# Submission to the Productivity Commission Inquiry into Housing Supply Regulations

Matthew Maltman and Theo Gibbons[1]

June 2026

[1]This submission is made in a personal capacity and should not necessarily be attributed to the e61 Institute.

# Contents

## Summary

Thank you for the opportunity to make a submission to this inquiry. Our key points are:

- **Target bans, not burdens.** The evidence that regulatory *bans* (chiefly zoning and other land-use controls that prohibit denser forms of housing) raise prices and reduce the quantity of housing built is extensive and replicated across multiple settings. The evidence that regulatory *burdens* (approval timelines, document length, reporting, the contributions process) materially restrict supply is, by contrast, thin. Reform effort should be allocated in proportion to evidence.
- **Approval times and approval rates contain almost no information about constraints on construction.** Both are shaped by what is allowed to be built in the first place. It does not matter how fast a council gets back to applicants if the answer is "no". High approval rates are not evidence that construction is unconstrained, because developers rarely lodge applications they know will be refused. The Commission should resist building its diagnosis or its evaluation framework around these metrics.
- **The reform with the strongest evidence base is legalising medium-density housing in established suburbs.** This works through two related levers: allowing more dwellings on a given parcel (minimum-lot-size and subdivision reform) and allowing more floorspace on a given parcel (height and floor-area-ratio reform). There are now at least half a dozen credible empirical studies showing that removing these constraints produces large increases in the quantity of housing built.
- **Supply shocks have large effects on prices, when measured correctly.** Claims that supply has "no effect on prices" rest on flawed models or loose heuristics, or look for price effects at a spatial scale too small for them to surface.
- **The remaining questions are partly normative.** Economics can establish, with confidence, how much supply restrictions raise prices and reduce the quantity of housing built. Whether the offsetting amenity or externality benefits justify that cost is partly a question of community preferences. But because the costs are now well documented and large, the burden of proof should sit with those defending the restrictions.

## Recommendations

In line with the request for a prioritised list of reforms in the interim report:

1. **Prioritise removing bans over reducing burdens.** The highest-value reforms legalise medium-density housing in established, well-located areas: broad upzoning, reform of minimum lot sizes and subdivision rules, and relaxation of height and floor-area limits. Process and burden reforms (faster approvals, lower fees, streamlined referrals) are worth pursuing, but they are second-order.
2. **Do not use approval rates or timelines as indicators of supply restrictiveness**. A more informative metric is the share of residential land on which medium-density housing is legal as-of-right. The Commission should assess states and territories against how far they have gone in *legalising* density, not how fast they process applications.
3. **Shift the burden of proof onto retained restrictions.** The costs of restrictive zoning are now among the most robust findings in applied microeconomics. Significant land-use restrictions should be justified by a demonstrated, quantified externality.

## About the authors

We are research economists at the e61 Institute. This submission is made in our personal capacity: the views expressed are our own and should not be attributed to e61.

Our research focuses on applying quasi-experimental methods to housing supply, zoning, and planning policy, and several of the studies discussed in this submission are our own work:

- Maltman and Greenaway-McGrevy (2025), "Going it alone: The impact of upzoning on housing construction in Lower Hutt," published in the *Journal of Housing Economics*, provides causal estimates of the effect of Lower Hutt's 2017 upzoning on construction and rents.
- Donovan and Maltman (2025), "Dispelling myths: Reviewing the evidence on zoning reforms in Auckland," published in *Land Use Policy*, reviews the evidence on the Auckland Unitary Plan and assesses the critiques made of it.
- Maltman (2025), "Zoning Reform and Construction Sector Productivity," a working paper documenting the link between New Zealand's zoning reforms and its divergent construction productivity performance.
- Behrens et al. (2026), a forthcoming paper providing the first causal evidence on zoning reform in an Australian city, studies a 2014 minimum-lot-size reform in Adelaide.

We would be happy to discuss any part of this submission, or the underlying research, with the Commission.

# 1. A framework for the inquiry

A clear framework matters, because it determines which policies look sensible and which evidence is credible.

Our framework lays out three key points. First, land-use restrictions raise prices and reduce the amount of housing built, relative to what would otherwise prevail. The question for policy is then whether, given this, each restriction has benefits large enough to justify keeping it.
Second, supply and quantity are distinct objects. The distinction matters because the policies worth pursuing are those that shift the supply curve outward.
Third, prices and quantities are general-equilibrium outcomes. The only way to truly know what moves them is with an experiment or a model.

