. SH is predominantly located either very close to the Auckland city CBD or at the outskirts of the city. THA covers the least amount of area.

Prior to the AUP, there were seven different plans that applied to different regions, which were

Figure 3: Auckland Region

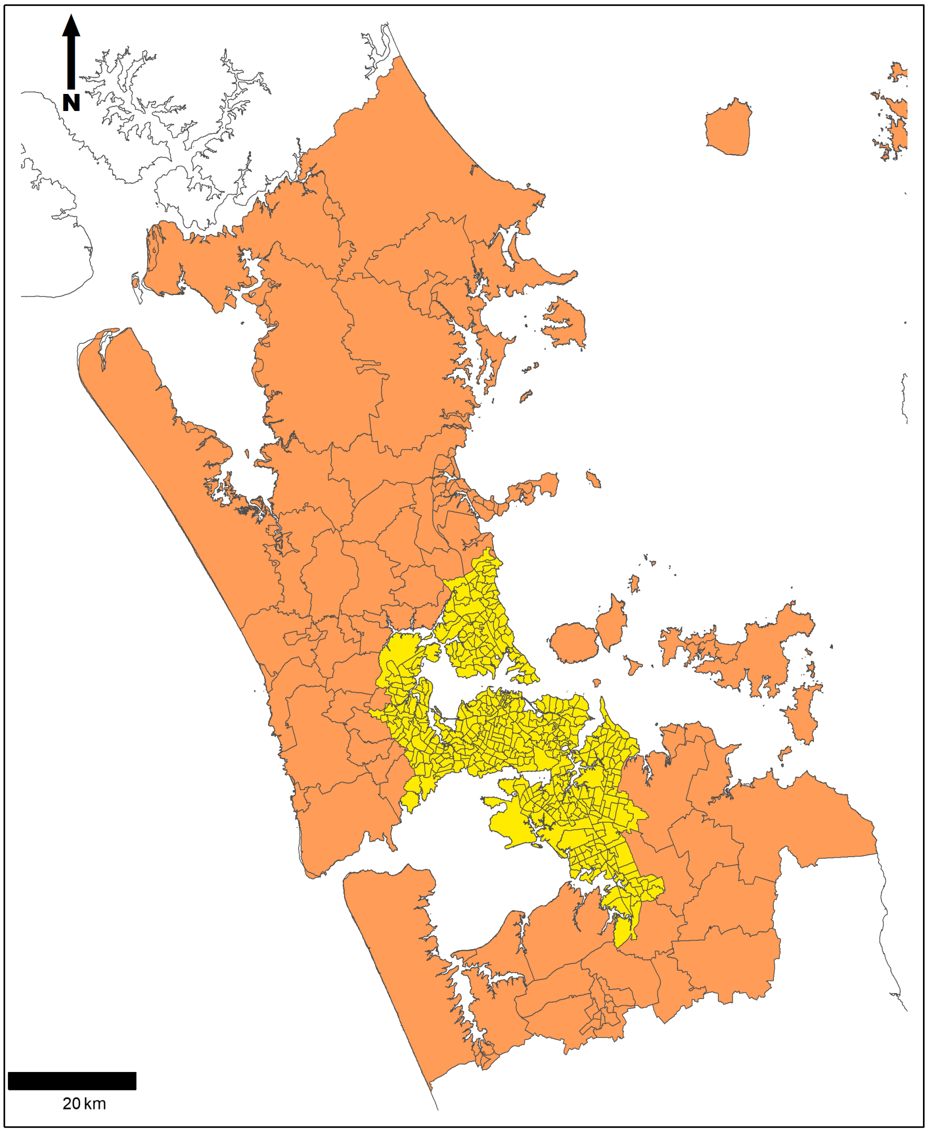

Notes: Auckland region (shaded) decomposed into Statistical Areas. Major urban area in yellow. Source: Greenaway-McGrevy and Phillips (2023).

Figure 4: Residential Zones of inner Auckland introduced under the Auckland Unitary Plan

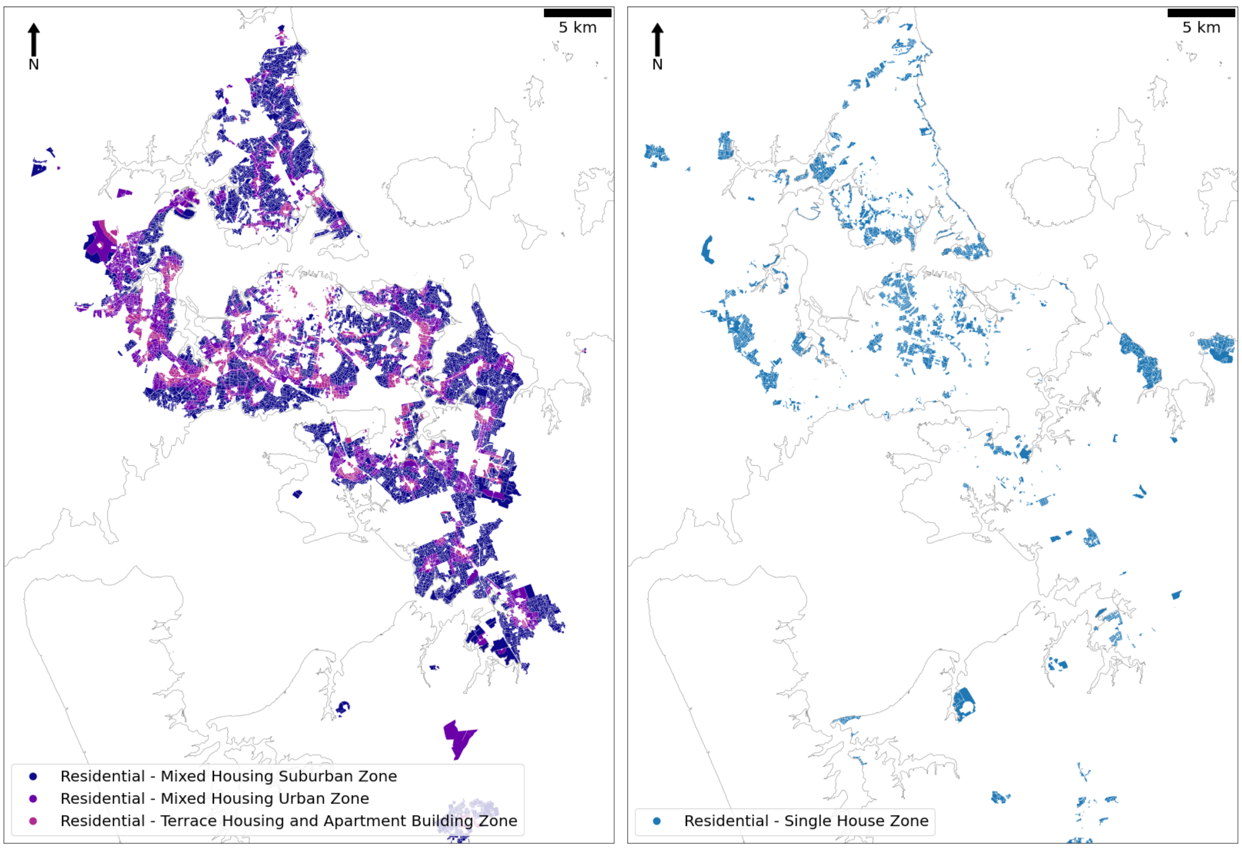

Source: Greenaway-McGrevy and Phillips (2023).

relics of the seven distinct and city councils that composed Auckland until 2010. Greenaway-McGrevy and Jones (2023) show that the vast majority – 98.7% – of land zoned as THA, MHU and MHS was upzoned in the sense that the FAR implied by site coverage and height restrictions increased. Nonetheless, using the same method as Greenaway-McGrevy and Jones (2023), we allocate these parcels to the non-upzoned group, and exclude any parcels that were rezoned to SH from semi-rural or rural. Between 75.1 and 78% of residential land was upzoned, in the sense that the FAR implied by site coverage and height restrictions was relaxed (Greenaway-McGrevy and Jones, 2023).

Our empirical design exploits differentials in land prices between upzoned and non-upzoned after the policy is announced. To model these differentials we adopt a DID framework, wherein the SH zone is used as the "control", while the MHS, MHU and THA zones are three different "treatment" groups. As discussed in the introduction, we note that outcomes in the control group do not represent a valid counterfactual for causal inference due to spillovers. Our empirical strategy requires accurate measures of *relative* differentials, and does not require these to be interpreted as treatment effects.

3.2 Data

Our dataset is based on residential dwelling valuations provided by the Auckland Council for the years 2011, 2014, and 2017. Ratings valuations are conducted every three years for the purposes of levying local government property taxes. The data include assessments of the value of built capital (called the "improvements value") and the land value of the property.[8] These are used to construct a measure of capital intensity of the dwelling structure that is employed in our empirical work.

Valuations are matched to individual land parcels obtained from a GIS database of land records. Valuations are matched to parcels based on the certificate of title of the valuation, which is a unique identifier that is also used in the GIS database on land parcels. This matching enables us to assign a geocoordinate to each dwelling in the dataset. Based on this geocoordinate, we can identify the residential zone of each dwelling. This information is used in our empirical work. We can also calculate the distance of the dwelling to various points of interest, including the CBD.

In our empirical work we use two measures of distance. First, we use Haversine (or straight line) distance. Second, we use Manhattan distance, which is based on road networks. Each measure comes with drawbacks. Because the coastal topography of Auckland is characterized by large bodies of water that mean that commuting routes are often not along straight lines. Haversine distance understates distances for commuting routes that must circumvent these large bodies of water. In addition, Auckland is characterized by three highways emanating from the CBD to the North, South and West, as well as several highways that link these arterials. The network offers faster commute times for households served by these highways. For such households, Manhattan distance overstates commuting times. Moreover, these highways often traverse the large bodies of water, meaning that highway placement is in part and endogenous response to the coastal topography.

[8]See Cooper and Namit (2021) for a discussion of the valuation methodology.

Estimation of the increase in floorspace from upzoning incorporates information on the proportion of land assigned to the different residential zones under the AUP. Estimation of these proportions is based on Statistical Areas (SAs), which are geographic areas designed and employed by Statistics New Zealand for the collection and organization of census data. SAs have a target population of between 2000-4000 people in cities (such as Auckland) and were drawn to reflect communities that interact socially and economically. They are analogous to census tracts in the US. We use the 2018 vintage of the SAs, of which there are 510 SAs in Auckland.[9]

In order to estimate aggregate increases in floorspace from upzoning, we also require information on the FARs of the existing dwelling stock prior to upzoning. Estimates of dwelling FARs are based on the valuation dataset, which report floor areas for each dwelling. Dwellings with exclusive ownership of land in the title (i.e. not joint ownership of land in the title) have a reported land area in the valuation dataset. Dwellings with joint ownership of land in the title do not have reported land areas. To address this issue, we join the valuation dataset with GIS information on individual parcels in Auckland. Because we have the geocoordinates of each dwelling (in addition to other locational information, such as address), we can identify the parcel of land associated with each title with a high degree of accuracy. Additional details on the joining algorithm are contained in the Appendix. Once joined, we can identify many-to-one matches, thereby identifying all other dwellings in the valuation dataset that share the same land parcel Based on this we can calculate the FAR for the entire parcel, which is then assigned to each dwelling associated with the parcel.

4 Empirics

In this section we examine the relative effects of upzoning on land prices using a multiperiod DID framework that compares outcomes in upzoned areas to non-upzoned areas. These results are then used in the following section to impute changes in housing supply.

Our model is informed by two insights from our theoretical framework described above. First, upzoning increases land prices relative to non-upzoned parcels. The impact of upzoning on the price of existing housing will therefore depend on land endowment of each property, which varies significantly in many real-world cities: Housing that is land intensive will appreciate in value relative to capital intensive housing. Our specification conditions on a measure of the capital intensity of housing when modeling the effects of upzoning of house prices, enabling us to easily obtain an estimate relative changes in land prices from the fitted model.

Second, differences in land price gradients between different zones reflect differences in optimal FARs (see condition (3) above). We adopt an empirical specification that is tractable and allows us to easily recover estimates of the differences in land price gradients between different zones from

[9]SAs were introduced in 2018, as the previous classification system had not been revised since 1992. The previous statistical geographies no longer reflect current land use and population patterns. The revision was also implemented in order to align the geographic unit standards with international best practice. Population data from the previous census (conducted in 2013) and associated projections were used in the design of the 2018 boundaries. For additional details, see https://www.stats.govt.nz/assets/Uploads/Retirement-of-archive-website-project-files/Methods/Statistical-standard-for-geographic-areas-2018/statistical-standard-for-geographic-areas-2018.pdf

the fitted model. These estimates are then used in the policy evaluation exercise in the following section.

4.1 Econometric Model Specification

Let $p_{i,t}$ denote price of house $i = 1, \ldots, n$ in period t, measured on a per square meter of flooring basis.Next, let $z_{i,t}$ denote the capital intensity of house i in period t. We use the site intensity ratio, which is the ratio of the capital value of the property to the total value (Clapp and Salavei, 2010; Clapp et al., 2012). In this case $z_{i,t} \in [0,1)$, where that case of $z_{i,t} = 0$ indicates that the parcel is equivalent to vacant land. Houses are mapped to residential zones using the longitude and latitude of their matched parcel, allowing us to identify which properties upzoned. We use $j = 0, 1, 2, 3$ to index the upzoned (treatment) and non-upzoned (control) groups. $j = 0$ indicates non-upzoned areas. Meanwhile $j = 1, 2, 3$ signify the upzoned groups ordered in increasing amounts of site development allowed (so $j = 1$ corresponds to MHS and $j = 3$ corresponds to THA).

Next we must determine the treatment date. We have data for 2011, 2014 and 2017, and use 2014 as the treatment date. Although this is a year after the announcement of the preliminary AUP, it is two years before the final version of the plan is announced. Moreover, as we explain below, it is apparent that there is only a very small pre-treatment trend between 2011 and 2014, indicating that there is not much of an announcement effect between 2013 and 2014.

Our empirical model is then

$$p_{i,t} = \delta_0' Z_i + \sum\nolimits_{j=1}^{m} \mathbf{1}_{i \in j} \beta_{0,j}' X_i + \sum\nolimits_{s=-\underline{T}, s \neq 0}^{\overline{T}} \mathbf{1}_{s=t} \delta_s' Z_i + \sum\nolimits_{s=-\underline{T}, s \neq 0}^{\overline{T}} \sum\nolimits_{j=1}^{m} \mathbf{1}_{s=t, i \in j} \beta_{s,j}' X_i + \varepsilon_{i,t}, \quad (8)$$

where $X_i = (1, z_i)'$ and $Z_i = (1, z_i, x_i, z_i x_i)'$, and where x_i denotes the (natural log of) distance to the CBD, and $z_i = z_{i,-1}$, i.e. site intensity prior to the policy implementation. The period index $t = -\underline{T}, \ldots, 0, \ldots, \overline{T}$, 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 have data for 2011, 2014 and 2017, and elect to use 2014 as the treatment date. Although this is a year after the announcement of the preliminary AUP in 2013, it is two years before the final version of the plan is announced in 2016. We therefore have one observation pre- and post-treatment, i.e.$\underline{T} = \overline{T} = 1$.

The empirical model permits the land gradients to change in each period, while zoning effects on house prices conditional on capital intensity is captured by the parameters in the vector $\{\beta_{s,j}\}_{s=1}^{\overline{T}}$. Let $\beta_{s,j} = \left(\beta_{s,j}^{(1)}, \beta_{s,j}^{(2)}\right)'$. Then, for houses located in zone j, the treatment effect in $t = 1$ on log prices conditional on a capital intensity measure of non-negative support z in is given by the linear function $\beta_{1,j}^{(1)} + \beta_{1,j}^{(2)} z$, where we anticipate $\beta_{1,j}^{(2)} < 0$and $\beta_{1,j}^{(1)} > 0$. For a house with zero measure of capital intensity (i.e. equivalent to vacant land), the treatment effect is $\beta_{1,j}^{(1)}$.

In fact, these additive constants for log land values $\left\{\beta_{s,j}^{(1)}\right\}_{j=1}^{3}$ capture differences in land price gradients for prices in levels. To see why, note that $p_{i,t}$ are logged prices, and that the model implies that (expected) land prices in period $t = 1$ are given by setting $z = 0$ so that

$$R_{1,j}(x) = e^{(\beta_{0,j}+\beta_{1,j})' X_i^0} \cdot e^{(\delta_0+\delta_1)' Z_i^0} = e^{(\beta_{0,j}+\beta_{1,j})' X_i^0} \cdot R_{1,0}(x)$$

where $Z_i^0 := (1, 0, x_i, 0)$ and $X_i^0 = (1, 0)$. Then the treatment effect $\beta_{t,j}^{(1)}$ for $t > 0$ is the change in the land price gradient differential $\ln(R_{t,j}(x)) - \ln(R_{t,0}(x))$ between $t = 0$ and t. To see this, we momentarily abstract from the time period notation, and note that

$$\frac{\frac{dR_j(x)}{dx}}{\frac{dR_0(x)}{dx}} = \frac{\frac{dR_j(x)}{dx}\frac{x}{R_j(x)}}{\frac{dR_0(x)}{dx}\frac{x}{R_0(x)}} \times \frac{R_j(x)}{R_0(x)}$$

Then taking logs we have

$$\begin{aligned} \ln\left(\frac{dR_j(x)}{dx}\right) - \ln\left(\frac{dR_0(x)}{dx}\right) &= \ln\left(\frac{dR_j(x)}{dx}\frac{x}{R_j(x)}\right) - \ln\left(\frac{dR_0(x)}{dx}\frac{x}{R_0(x)}\right) + \ln(R_j(x)) - \ln(R_0(x)) \\ &= \ln(R_j(x)) - \ln(R_0(x)) \end{aligned}$$

where the second equality follows from the fact that the elasticity of the log land price gradient in zone j compared to zone 0 with respect to distance is constant in the empirical model, i.e. $\ln\left(\frac{dR_j(x)}{dx}\frac{x}{R_j(x)}\right) = \ln\left(\frac{dR_0(x)}{dx}\frac{x}{R_0(x)}\right)$. Re-introducing the time period notation, and examining the *change* in land price gradient differentials, we have

$$\begin{aligned} & \ln\left(\frac{dR_{t,j}(x)}{dx}\right) - \ln\left(\frac{dR_{0,j}(x)}{dx}\right) - \left[\ln\left(\frac{dR_{t,0}(x)}{dx}\right) - \ln\left(\frac{dR_{0,0}(x)}{dx}\right)\right] \\ = \; & \ln(R_{t,j}(x)) - \ln(R_{0,j}(x)) - [\ln(R_{t,0}(x)) - \ln(R_{0,0}(x))] \end{aligned}$$

Substituting in (8) and taking expectations we have

$$\mathrm{E}\left[\ln\left(\frac{dR_{t,j}(x)}{dx}\right) - \ln\left(\frac{dR_{0,j}(x)}{dx}\right) - \left[\ln\left(\frac{dR_{t,0}(x)}{dx}\right) - \ln\left(\frac{dR_{0,0}(x)}{dx}\right)\right]\right] = \beta_{t,j}^{(1)} \tag{9}$$

Thus the treatment effect $\beta_{t,j}^{(1)}$ provides an estimate of the *change* in difference between the land price gradient in zone j (upzoned areas) and the land price gradient in zone 0 (non-upzoned areas).