## 1.1 Land-use restrictions raise prices and lower quantities

Australia has accumulated a set of land-use and related restrictions that mean less housing is supplied at every price level, and that supply responds less when demand changes. The general-equilibrium consequence is straightforward: prices are higher, and fewer homes are built, than would otherwise prevail.

This is usually described as a housing "shortage". The recognition that Australia needs far more housing is welcome, but the label is not economically accurate, and the imprecision matters.

A shortage is excess demand at the current price. It arises in two circumstances. The first is after a shock to supply or demand, where prices have not yet adjusted: a disruption to oil supply, for instance, leaves petrol "short", meaning prices climb toward a new market-clearing level. The second is long-run, and requires something that prevents prices from adjusting. Social housing is in permanent shortage because it is, by definition, made available below market price. Rapid-antigen tests were in shortage during the pandemic because social and political pressure stopped pharmacies from raising prices until the market cleared.[2]

Neither circumstance describes the Australian housing market, where prices are free to adjust. The issue is precisely that prices have risen faster than wages for two decades and are above marginal cost.

The Australian literature has repeatedly tried to put a number on the "shortage" and

[2] In both of these circumstances there are good reasons why we may not want prices to rise and the market to "clear".

reached widely divergent conclusions, from large undersupply to outright surplus, despite using essentially the same data (Productivity Commission 2022). The divergence between these studies is due to assumptions, not data. These studies generate a gap between supply and demand only by fixing household formation at some assumed level. That is not how the market behaves: people form smaller households when housing is cheap, and larger ones when it is more expensive. Once household formation is allowed to respond to price, the apparent shortage dissolves. The market is clearing, but at a worse standard of living for Australian households. The Commission itself has previously rejected the shortage label as neither useful nor economically accurate for this reason (Productivity Commission 2022).

This is not semantics: it changes the central question we ought ask. That question is not "how many homes are we missing?", and therefore not "is reform X sufficient to close that gap?". That has no well-defined answer, just as there is no well-defined answer to how many cars or cafés Australia needs. The objective is not to reach a particular level of construction and then stop. It also does not matter, as some argue, if we already build "a lot". The right questions are: **how much does a policy raise the price, and reduce the quantity, of housing? And, is the policy worth it?** For zoning and planning regulations, the empirical answer to the first question, set out below, is "a great deal." The answer to the second is partly normative, and Section 8 returns to it.

Figure 1: The effect of land-use restrictions on housing supply.

This is not an argument for removing all restrictions. Many serve real purposes, and removing one can impose genuine costs through lost amenity, congestion, or other local externalities. The relevant test for each regulation is whether those costs exceed the benefit of the housing it prevents. Section 8 returns to how we might think about this trade-off.

## 1.2 Supply is not quantity

The single most important distinction in this submission is between supply and quantity, because conflating the two is the source of many of the mistaken arguments the Commission will hear.

Supply is a curve. Quantity is an outcome. Supply is how much housing gets produced at *each price*. Quantity is the number of dwellings actually built: a general-equilibrium outcome, the intersection of supply and demand, which moves when supply moves, when demand moves, or when both do. Supply-side reform shifts the supply curve outward, so that more housing is built at any given price.

It follows that the quantity of housing built, on its own, tells you almost nothing about supply conditions. That a lot of housing was built does not mean supply is unrestricted, since strong demand also produces high construction.

Nearly every popular mistake in this debate comes from reading a quantity as supply. The error appears in two forms, one temporal and one cross-sectional:

- **Over time.** A stretch of high construction is taken as proof that restrictions cannot be binding. But quantity can be high simply because demand was strong over a long period of time.
- **Across areas.** An area that builds a lot is assumed to be lightly regulated. But demand varies across space just as it varies over time, so an area can build a great deal because demand there is intense, not because its settings are permissive.

In both cases the quantity built, whether over time or across areas, is a general-equilibrium outcome, and on its own reveals little about the position of the supply curve.

The same mistake, pushed one step further, reappears on the price side as "we built a lot of housing and prices did not fall, so supply does not affect prices." We deal with that version in the next subsection and in Section 5, because it requires one extra idea: the spatial scale at which a price effect can be detected.

The distinction matters for the inquiry in two concrete ways. It determines which evidence about reform can be trusted, and it determines which indicators an evaluation framework should rely on. Both are developed below.

### 1.3 Knowing what works requires a counterfactual

Every claim about whether a supply-side policy works is a claim about a counterfactual: what prices, rents and quantities would have been under different settings. You cannot judge a reform, or the "state of supply", from what happened alone. You have to ask what would have happened otherwise. The single most useful habit in this area is to keep asking, of any fact: compared to what? The Commission may hear claims that reforms to supply-side regulations will not make much of an impact, or that supply has little or no effect on prices, and those making them will often not ask themselves this question.[3]

A credible counterfactual can come from only two places: a model, or an experiment. The gold-standard versions are an estimated structural model of the market or a clean quasi-experiment, and the rest of this submission leans on them. But even a simple mental version of these two tools can discipline reasoning, and keeping those in mind is most of what separates careful analysis from the traps that follow.

The simple model is the supply-and-demand diagram itself. Faced with a change in prices and quantities, it tells you what each force would have produced: quantities rising while prices fall, or grow only slowly, is more consistent with a supply shift, while quantities and prices rising together is more consistent with a demand shift. The model is what supplies the counterfactual here: absent a supply shift, you would not see quantities up and prices down together. This simple mental model is often only useful as a first pass, since both curves often move at once, but it is more useful than looking at them separately.

A simple experiment is a comparison. Take a place changed by a policy[4] and ask what you would have expected without it, which in practice means comparing it both to its own prior trajectory and a comparable set of areas. This provides something like an area-level fixed effect: it holds roughly constant the area's desirability, its position in the national labour market, and other demand-side characteristics that do not change

[3] The claim that supply has little or no effect on prices can only hold if the demand curve is horizontal: only if buyers would absorb any additional quantity at today's price and pay no less for it. As discussed below, that is only true for analysis at a local or hyperlocal level, where it is easy to substitute between areas. However, nationally the price elasticity of housing demand is low, so the demand curve is steep.

[4] Or somewhere hit by another exogenous shock, if the object of interest is the elasticity of supply rather than the effect of a particular reform.

sharply over the study period.[5] To run this comparison you need all three of these elements: a distinct shock, comparison to prior trends, and comparison to comparable areas.

Applied microeconometric quasi-experiments are the same idea executed more rigorously. For example, studies have compared upzoned cities against other areas. In both Auckland and Lower Hutt for instance, we have also observed both an increase in quantity and a decrease in prices relative to a counterfactual, consistent with predictions from our simple mental model about how a supply shock behaves. We set out the quantity effects of these reforms in Section 4 and the price effects in Section 5.

Almost all bad analysis falls into one of three traps, every one of them a failure to specify a credible, relevant counterfactual: they offer none at all, an irrelevant one, or one measured at the wrong scale.

**The first trap: no counterfactual.** The quantity of housing built is the joint product of demand and supply, so on its own it cannot tell you whether the market is working, whether a restriction binds, or what is driving prices. The starkest version is to read causation straight off a raw relationship between the two: regressing price on quantity, or quantity on price, or comparing prices and quantities across areas, when construction may be high somewhere simply because demand there is high.[6] This is the error behind the two claims in Section 1.2: that strong national construction proves restrictions cannot bind, and that a high-building area must be lightly regulated.

**The second trap: an irrelevant counterfactual.** A common argument identifies something broadly true about the world but carries an implicit counterfactual that is of no use to policy. Saunders and Tulip (2019), for example, find that the long decline in real interest rates explains much of the growth in Australian house prices. This is a useful and correct observation, but the question is what to do with it.

The demand curve has shifted out persistently over this period: real interest rates fell, population grew strongly,[7] and real incomes rose. But recording these facts says nothing about what policy should do. The mistake is to take them as given and then treat the supply curve as fixed. The supply curve is the object that determines how a demand shift translates into prices versus quantities: had it been more elastic, and shifting

[5] In the long run a city's fundamentals may shift, but over the horizons studied, comparison against its own prior trajectory and against similar cities holds demand-side factors roughly constant.

[6] Reporting such correlations is perfectly fine as motivation or descriptive statistics. The problem arises when they are treated as evidence for how supply and prices actually relate.