4.2 Results

In fitting the model, we exclude dwellings with land area in excess of 10 hectares, floor areas in excess of 2000m^2, and site coverage areas in excess of 1000m^2. These filters guard against incorrectly coded data and outliers impacting our findings.

Figure 5 plots $\hat{\beta}_{1,j}^{(1)} + \hat{\beta}_{1,j}^{(2)} z$ against $z \in [0, 1]$ for $j = 1, 2, 3$. These are the zoning effects on log prices, conditional on site intensity. Note that in each case, the linear function in z is positive for $z = 0$ and decreasing in z, confirming our prediction that the impact of upzoning on relative land values is increasing in land intensity (or, equivalently, decreasing in capital intensity). Interestingly, the empirical model also implies that dwellings that exceeded a threshold level of site intensity depreciated in relative value, at least in the MHS and THA zones. For the THA zone,

Figure 5: Changes in Land Gradient Differentials from Upzoning

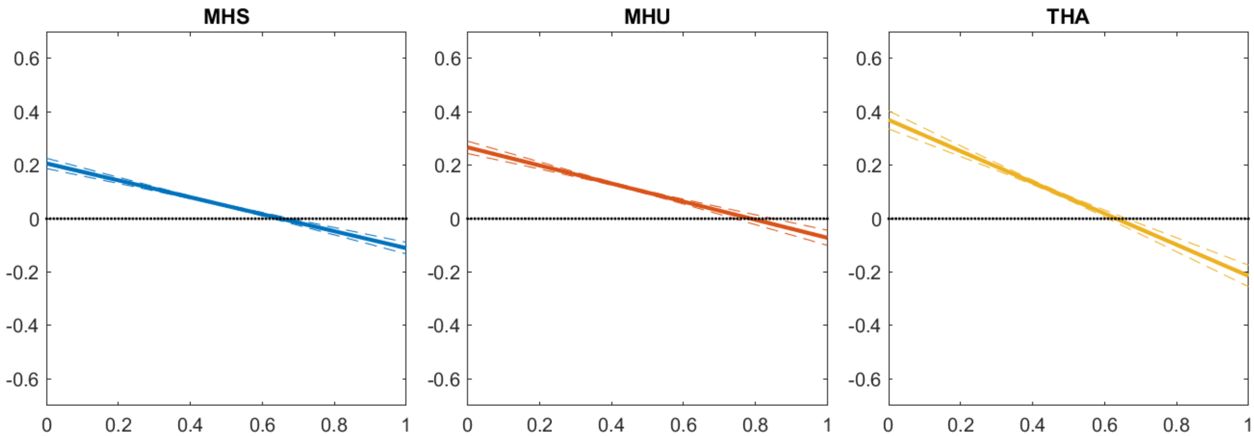

Notes: y-axis is the change in log dwelling price in upzoned areas (MHS, MHU or THA) relative to non-upzoned areas conditional on site intensity of the property (x-axis). White (1980) standard errors.

this occurred at a site intensity of approximately 0.6.

Evaluating the linear functions at $z = 0$ yields estimates of the relative impact of upzoning on log land prices. The empirical estimates indicates that upzoning increased log land prices by 0.369 on average in upzoned to THA compared to non-upzoned, 0.265 in MHU, and 0.205 in MHS. The magnitude of these effects correspond to the ordinal ranking of the amount of intensification permitted in each zone. Under (9) these estimates correspond to the change in (non-logged) land price gradients in each of the three zones.

In the Appendix, we plot the same function in the pre-treatment period in Figure 7, showing that the function is near flat. The point estimates are slightly below zero for $z = 0$ for all zones, indicating a a slight, albeit negligible upward trend in land prices prior to treatment. This is unsurprising as our treatment date is 2014, one year after the first announcement of the policy. This means that our estimates of the increase in land price gradient – and thus floorspace capacity – are conservative. We would find larger effects if we used 2011 as the treatment date.

There is a notable discrepancy between the empirical model and the theoretical model that is worth commenting on before proceeding. The theoretical model implies that gradient differentials between zone j and all zones $k > j$ are only constant for distances $x < \hat{x}_j$, where recall that $\hat{x}_j$ denotes the distance at which the FAR restriction no longer binds. In contrast, the empirical model imposes constancy for all x. To examine the extent to which this restriction in the empirical model impacted our results, we estimated a model specification of the form

$$p_{i,t} = \delta_0' Z_i + \sum_{s=-\underline{T}, s\neq 0}^{\overline{T}} \mathbf{1}_{s=t} \delta_s' Z_i + \sum_{j=1}^{m} \mathbf{1}_{i\in j, x_i < \hat{x}_j} \beta_{0,j}' X_i + \sum_{s=-\underline{T}, s\neq 0}^{\overline{T}} \sum_{j=1}^{m} \mathbf{1}_{s=t, i\in j, x_i<\hat{x}_j} \beta_{s,j}' X_i + \varepsilon_{i,t}$$

where $x_i < \hat{x}_j$ in the indicator function adds an additional condition for the treatment indicator to be one. Namely, that the parcel is less than $\hat{x}_j$ from the CBD. Here $\{\hat{x}_j\}_{j=1}^m$ are additional constraints that are estimated by minimizing the OLS function. However, we found the fitting the model to the data yields estimates of $\{\hat{x}_j\}_{j=1}^m$ that are very near or else correspond to the furthermost parcel within each zone. Thus our original restriction on the empirical model appears to have little to no impact on our findings.

In other applications, where the FAR restriction is non-binding over a large proportion of the city, more flexible regression functions can be applied to model land price gradients.

4.2.1 Alternative Specifications

We also consider different specifications to (8) and estimation methods to test the robustness of our results.

First, we consider using Manhattan distance to the CBD instead of Haversine distance. This approach uses extant road networks to identify the shortest route to the CBD from the land parcel.

Second, we augment the model (8) with a set of individual-parcel control variables, interacting the control variables with the period fixed effects so that the coefficients on the controls can vary over time:

$$\begin{aligned} p_{i,t} \quad = \quad & \delta_0' Z_i + \sum\nolimits_{j=1}^{m} \mathbf{1}_{i\in j} \beta_{0,j}' X_i + \sum\nolimits_{s=-\underline{T}, s\neq 0}^{\overline{T}} \mathbf{1}_{s=t} \delta_s' Z_i + \sum\nolimits_{s=-\underline{T}, s\neq 0}^{\overline{T}} \sum\nolimits_{j=1}^{m} \mathbf{1}_{s=t, i\in j} \beta_{s,j}' X_i + \\ & \gamma_0' W_i + \sum\nolimits_{s=-\underline{T}, s\neq 0}^{\overline{T}} \mathbf{1}_{s=t} \gamma_s' W_i + \varepsilon_{i,t}, \end{aligned} \tag{10}$$

where W_i denotes the vector of controls. We select controls based on locational factors that may influence selection of a given parcel into treatment (i.e. upzoning). Residential areas close to main arterials and transportation network access points were targeted for upzoning. Changes in transport network congestion between 2014 and 2017 may therefore manifest as a differential price impact between treatment and controls areas as many upzoned areas tend to be proximate to transportation network access points. In addition, many of the inner suburbs of Auckland were not upzoned and their residential buildings were protected under character preservation provisions. Our set of controls includes dummy variables for parcels that are within 1km of motorway access points, 1km of public rapid transit stops, and 4km of the CBD. Each variable is standardized prior to being included in the regression.

Third, we estimate the model with dwelling-level fixed effects to account for time-invariant confounding factors. This entails replacing the first two terms on the right hand side on (8) with an individual level fixed effect. Including fixed effects in the regression does not change the point estimates of $\beta_{s,j}$ for $s \neq 0$, because these parameters capture time-differences and the fixed effects are time invariant. However, the fixed effects tighten the standard errors considerably.

Finally, to account for potential measurement error in site intensity, we estimate (8) using the FAR of the dwelling as an instrument. This should be correlated with site intensity, but not any random measurement error in the regressor.

Table 1 reports the estimated uplift in land prices gradients for the different zones (i.e. $\left\{\hat{\beta}_{1,j}^{(1)}\right\}_{j=1}^m$) under various model specifications. These include different measures of distance (Haversine and Manhattan), specifications with and without controls, and OLS and IV estimation. In our analysis

Table 1: Relative Increase in Land Price Gradients

Zone	(A)	(B)	(C)	(D)	(E)	(F)	(G)	(H)
MHS	0.205	0.193	0.221	0.223	0.189	0.174	0.200	0.201
MHU	0.265	0.250	0.329	0.331	0.248	0.230	0.305	0.304
THA	0.369	0.355	0.454	0.452	0.349	0.332	0.420	0.416

Notes: Tabulated entries are log land price changes upzoned areas (MHS, MHU or THA) relative to non-upzoned areas. (A) denotes OLS estimates from (8) using Haversine distance to the CBD; (B) denotes Haversine distance with controls included in the regression; (C) denotes Haversine distance from IV estimation of (8); (D) denotes Haversine distance with controls included in the regression and IV estimation. (E)–(H) repeat (A)–(D) but use the Manhattan measure of distance.

to follow, we will present results for all eight estimates presented in Table 1.

Among these approaches, our preferred approach is IV estimation with controls, fixed effects, and Manhattan distance. IV accounts for measurement error in site intensity, which would generate attenuation bias, leading us to underestimate land price gradient differentials.Including controls accounts for endogenous upzoning selection through covariates. Manhattan distance accounts for travel via extant road networks, which, as discussed earlier, is important for a topography such as Auckland's, with its numerous water bodies. Including fixed effects tightens the standard errors.