[7] Migration is partly a policy choice, and some argue migration settings should be coordinated with housing policy. That is a separate debate from the one here. Whatever migration policy is chosen, the supply curve still determines how much any given level of demand growth shows up in prices rather than in quantities.

outward over the same period, the same demand would have produced more housing and lower prices than eventuated. We cannot undo a global decline in interest rates, and would not wish to reverse rising incomes. The counterfactual in which rates never fell illuminates the past but is irrelevant to policy.

A policymaker told 'lower interest rates have driven past price growth' has learned nothing they can act on.

**The third trap: the wrong scale.** The first two traps apply to quantities and prices alike, whereas the third is specific to prices. A within-city experiment can identify a quantity effect, as we do in Campbelltown (Behrens et al. 2026), where building rose sharply after lot-size reform, but it cannot recover the price effect of that reform, because in a city households move freely between suburbs and so the price effects from new supply spread across the whole metropolitan area. This is what the literature finds and what structural models predict: the affordability benefits show up at the city level and accrue to households outside the rezoned area, through these migration spillovers (Rollet 2025).[8]

A thought experiment makes this concrete. Suppose we removed all zoning restrictions in the Eastern Suburbs of Sydney. The East would see a large increase in construction while the rest of the city would see little change in building, but both would see a similar change in prices, because they sit in the same integrated market. Even if the East's prices dropped slightly more than elsewhere, that gap would drastically understate the effect of the policy. Measuring prices in the rezoned area alone therefore tells you almost nothing. A credible price effect needs a model or a shock that moves an entire market, as we discuss in Section 5.

## 2. The central distinction: target bans, not burdens

*Addresses Terms of Reference matters 1 (approval processes) and 2 (land-use controls).*

Costly regulatory settings, in any domain, fall into two broad categories.

**Burdens** are settings where government still permits an activity but makes it slower, costlier, or more cumbersome: for example, long permit wait times, complex documen-

[8]There is credible evidence that supply shocks move prices at a hyperlocal level (Asquith et al. 2023; Li 2022; Mense 2025; Pennington 2021). But this literature measures prices or rents *relative to nearby comparators* and finds effects that decay to zero over very short distances. These are relative, local estimates, not general-equilibrium claims about prices across the whole market, which is the policy-relevant question. The asymmetry is partly methodological: there are good applied-microeconometric methods for detecting quantity spillovers within a city, but as yet no equally reliable methods for detecting price spillovers within one. Developing them would be a useful area for future research.

tation, heavy reporting, and the administrative process around developer contributions. Burdens operate on the intensive margin: activity still occurs, it just costs more, so you get somewhat less of it.

**Bans** are settings that prohibit an activity outright, or create enough uncertainty that a project realistically never proceeds: for example, zoning rules that make medium-density housing illegal across most residential land, minimum lot sizes, and height limits. Bans operate on the extensive margin: when they bind, profitable and welfare-enhancing projects simply do not happen.

Historically, almost all of the public conversation about housing supply regulation, and almost all of what industry raises in consultations, has concerned burdens. That emphasis is wrong, for two reasons.

**First, the measured cost of burdens is small, and the measured cost of bans is large.** Where researchers have tried to estimate per-dwelling costs, the difference is stark (Figure 2). The administrative cost of the development-approval system is on the order of $8,000–$15,000 per dwelling, and even the broadest estimate, the cost of *all* regulation other than zoning, is around $64,000 (Productivity Commission 2025). The "zoning tax," the gap between dwelling prices and the marginal cost of supply, attributable to land-use restrictions that prohibit development, ranges from around $460,000 for Sydney apartments to $660,000 for Sydney detached houses (Coates et al. 2025; Jenner and Tulip 2020; Kendall and Tulip 2018; Tulip 2022). None of these figures are exact or directly comparable, and they all rest on strong assumptions. But that even the broadest burden estimate remains several times smaller than the zoning tax tells you bans matter far more.