Results based on our preferred specification appear in column (H). This shows that upzoning increased log land prices by 0.416 in upzoned to THA compared to non-upzoned, 0.304 in MHU, and 0.201 in MHS. These are somewhat higher than the estimates based on OLS (in column (F)), which is consistent with measurement error causing attenuation bias in the OLS estimator. We also note that the estimates based on Haversine distance are somewhat larger, which is consistent with the Manhattan distance exceeding Haversine distance for each property.[10]

We also depict the change in price gradient under our preferred empirical specification in Figure 6. Analogous changes in price gradients for other specifications can be found in the Appendix.

5 Assessing the Impacts of Upzoning in Auckland

In this section we interpret the relative changes in land prices from upzoning through the lens of theoretical model to assess the effects of the zoning reform. Following Bertaud and Brueckner (2005), the impacts of rezoning are assessed by comparing two static equilibria: one with the old LURs and one with the new LURs.Bertaud and Brueckner (2005) consider a hypothetical policy in which height restrictions in Bangalore are relaxed. We consider a real-world policy.

We model upzoning as a policy that shifts the spatial equilibrium attained under a set of restrictive LURs towards another spatial equilibrium attained under a less restrictive set of LURs. By comparing static equilibria under the two different sets of restrictions we can anticipate the effects

[10]Manhattan distance can be thought of as an approximate enlarging dilation of Haversine distance, with a scaling factor that varies by location.

Figure 6: Changes in Land Gradient Differentials from Upzoning, Instrumental Variables Estimation

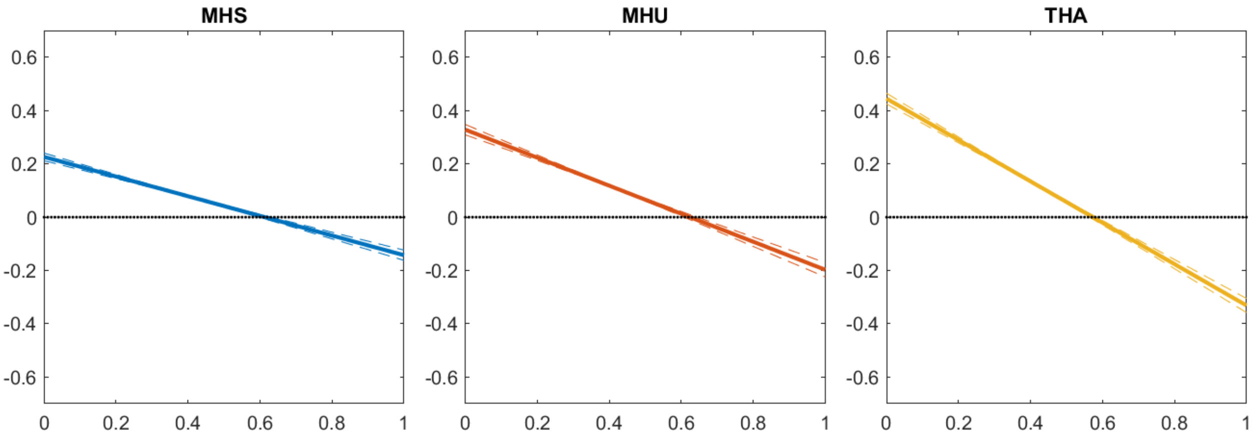

Notes: y-axis is the change in log dwelling price in upzoned areas (MHS, MHU or THA) relative to non-upzoned areas conditional on site intensity of the property (x-axis). Fitted values obtained from IV estimation of (8). White (1980) standard errors. Top: OLS estimation. Bottom: IV estimation with controls.

of upzoning on urban development patterns and the housing market. Although this comparison abstracts from the process of redevelopment, which can reasonably be expected to play out over several decades, it nonetheless provides a framework for understanding how upzoning affects outcomes such as land and house prices and urban development more broadly if it can be reasonably expected that the city transitions from one static equilibria towards the other after the policy.

We proceed in two steps. First, we calculate the increase in floorspace supply from the reform. This is achieved by combining estimates of the relative increase in upzoned land price gradients with geographical information on the amount of residential land assigned to each zone and the FARs of extant dwellings.Second, we back out the implied change in equilibrium prices in response to an increase in dwelling floorspace supply. To do so, we use a set of plausible estimates of housing price elasticities of demand, taken from the extant literature. We discuss each step in turn.

5.1 Changes in floorspace supply

Viewed through the lens of the theoretical model, the changes in the relative land price gradients estimated in the previous section tell us about anticipated increases in supplied floorspace.Under (3) we have

$$\frac{h_j(x)}{h_0(x)} = \frac{\frac{dR_j(x)}{dx}}{\frac{dR_0(x)}{dx}}$$

Introducing the time subscript, from (9) it follows that

$$\mathrm{E}\left[\ln\left(h_{t,j}(x)\right) - \ln\left(h_{0,j}(x)\right) - \left[\ln\left(h_{t,0}(x)\right) - \ln\left(h_{0,0}(x)\right)\right]\right] = \beta_{t,j}^{(1)} \tag{11}$$

The estimates $\left\{\hat{\beta}_{1,j}^{(1)}\right\}_{j=1}^{3}$ therefore provide us with estimates of the increase in FARs in the three upzoned areas (MHS, MHU and THA) relative to non-upzoned areas.

We assume that $\ln\left(h_{0,j}\left(x\right)\right)-\ln\left(h_{0,0}\left(x\right)\right)=0$, i.e. there is no change in the FAR of non-upzoned areas for the set of distances x considered in the analysis. Given an estimate of $h_{0,j}\left(x\right)$ we can obtain the absolute increase in FAR in zone j as $\left(e^{\hat{\beta}_{1,j}}-1\right)h_{t,0}\left(x\right)$.

We can combine the implied increase in floorspace with (i) the proportion of land assigned to each zone, and (ii) the existing floorspace in each zone, to calculate an overall estimate of floorspace increase. We calculate the log change in floorspace capacity as follows:

$$ln\left(\frac{\sum_{j=1}^{m_z}e^{\hat{\beta}_{j,1}}a_jh_j}{\sum_{j=1}^{m_z}a_jh_j}\right) \tag{12}$$

where a_j denotes the amount of land in zone j and h_j denotes the current (pre-upzoning) FAR that developers build to in zone j. $e^{\hat{\beta}_{j,1}}$ is the increase in FARs in zone j as estimated from our DID analysis. (Note that for non-upzoned areas, $\hat{\beta}_{j,1}=0$.)

For the purpose of computing the increase in floorspace supply, we expand the set of zones to include two additional groups: *Semi-Rural* and *Rural*. Semi-rural is composed of *Large Lot* and *Rural and Coastal Settlement*, which are peri-urban zones at suburban fringes of the city. Their LURs are also tabulated in Table 6 in the Appendix, and demonstrate that they are very low density zones. We also include rural zones because commuters can live in dwellings on rural land, albeit in dwellings with significantly lower capital intensities than those built in residential areas due to even more restrictive LURs. Allowing residents to commute in from rural zones is important because urban growth boundaries have been found to constrain urban growth in Auckland (Grimes and Liang, 2009). Excluding these zones from the analysis would significantly bias the increase in floorspace capacity upwards.[11]

The increase in floorspace supply depends on the geographic delineation of the city. In the AMM model, city size is an endogenous outcome, delimited by the marginal commuter that is indifferent between commuting towards the CBD and working in the rural economy. Commuting zones are therefore more analogous to cities than administrative boundaries.

The Auckland region does not comprise a single commuting zone according to Statistics New Zealand.[12] While there is a large commuting zone centered on the CBD and central isthmus, much of the rural areas are not part of any commuting zone, and there is a second, smaller commuting zone around the township of Warkworth, approximately 50 km North of the CBD. Thus, in calculating the increase in floorspace supply, we aggregate over all of the commuting zone of central Auckland.This excludes the Barrier and Gulf Islands from the analysis, as well as several rural SAs and the SAs associated with the Warkworth commuting zone. However, we will also present results for smaller city sizes based on both Haversine and Manhattan distances to the CBD. A city radius

[11] We include the Unitary Plan "Future Urban" zone in "Rural" and the "Hauraki Gulf Islands" zone in "Large Lot".

[12] Statistics New Zealand defines Functional Urban Areas (FUAs) on the basis of commuting patterns. We use the 2018 Auckland FUA.

of 50km in Haversine distance is sufficient to capture all of the Auckland commuting zone. A radius of 75km is required for Manhattan distance.

We straightforwardly calculate land areas $\{a_j\}_{j=1}^{m}$ using GIS information on zones. Table 2 exhibits the proportion of residential land assigned to each zone, as well as the total residential land area, for different city sizes, using both Haversine and Manhattan distances to the CBD. We use 2018 SAs as the geographic unit of analysis and the centroid of the SA is used to generate measures of distance to the CBD.

We use data on parcel-level FARs for recent builds to estimate $\{h_j\}_{j=1}^{m}$. We use recent builds because these are more likely to reflect the FARs that developers build to under the market and regulatory conditions prior to upzoning. The Appendix details the methodology adopted for assigning FARs to each dwelling in the RV dataset. We use the 2017 RVs to calculate the FARs. Although the valuation occurred in July 2017, which is nine months after the AUP became operational, we anticipate that only a handful of buildings would have been completed under the new LURs. We also constrain the calculations of FARs to recently built dwellings, which are more likely to reflect the LURs in place prior to the policy change. Because the RV dataset only contains the decade of build, not the year, we constrain our time frame to the 2000s and to the 2010s.

Table 3 reports the median FARs in each zone for different city radii. We use median FARs as arithmetic means are larger than medians, indicating that the distribution is positively skewed and possibly affected by large outliers.

Several patterns deserve comment. First, FARs are somewhat higher in zones targeted for upzoning (THA, MHU, MHS) compared to SH.[13] This observation is, however, inconsequential to our empirical strategy, which only requires that there is no change to FARs in the control group from the policy. Second, FARs tend to decrease between 10 and 20km, thereafter there is less variation. This is consistent with greater demand (and supply) for floorspace and dwellings close to the city center. In unreported results, using all buildings to compute the FAR resulted in similar median FARs across the THA, MHU, MHS and SH zones.

Table 4 exhibits the increase in floorspace capacity using eq. (12) for various city sizes.

Columns (A) through (D) correspond to capacity increases under the Haversine measure of distance. Column (A) uses the changes in land price gradients obtained from OLS estimation of the DID model (given in column (A) of Table 1). (B) denotes OLS estimation of the DID model augmented with controls. Columns (C) and (D) denote IV estimation of the DID model and the DID model with controls. Columns (E)–(H) repeat the specifications for (A)–(D) but use Manhattan distance instead of Haversine distance.

Column (H) is our preferred estimate based on Haversine Distance and IV estimation of the DID model augmented with controls. It shows that the projected increases in floorspace capacity decreases as the radius of the city expands. For example, for a city radius of 30km, the increase in floorspace is 22.9%. By 75km, which includes all of the Auckland commuting zone, the increase is 18.0%. This pattern is due to the decrease in the proportion of upzoned residential land as the city radius expands (see Table 2 above). This is our preferred estimate of the aggregate increase in

[13]Interestingly this pattern is not evident when we do not limit the sample to new dwellings.

Table 2: Residential Land Area and Proportion of Zones for Different City Sizes

Haversine Distance

City Size (km radius)	Area	Population	Dwellings	THA	MHU	MHS	SH	Semi-Rural	Rural
				Proportion of residential land in each zone					
10	101.0	562.4	193.2	0.108	0.231	0.471	0.189	0.001	-
20	422.4	1212.2	383.0	0.054	0.147	0.262	0.101	0.061	0.375
30	986.8	1438.3	450.2	0.025	0.072	0.138	0.065	0.047	0.652
40	1731.6	1498.7	470.0	0.014	0.043	0.083	0.044	0.036	0.780
50	2366.9	1534.5	482.1	0.011	0.032	0.063	0.033	0.028	0.833
60	2366.9	1534.5	482.1	0.011	0.032	0.063	0.033	0.028	0.833

Manhattan Distance

City Size (km radius)	Area	Population	Dwellings	THA	MHU	MHS	SH	Semi-Rural	Rural
				Proportion of residential land in each zone					
15	117.7	644.7	218.9	0.123	0.263	0.446	0.165	0.003	-
30	536.3	1316.9	409.7	0.044	0.124	0.234	0.084	0.053	0.461
45	1326.4	1468.4	459.6	0.018	0.056	0.107	0.054	0.042	0.723
60	2110.8	1523.3	478.2	0.012	0.036	0.070	0.037	0.031	0.815
75	2366.9	1534.5	482.1	0.011	0.032	0.063	0.033	0.028	0.833

Notes: Radius is Haversine or Manhattan distance from the CBD to the city edge. For a given radius, the tabulated data are calculated based on the set of statistical areas (SAs) that have a centroid within the specified radius of the CBD. For example, for the 10km Haversine city radius, only SAs with centroids within 10km of the CBD are included in the calculations. Only residential and rural land areas are tabulated. Population and Dwellings are based on March 2018 census data.

Table 3: Median FARs in different Residential Zones prior to Upzoning

Haversine Distance						
City Size (km radius)	THA	MHU	MHS	SH	Semi-Rural	Rural
10	0.607	0.499	0.478	0.459	0.238	
20	0.583	0.402	0.420	0.426	0.122	0.020
30	0.583	0.397	0.400	0.382	0.109	0.022
40	0.583	0.415	0.393	0.374	0.111	0.019
50	0.583	0.415	0.385	0.370	0.101	0.018

Manhattan Distance						
City Size (km radius)	THA	MHU	MHS	SH	Semi-Rural	Rural
15	0.566	0.480	0.477	0.446	0.298	
30	0.582	0.400	0.400	0.417	0.120	0.022
45	0.581	0.415	0.396	0.382	0.109	0.019
60	0.583	0.415	0.389	0.371	0.108	0.016
75	0.583	0.415	0.385	0.370	0.101	0.016

Notes: FARs are floor to land area ratios and are calculated for dwellings completed between 2000 and July 2017. Radius is distance from the CBD to the city edge. For a given radius, the tabulated data are the median FARs of dwellings that have a geocoordinate within the specified radius. For example, the median FAR among dwellings in the THA zone and within 40km of the CBD is 0.590 when the Haversine measure of distance is used.

Table 4: Percent Increase in Floorspace for Different City Sizes

Haversine Distance					Manhattan Distance				
City Size (km radius)	(A)	(B)	(C)	(D)	City Size (km radius)	(E)	(F)	(G)	(H)
10	23.63	22.14	28.00	28.17	15	22.62	20.82	25.94	25.95
20	22.27	20.86	26.40	26.57	30	20.09	18.49	22.92	22.94
30	19.84	18.59	23.47	23.63	45	17.31	15.92	19.74	19.77
40	18.06	16.92	21.38	21.53	60	16.06	14.78	18.31	18.34
50	17.21	16.12	20.36	20.50	75	15.81	14.55	18.02	18.05

Notes: Tabulated figures are percent increases floorspace. Radius is distance from the CBD to the city edge. (A) denotes the capacity estimate using the DID model without controls; (B) denotes the capacity estimate using the DID model with controls; (C) denotes the capacity estimate using IV estimation of the DID model; (D) denotes the capacity estimate using IV estimation the DID model with controls. Columns (A)–(D) use Haversine distance. Columns (E)–(H) repeat the specifications of (A)–(D) but use Manhattan distance instead of Haversine distance. Median FARs are used to calculate the increase in floorspace.

capacity from upzoning.

5.2 Changes in housing costs

We translate the implied increase in housing supply into a reduction in dwelling prices using plausible estimates of the own-price elasticity of housing demand, which we source from the existing literature. This approach assumes that future housing demand only depends on price changes, which is concordant with the closed city assumption of fixed population.

Most estimates lie between a range of -0.5 to -0.8, indicating housing demand is price-inelastic. Polinsky and Ellwood (1979) produce an estimate of -0.7 and Hanushek and Quigley (1979) produce experimental estimates of -0.64 for Pittsburgh and -0.45 for Phoenix. Albouy et al. (2018) conclude that the price elasticity is close to two-thirds (in magnitude), though they find estimates of anywhere between -0.83 and -0.49 depending on the model specification. Unfortunately we were unable to source peer-reviewed price elasticity estimates for New Zealand from the extant literature.

We consider price elasticities of -0.5, -0.65 and -0.8.[14] The upper and lower bounds reflect the range of estimates obtained by Albouy et al. (2018), while -0.65 lies in the middle of the upper and lower bound and also approximates their preferred estimate of two-thirds. Combining the price elasticities of demand with the estimated percent increases in floorspace capacity yield an estimate of the implied percent changes in dwelling prices.[15] The larger the elasticity (in magnitude), the smaller the implied change in house prices. We use the implied increases in floorspace capacity given in columns (C), (D) (G) or (H) from Table 4, which correspond to IV estimates.

Table 5 exhibits the results. Panel (A) presents results for OLS estimation of the DID model without controls and uses Haversine distance. Panel (B) is based on OLS estimation of the DID model with controls and uses Haversine distance. Panels (C) and (D) replicate panels (A) and (B) but are based on IV estimation. (E) through (H) replicate (A) through (D) but with Manhattan measure of distance.

Our preferred estimate is given in panel (H) as it corresponds to IV estimation of the DID model with controls and Manhattan distance. At the city radius of 75km (Manhattan distance), which encapsulates the commuting regions of Auckland, this model implies a 29.1% reduction in house prices using the middle estimate of the demand elasticity parameter (i.e.,-0.65). The upper bound estimate corresponds to a 39.3% reduction in house prices, while the lower bound corresponds to a 23.0% reduction in house prices.

Price decreases generally get larger as the radius of the city gets smaller. This reflect the fact that proportionately more land area is upzoned as the radius contracts towards the CBD. For example, if the radius of Auckland is thought to have a radius of 30km (Manhattan distance), our preferred model specification implies a decrease in house prices of 37.4%. However, it is important to note that this conclusion supposes that the size of Auckland is much smaller,and that housing

[14] Of course, the Cobb-Douglas utility in the AMM model implies an elasticity of -1.

[15] Letting g denote the percent increase in housing stock, the implied percent decrease in dwelling prices is $\left(e^{ln(1+g/100)/\hat{\beta}} - 1\right) \times 100$

located further from 30km is not considered to be part of Auckland city.[16] We focus the impact of the policy on house prices when the entirety of the Auckland commuting zone is included in the analysis, and therefore use 75km as our cutoff (or 50km Haversine distance). As the table shows, this results in a more conservative estimate of the price impacts of the policy.

6 Conclusion

This paper sets out a framework for examining the long run impacts of upzoning on urban development. Our approach develops the AMM model to allow for variation in LURs between residential zones, mimicking urban planning in practice. We show how relative changes in land price gradients between different zones after zoning reforms can be used to infer changes in equilibrium development patterns from the policy change, including the increase in optimal floorspace from redevelopment.

We use the model to evaluate the impact of a recent upzoning reform implemented in Auckland, New Zealand. The application makes use of detailed information on the geographic location of different residential zones, the land use regulations that apply with each zone, and information on the individual dwellings comprising the housing stock prior to upzoning. We use this information to estimate relative changes in property values, including land prices, in upzoned and non-upzoned areas using a conventional difference-in-differences framework. Interpreting these results through the lens of our AMM model, the changes in land price gradients are consistent with a 33.2% increase in floorspace capacity. Using plausible estimates of the house price elasticity of demand from the extant literature, this translates into a reduction in dwelling prices of between 26.2% and 45.2%, relative to the counterfactual of no upzoning.