Figure 2: Estimated cost per dwelling of regulatory burdens versus zoning taxes

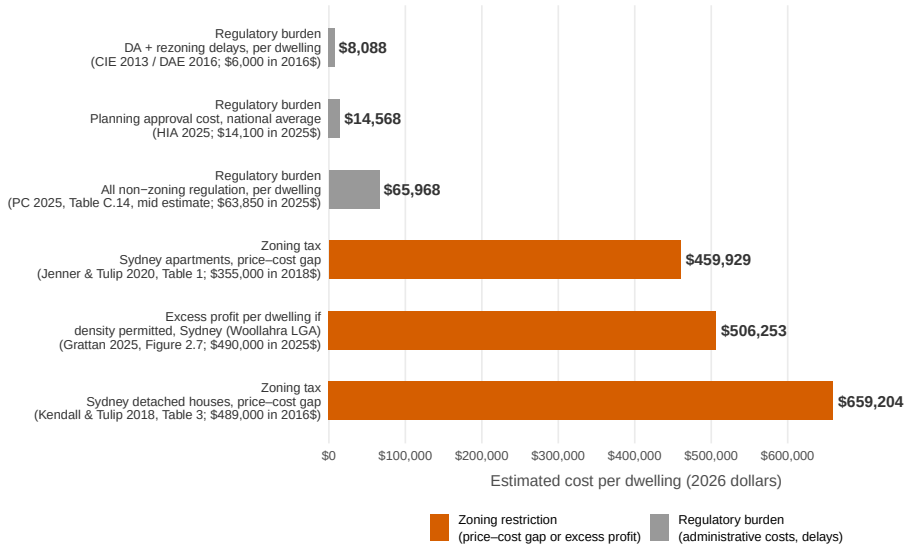


"Regulatory burden" covers the cost of regulation other than zoning, from the administrative cost of approvals to all non-zoning regulation combined; "zoning tax" is the gap between dwelling prices and the marginal cost of supply; "excess profit" is the potential profit above normal developer margins if density were permitted. Sources: Centre for International Economics (2013); Deloitte Access Economics (2016); Housing Industry Association (2025); Productivity Commission (2025) (Table C.14, mid estimate, all categories except zoning and land purchase); Jenner and Tulip (2020); Coates et al. (2025); Kendall and Tulip (2018). All values CPI-adjusted to 2026 dollars using ABS All Groups CPI (Cat. 6401.0).

**Second, the empirical evidence runs entirely one way.** There are now at least half a dozen credible studies, using real-world policy changes and careful counterfactuals, finding that removing bans leads to large increases in the quantity of housing built. Section 4 reviews them: the Australasian natural experiments in Auckland, Lower Hutt, and Campbelltown, and corroborating international studies from Zurich, São Paulo, and New York.

By contrast, there are essentially no studies showing that speeding up approvals, shortening planning documents, or reducing engagement costs produces a construction response in anything like the same ballpark.[9]

This asymmetry deserves weight precisely because there are strong political incentives to demonstrate that burden-reduction works. Governments experiment with burden-reduction regularly (fee waivers, standardising approval pathways, etc.), so the natural

[9]Two recent Los Angeles studies do find supply effects from the approval process (Gabriel and Kung 2023; Soltas and Gruber 2026). Their mechanism largely operates through the cost of *discretionary* approval and its associated uncertainty, which functions as a probabilistic ban. The policy response their own authors draw out, a shift to as-of-right permitting, is itself the removal of a ban. They are also estimates from a uniquely discretionary US planning setting.

experiments exist. If these reforms produced construction responses comparable to those documented for upzoning, the evidence would be difficult to miss. The most likely explanation is that burden-reduction, while worthwhile, is second-order to the removal of binding restrictions on what can be built.

There is also a structural reason why industry consultation tends to emphasise burdens over bans. Existing firms can describe the costs they face today, but they cannot describe the counterfactual market that reform would create. The firms best placed to describe the gains from lifting a ban are those that would enter or scale up after the reform. By definition, they are not at the table.

None of this is an absolutist claim: it is simply where the weight of the evidence currently sits, and new research could shift it. Approval processes should still be improved where possible. But reform costs political capital, and that capital should be spent first on the bans, where the evidence is strongest and the payoff is largest.

## 3. Approval times and approval rates are poor metrics, because they reflect what is allowed to be built

*Addresses Terms of Reference matter 1 (approval processes).*

Much of the supply-side debate, on both sides, fixates on two statistics: approval timelines and approval rates. Planning bodies cite high approval rates (often 90 per cent or more) as evidence that planning is not a barrier to new supply (Local Government NSW 2024), while industry cites long timelines (180-plus days in some councils) as evidence that it is (NSW Department of Planning, Housing and Infrastructure 2024). Both statistics are shaped by what the rules permit in the first place, and the problem is fundamentally one of numerators and denominators.