[16]Note that the Table should not be understood as implying that dwelling price reductions are generally larger towards the center of the city. In fact, under the AMM model with Cobb Douglas utility, the reduction in house prices is proportionately that same for all $x \leq \bar{x}$. The Table shows price decreases for different values of $\bar{x}$.

Table 5: Decrease in Dwelling Prices for Different City Sizes

Haversine Distance

City Size (km radius)	(A) price elasticity =			(B) price elasticity =			(C) price elasticity =			(D) price elasticity =		
	-0.8	-0.65	-0.5	-0.8	-0.65	-0.5	-0.8	-0.65	-0.5	-0.8	-0.65	-0.5
20	28.6	36.2	49.5	26.7	33.8	46.1	34.0	43.4	59.8	34.2	43.7	60.2
30	25.4	32.1	43.6	23.8	30.0	40.6	30.2	38.3	52.4	30.4	38.6	52.8
40	23.1	29.1	39.4	21.6	27.2	36.7	27.4	34.7	47.3	27.6	35.0	47.7
50	22.0	27.7	37.4	20.5	25.9	34.8	26.1	33.0	44.9	26.2	33.2	45.2

Manhattan Distance

City Size (km radius)	(E) price elasticity =			(F) price elasticity =			(G) price elasticity =			(H) price elasticity =		
	-0.8	-0.65	-0.5	-0.8	-0.65	-0.5	-0.8	-0.65	-0.5	-0.8	-0.65	-0.5
30	25.7	32.5	44.2	23.6	29.8	40.4	29.4	37.4	51.1	29.5	37.4	51.1
45	22.1	27.8	37.6	20.3	25.5	34.4	25.3	31.9	43.4	25.3	32.0	43.4
60	20.5	25.7	34.7	18.8	23.6	31.7	23.4	29.5	40.0	23.4	29.6	40.0
75	20.1	25.3	34.1	18.5	23.2	31.2	23.0	29.0	39.3	23.0	29.1	39.3

Notes: Table entries are percent decreases in dwelling prices on a per unit of flooring basis. Radius is distance from the CBD to the city edge. Panels (A) and (B) are OLS estimation of increased floorspace with and without controls. Panels (C) and (D) are IV estimation of increased floorspace with and without controls. Panels (E) through (H) are defined as (A) through (D) but are based on Manhattan distance. Price decreases are generally larger for smaller city sizes because the proportion of upzoned land decreases as the radius of the city increases.

7 Appendix

7.1 Additional Tables and Figures

Table 6: Summary of Land Use Regulation by Residential Zone under the Unitary Plan

Regulation	Terraced Housing Apartments	Mixed Housing Urban	Mixed Housing Suburban	Single House	Large Lot	Rural & Coastal Settlement
Max. height	16m (five storeys)	11 to 12m (three storeys)	8 to 9m (two storeys)	8 to 9m (two storeys)	8 to 9m (two storeys)	8 to 9m (two storeys)
Height in relation to boundary	3m up + 45° recession plane	3m up + 45° recession plane	2.5m up + 45° recession plane	2.5m up + 45° recession plane	does not apply*	2.5m up + 45° recession plane
Setback (side and rear)	0m	1m	1m	1m	6m	1m
Setback (front)	1.5m	2.5m	3m	3m	10m	5m
Max. site coverage (%)	50%	45%	40%	35%	lesser of 20% or $400m^2$	lesser of 20% or $400m^2$
Max. impervious area (%)	70%	60%	60%	60%	lesser of 35% or $1400m^2$	lesser of 35% or $1400m^2$
Min. dwelling size (1 bedroom)	$45m^2$	$45m^2$	$45m^2$	n/a	n/a	n/a
Max. dwellings per site	does not apply	3	3	1	1	1
Min. Lot Size (subdivision)	$1200m^2$	$300m^2$	$400m^2$	$600m^2$	$2500m^2$	$4000m^2$

Notes: Tabulated restrictions are 'as of right' and can be exceeded through resource consent notification. Height in relation to boundary restrictions 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. Planners have discretion in setting height in relation to boundary and setbacks in the large lot zone. The regulations "[r]require development to be of a height and bulk and have sufficient setbacks and open space to maintain and be in keeping with the spacious landscape character of the area". Impervious area is the area under the dwelling and structures such as concrete driveways that prevent rainwater absorption into the soil.

Figure 7: Pre-trend effects of upzoning on land price gradients

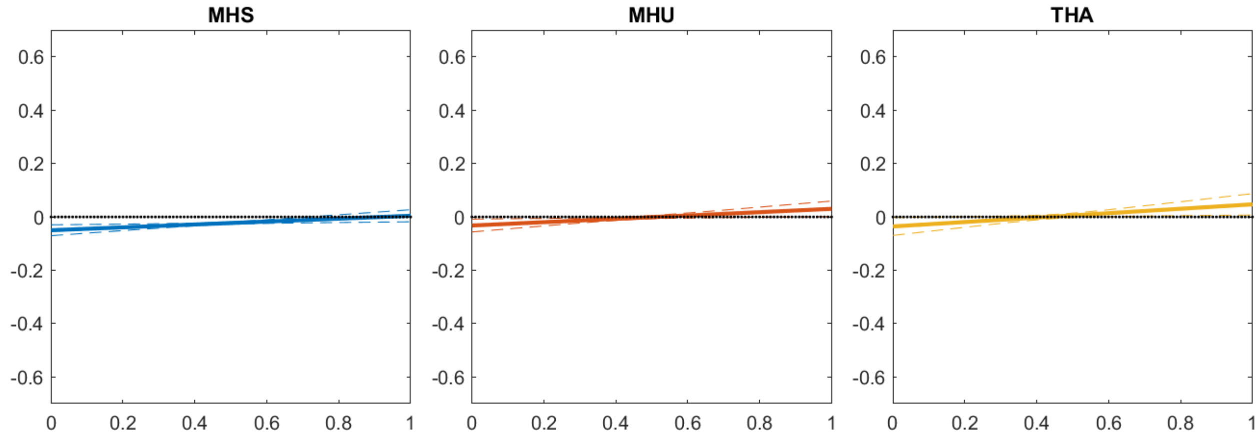

Notes: Pre-trends conditional on site intensity in period prior to treatment ($t = -1$). Outcome is change in log dwelling price between 2014 and 2017. Based on fitted model (8). IV estimation of (8) including control variables. White (1980) standard errors.

7.2 Alonso-Muth-Mills Model

Our specification follows Bertaud and Brueckner (2005). Households have preferences described by

$$U(C, H) = C^{1-\alpha} H^{\alpha},\ \alpha \in (0, 1)$$

where C is the consumption numeraire (price set to unity) and H is housing floorspace. Their budget constraint is

$$C = W - tx - P(x)\, H$$

where x denotes the distance of their house to the CBD, t is the per distance cost of commuting, and $P(x)$ is dwelling rent at distance x. Travel costs t are comprised of a pecuniary cost g and an opportunity cost of time $\upsilon \times W$ for some $\upsilon \in (0, 1]$, such that $t = g + \upsilon W s^{-1}$, where s is the speed of commute.

Housing floorspace H is produced using capital K and land L as

$$H(K, L) = AK^{\gamma} L^{1-\gamma},\ \gamma \in (0, 1)$$

where L is land and K is capital K. Because production is constant returns to scale, we can define housing per unit of land as

$$h = \frac{H}{L} = A \left(\frac{K}{L} \right)^{\gamma} = Ak^{\gamma}$$

where $k = K/L$ is the capital to land ratio. Developer profits per unit of land are then

$$P(x)\, h(k) - p_K k - R(x)$$

where $R(x)$ is per unit land rent at x. Developers profit-maximise and earn zero profits. Under

Figure 8: Changes in Land Gradient Differentials from Upzoning

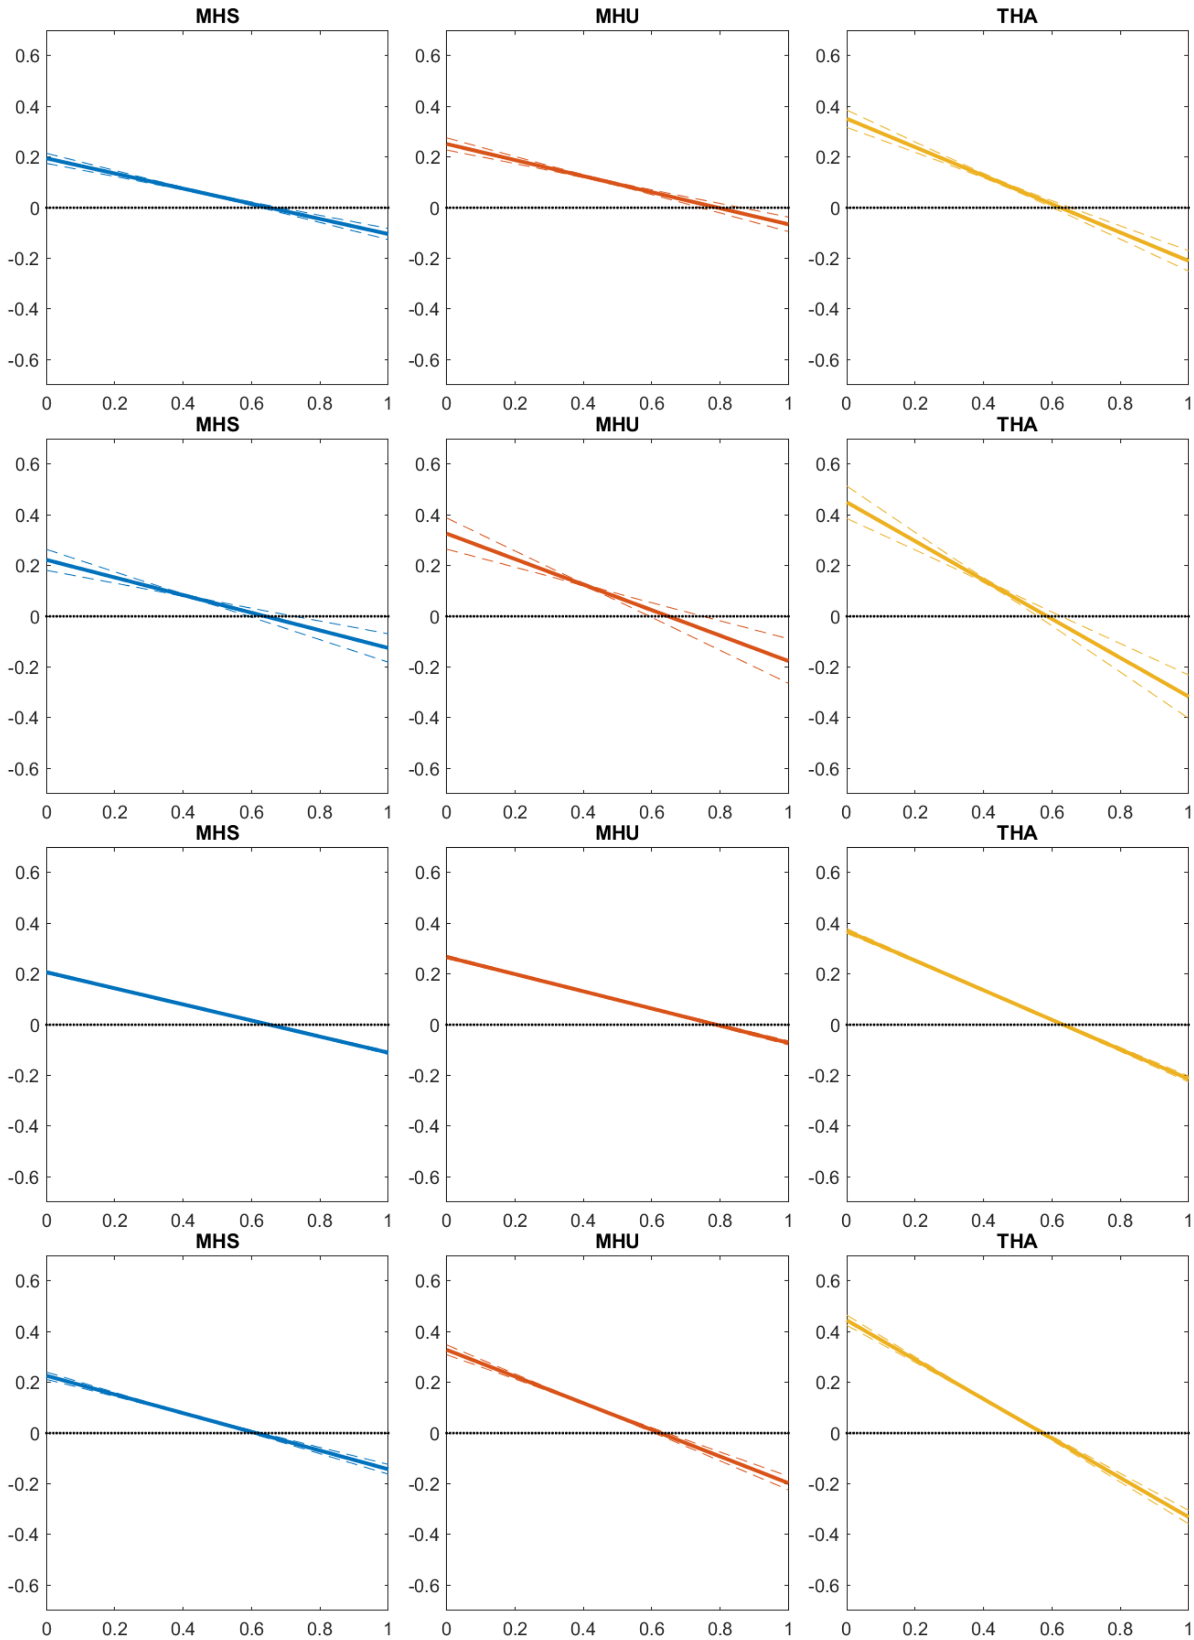

Notes: Robustness Checks. From top panel to bottom panel: OLS with controls; IV estimation; OLS with fixed effects; IV with controls and fixed effects.

Figure 9: Changes in Land Gradient Differentials from Upzoning using Manhattan Distance

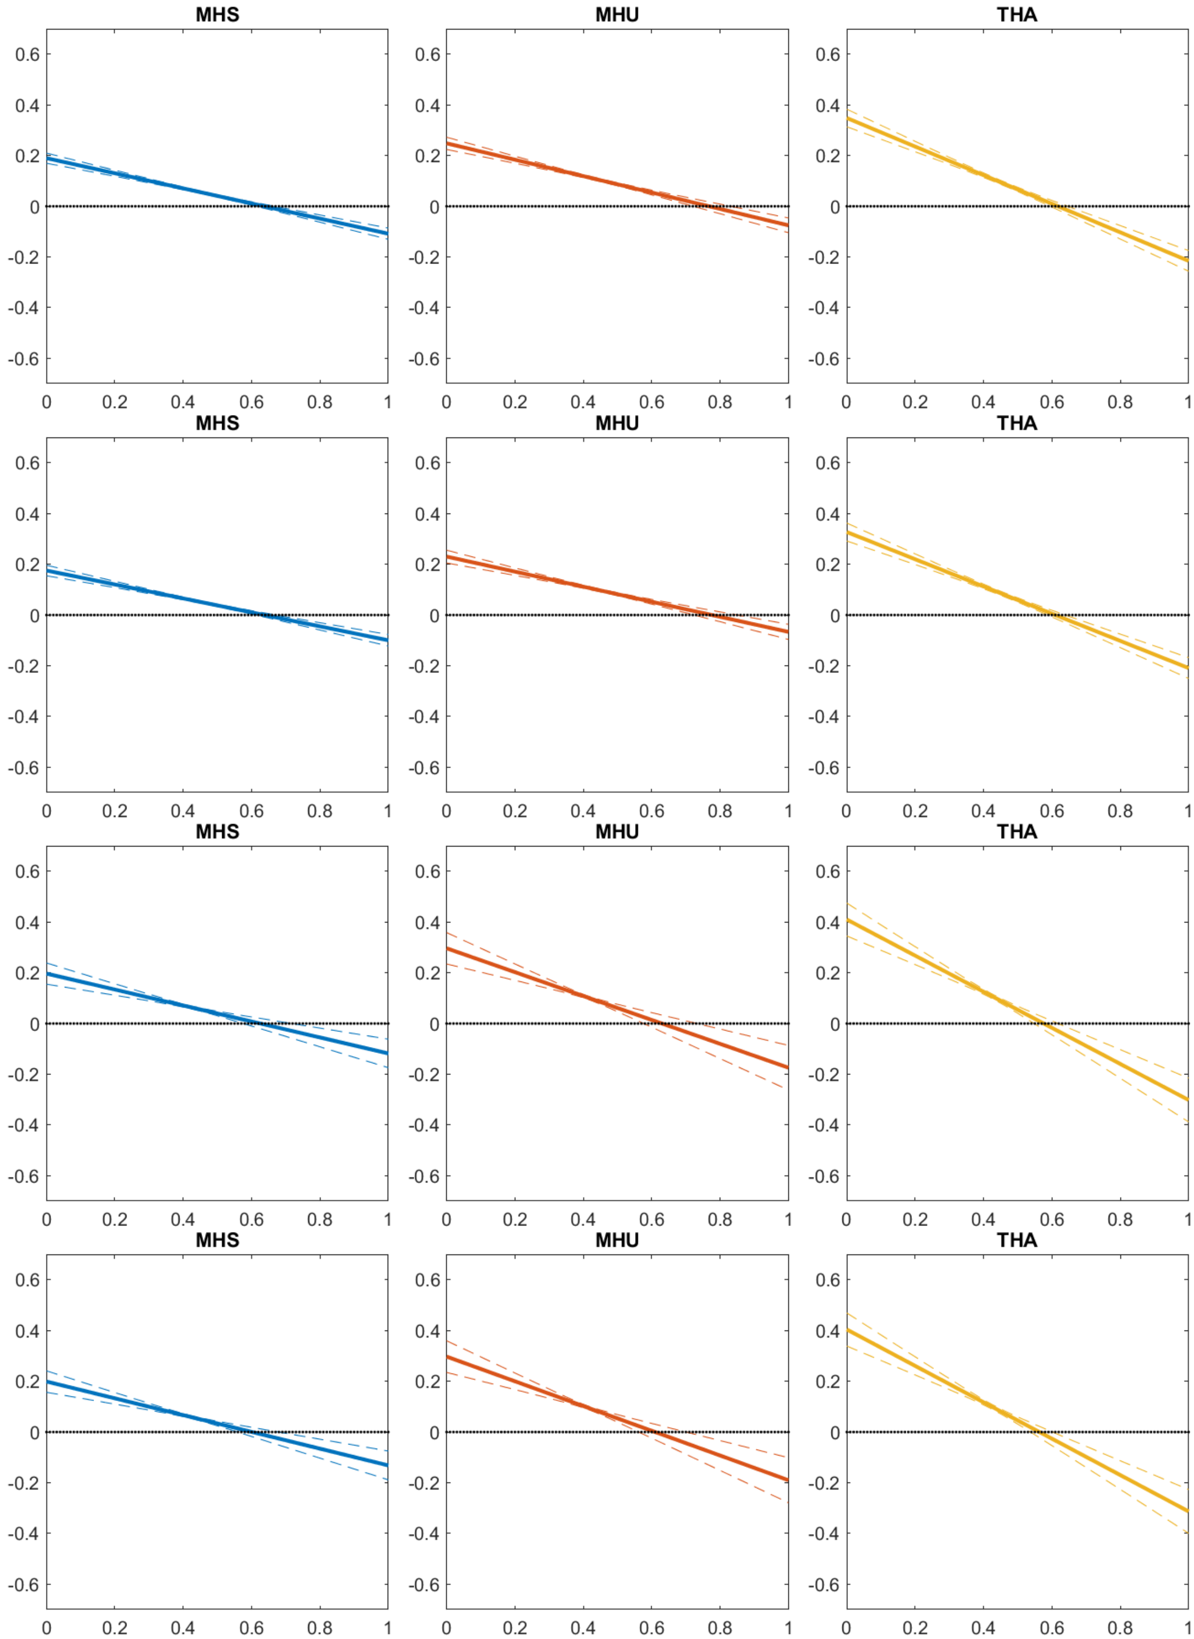

Notes: Manhattan Distance. From top panel to bottom panel: OLS; OLS with controls; IV estimation; IV with controls.

these assumptions, h, H and k depend on distance to the city center x, and are henceforth denoted as $h(x)$, $H(x)$ and $k(x)$.