The approval rate is a ratio: approvals divided by applications. When a council bans medium-density housing across its residential land, no rational developer lodges an application to build townhouses. This affects both the numerator *and* the denominator of an approval rate. As a result, the applications that do get lodged are overwhelmingly those expected to succeed under the current regulations. The result can easily be a high approval rate in that council. Consider a jurisdiction approving 50 per cent of 1,000 applications versus one approving 99 per cent of 100: there is no way to tell from these numbers alone which is more constrained by policy.

The same logic applies to timelines. A restrictive council may process a small number of straightforward, compliant applications quickly, while a council that permits more

ambitious development receives more complex applications that take longer, so the headline number may *negatively* correlate with supply restrictions.

These metrics are therefore, at best, uninformative about supply-side regulations, and at worst actively misleading.

The clearest demonstration comes from a genuine supply-side reform that led to an increase in construction. In Campbelltown, minimum lot sizes were cut from 500 square metres to as low as 150 square metres in 2014. What happened to these two metrics? Approval timelines did not fall - if anything they rose - as the volume and complexity of applications increased. Approval rates also barely moved. What changed was the *number* of applications, as the ban was lifted. This is what drove the increase in construction.

Figure 3: DAs in Campbelltown around the 2014 minimum-lot-size reform.

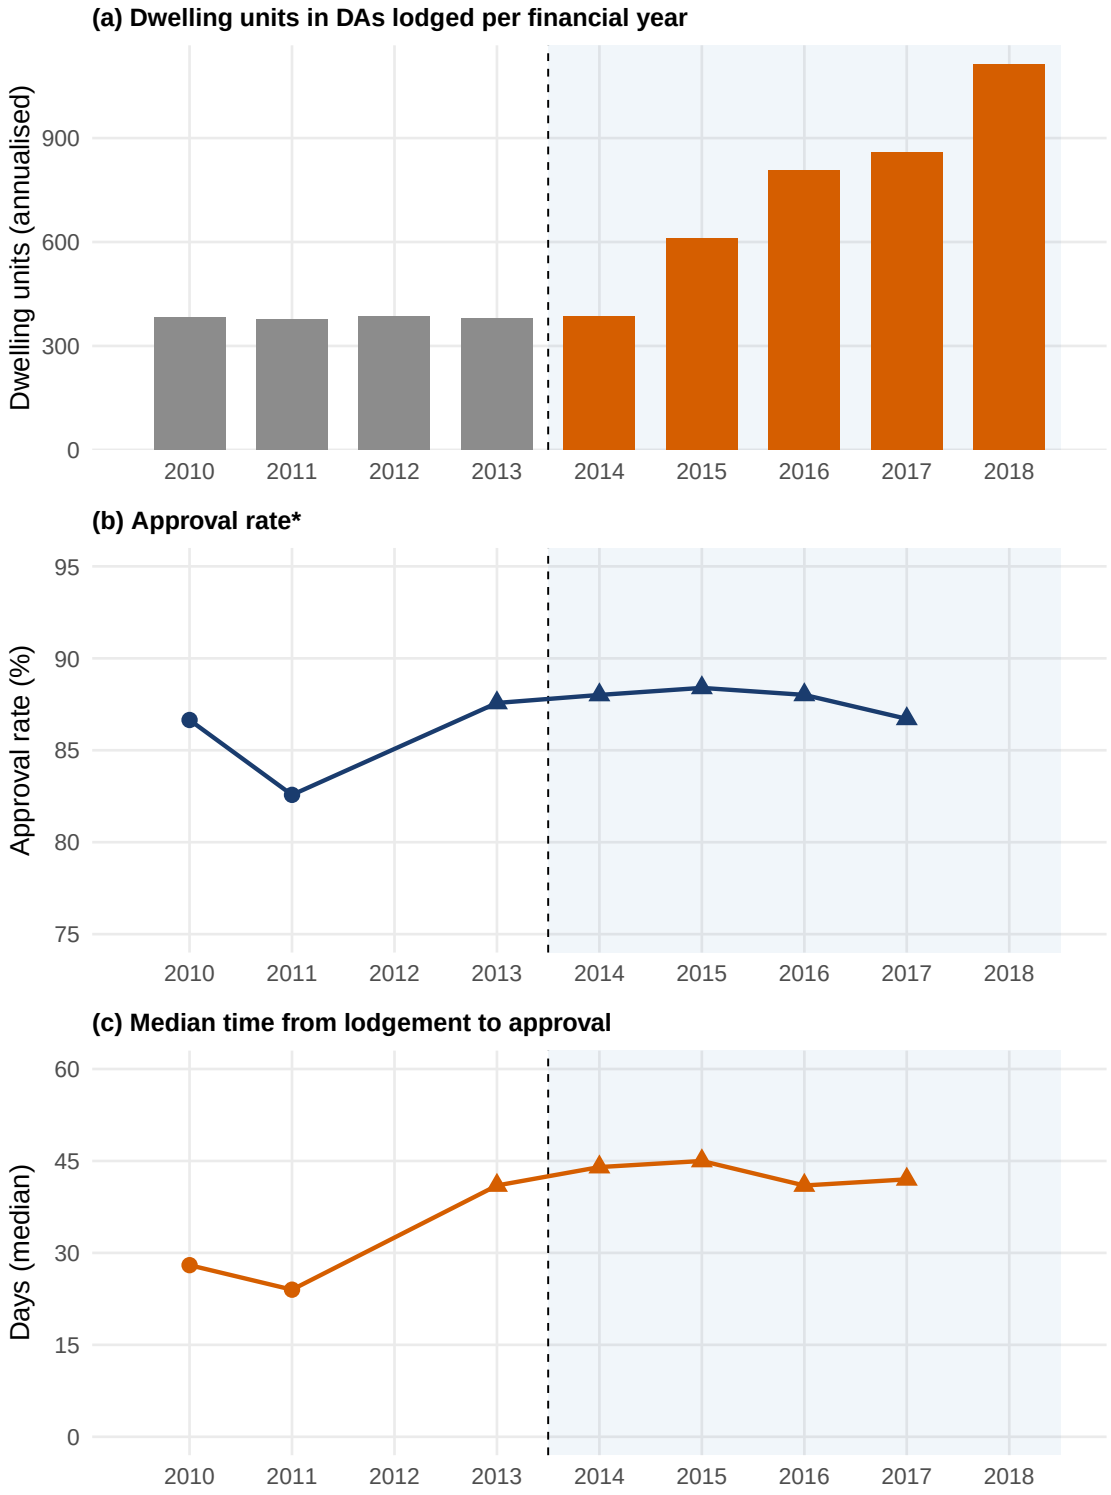


Panel (a) shows estimated dwelling units in new-dwelling DAs lodged with Campbelltown council, annualised from months with complete register coverage. Dashed line marks the 2013 reform; blue shading the reform period. Panels (b) and (c) use years with complete register coverage only (2012 excluded). Approval rate is the share of lodged DAs that appear in the approved register. Source: Behrens et al. (2026). Preliminary results from unreleased work, subject to change.