Three equilibrium conditions solve the standard model without restrictions on $h(x)$. First, utility is equalized at all locations x. Second, land rent at the edge of the city $\bar{x}$ is equal to exogenous agricultural rents $\bar{R}$, $R(\bar{x}) = \bar{R}$. Finally, the city population N fits into the area of the city, $\int_0^{\bar{x}} \theta h(x)/H(x)\, x dx = N$, where we normalize one worker per dwelling.

In equilibrium, utility, dwelling prices and floorspace depend only on $\bar{x}$. Equilibrium utility $\bar{U}$ is given by

$$\bar{U} = \frac{\alpha^{\alpha}\gamma^{\alpha\gamma}A^{\alpha}(1-\gamma)^{\alpha(1-\gamma)}(W-t\bar{x})(1-\alpha)^{(1-\alpha)}}{\bar{R}^{\alpha(1-\gamma)}p_K^{\alpha\gamma}}$$

Then by solving for the Hicksian demand functions we have

$$P(x) = \frac{(W-tx)^{\frac{1}{\alpha}}(1-\alpha)^{\frac{1-\alpha}{\alpha}}\alpha}{\bar{U}^{\frac{1}{\alpha}}} \tag{13}$$

and

$$H(x) = \frac{\bar{U}^{\frac{1}{\alpha}}}{(W-tx)^{\frac{1-\alpha}{\alpha}}(1-\alpha)^{\frac{1-\alpha}{\alpha}}} \tag{14}$$

$h(x)$ and $R(x)$ can then be solved for in each zone $k = 1, \ldots, m_z$ using the zero profit condition and the restriction on $h(x)$ that applies in each zone.

7.3 Estimating FARs of extant dwellings prior to upzoning

The ratings valuation (RV) dataset contains information on floorspace for each dwelling. To obtain a land area from each of these dwellings, we first match the RV to land parcels using the process described below.

Procedure for matching RVs to LINZ parcels. We use the following steps for each RV.

(i) Check whether the title of the RV matches to a LINZ parcel. Stop if there is a match. If not, proceed to the next step.

(ii) Assign the RV a geocoordinate using its address. Identify the the LINZ parcel of the geocoordinate. If the number and first word of the LINZ parcel address matches the address of the RV, we have identified the LINZ parcel and stop. If the is no match, proceed to the next step.

(iii) Find all the LINZ parcels within a 2km radius of the geocoordinate of the address. Within this set, we check whether the number and first word of the LINZ parcel address matches the address of the RV. If there is match, we stop. If not, we proceed to the next step.

(iv) Do a text field search among all addresses in the LINZ dataset. Within this set, we check whether the number and first word of the LINZ parcel address matches the address of the RV, and whether the name of the suburb of the RV address matches that of the LINZ parcel.

If there is match, we stop. If not, there is no match, and the RV does not have a parcel assigned to it.

97.3% of all RVs are matched under step 1. A further 2.3% are matched at step 2; 0.2% at step 3; and 0.09% at step 4

Procedure for apportioning floorspace to parcels. The ratings valuation (RV) dataset contains information on floorspace for each dwelling. Land area is also included for dwellings that have exclusive land ownership on the title. For these properties it is straightforward to calculate the FAR of the property. High density dwelling forms and single family homes on cross leases do not have a reported land area. To obtain a land are from each of these dwellings, we match the dwelling to land parcels using the following process. For many-to-one matches (i.e., several dwellings assigned to the same land parcel, as would be expected for apartment buildings), we assign all dwellings the same FAR, which is calculated as the total floor area of all dwellings assigned to the parcel divide by the land are of the parcel.

References

ALBOUY, D., G. EHRLICH, AND M. SHIN (2018): "Metropolitan land values," *Review of Economics and Statistics*, 100, 454–466. 4, 29

ALONSO, W. (1964): *Location and Land Use*, Cambridge: Harvard University Press. 3, 5

ARNOTT, R. J. AND J. G. MACKINNON (1977): "Measuring the costs of height restrictions with a general equilibrium model," *Regional Science and Urban Economics*, 7, 359–375. 3, 4, 5, 6, 9, 11

BAUM-SNOW, N. (2007): "Suburbanization and transportation in the monocentric model," *Journal of Urban Economics*, 62, 405–423. 6

BERTAUD, A. AND J. K. BRUECKNER (2005): "Analyzing building-height restrictions: Predicted impacts and welfare costs," *Regional Science and Urban Economics*, 35, 109–125. 3, 4, 5, 6, 7, 9, 11, 23, 34

CLAPP, J. M., K. S. BARDOS, AND S. K. WONG (2012): "Empirical estimation of the option premium for residential redevelopment," *Regional Science and Urban Economics*, 42, 240–256. 19

CLAPP, J. M. AND K. SALAVEI (2010): "Hedonic pricing with redevelopment options: A new approach to estimating depreciation effects," *Journal of Urban Economics*, 67, 362–377. 19

COOPER, G. C. AND K. NAMIT (2021): "City with a billion dollar view," *New Zealand Economic Papers*, 55, 19–37. 17

DURANTON, G. AND D. PUGA (2015): "Urban Land Use," *Handbook of Regional and Urban Economics*, 5, 467–560. 3, 5, 12

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, 4

FREEMARK, Y. (2019): "Upzoning Chicago: Impacts of a Zoning Reform on Property Values and Housing Construction:," *Urban Affairs Review*, 56, 758–789. 2

GLAESER, E. (2008): *Cities, Agglomeration, and Spatial Equilibrium*. 5

GLAESER, E. L. AND J. GYOURKO (2003): "Building restrictions and housing availability," *Economic Policy Review*, 21–39. 2

GRAY, M. N. (2022): *Arbitrary Lines: How Zoning Broke the American City and How to Fix It*, Island Press. 2

GREENAWAY-MCGREVY, R. AND J. A. JONES (2023): "Can zoning reform change urban development patterns? Evidence from Auckland," . 2, 3, 14, 17

——— (2024): "Mode Choice and the Effects of Rapid Transit Improvements on Private Vehicle Use and Urban Development," *Journal of Transportation Economics and Policy*. 6

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, 959–976. 2, 4

Greenaway-McGrevy, R. and P. Phillips (2023): "The Impact of Upzoning on House Prices and Urban Development in Auckland," *Journal of Urban Economics*. 2, 3, 4, 15, 16

Grimes, A. and Y. Liang (2009): "Spatial determinants of land prices: Does Auckland's Metropolitan urban limit have an effect?" *Applied Spatial Analysis and Policy*, 2, 23–45. 25

Hanushek, E. A. and J. M. Quigley (1979): "The dynamics of the housing market: A stock adjustment model of housing consumption," *Journal of Urban Economics*, 6, 90–111. 29

Imbens, G. W. and J. M. Wooldridge (2009): "Recent developments in the econometrics of program evaluation," *Journal of Economic Literature*, 47, 5–86. 4

Kulish, M., A. Richards, and C. Gillitzer (2012): "Urban Structure and Housing Prices: Some Evidence from Australian Cities*," *Economic Record*, 88, 303–322. 4, 5, 6, 9, 10, 11

Limb, M. and C. K. Murray (2022): "We zoned for density and got higher house prices: Supply and price effects of upzoning over 20 years," *Urban Policy and Research*, DOI: 10.1080/08111146.2022.2124966. 2

Manville, M., P. Monkkonen, and M. Lens (2020): "It's Time to End Single-Family Zoning," *Journal of the American Planning Association*, 86, 106–112. 2

Mills, E. S. . (1967): "An Aggregative Model of Resource Allocation in a Metropolitan Area," *The American Economic Review*, 57, 197–210. 3, 5

Muth, R. F. (1969): *Cities and Housing.*, Chicago: University of Chicago Press. 3, 5

Peng, X. E. (2023): "The Dynamics of Urban Development : Evidence From Zoning Reform in New York," *Working Paper*. 2

Polinsky, A. M. and D. T. Ellwood (1979): "An Empirical Reconciliation of Micro and Grouped Estimates of the Demand for Housing," *The Review of Economics and Statistics*, 61, 199–205. 29

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, 223–248. 2

Schill, M. H. (2005): "Regulations and Housing Development: What We Know," *Cityscape*, 8, 5–19. 4

Wegmann, J. (2020): "Death to Single-Family Zoning. . . and New Life to the Missing Middle," *Journal of the American Planning Association*, 86, 113–119. 2