This can be put practically. In Woollahra, a development application takes around 184 days to be determined (NSW Department of Planning, Housing and Infrastructure 2024), while subdivision is not permitted across much of the area (Woollahra LEP 2014, clause 6.5). The burden adds cost and uncertainty to the few projects that can proceed. The ban eliminates an entire class of projects altogether. Approval-timeline statistics capture only the first issue and are silent on the second.

More broadly, if we believe a regulation serves a clear purpose, whether managing flood risk, protecting heritage, or ensuring structural safety, then we should let government apply it properly. That takes time, and it is not obvious that faster is always better: cutting approval timelines to zero, or approving everything, would defeat the purpose of having regulations at all. The goal should not be to minimise the time government spends administering regulations we want applied, it should be to ask whether the right things are being regulated in the first place.

The practical implication for the Commission is direct. A "prioritised list of reforms" should not be organised around faster approvals, and an evaluation framework should not use approval rates or timelines as indicators of regulatory performance. The relevant question is *what* the rules permit to be built.

## 4. The reform with the strongest evidence: legalise medium-density housing in established areas

*Addresses Terms of Reference matter 2 (availability and use of land for housing).*

Legalising medium-density housing in established residential areas is the most effective lever available to governments. This section reviews the applied studies that demonstrate this, which reach the same conclusion as recent Grattan Institute work (Coates et al. 2025).

The reforms work through two overlapping channels:

1. **More dwellings on a given parcel**, via reform to minimum lot sizes and subdivision rules.
2. **More floorspace on a given parcel**, via reform to height limits and floor-area ratios.

These are not mutually exclusive, and the best reforms enable both.

### Auckland: the Unitary Plan (2016)

In 2010, the New Zealand Government forcibly amalgamated eight Auckland councils into one and directed it to produce a unified land-use framework. That process led to the Auckland Unitary Plan (AUP) - drafted in 2013 and made operative in 2016 - which replaced a patchwork of legacy zoning rules with a consistent, region-wide system aimed at enabling more housing close to jobs and transport.

The AUP's main innovation was its large-scale upzoning of suburban land. Where most of Auckland had previously been restricted to detached single-family dwellings, around three-quarters of its residential areas were rezoned to allow townhouses, duplexes, and small apartment blocks, with greater height and density permitted along transport corridors.

The construction response was so large that it is visible in the raw, aggregate data: per-capita construction matched the national rate before 2016 and exceeded it in every subsequent year (Figure 4). Almost all of the additional dwellings were townhouses and terraced housing in newly upzoned areas. The formal estimates construct credible counterfactuals to estimate the causal impact of these changes. Greenaway-McGrevy and Phillips (2023) find the reform produced a genuine development boom, around 22,000 additional dwellings permitted in the first five years, equal to about 4 per cent of the existing stock. Subsequent analysis that uses less conservative assumptions and a longer time period puts this number at 52,200 (Greenaway-McGrevy 2026).

Figure 4: Dwelling consents per 1,000 residents: Auckland versus non-upzoned New Zealand.

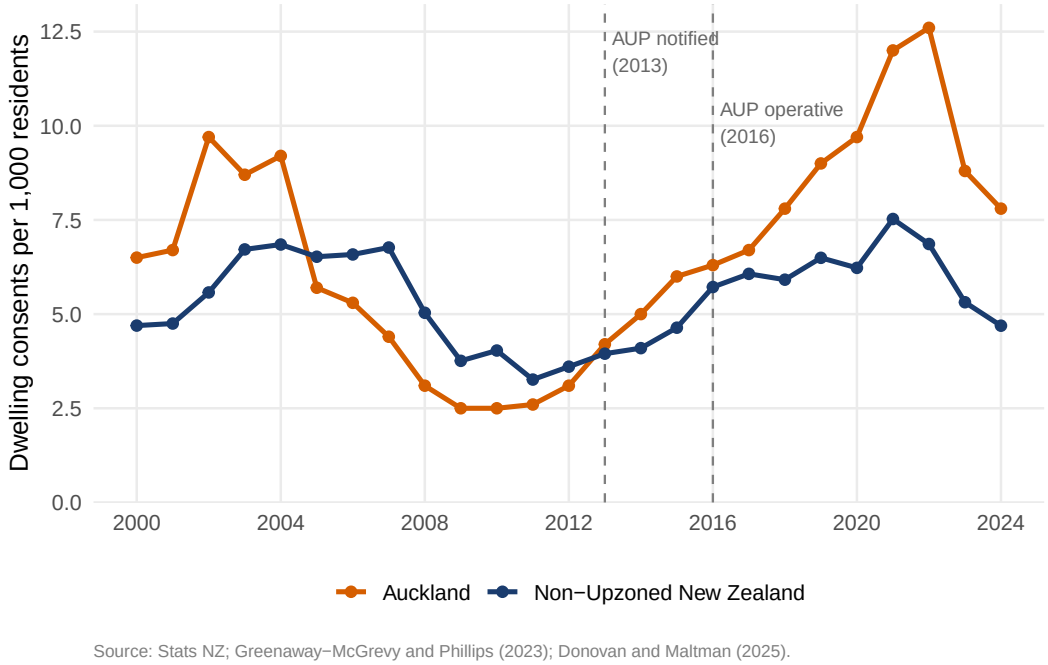


Source: Stats NZ; Greenaway-McGrevy and Phillips (2023); Donovan and Maltman (2025).

The Auckland findings have attracted critiques, primarily in blog posts and on social media, that have occasionally been cited in formal planning processes. These critiques have been addressed in detail[10], and the opinion of the economics community is heavily one-sided: a New Zealand Association of Economists survey found 96 per cent of respondents agreeing that land-use restrictions reduce the quantity of housing built and 94 per cent agreeing they worsen affordability (Wesselbaum and Dela Cruz 2024).

## Lower Hutt: upzoning (2017)

Lower Hutt provides a particularly clean test of the bans-versus-burdens distinction because, unusually, it ran both experiments. In 2012, facing a shrinking population, the council introduced a Development Stimulus Package that waived building consent fees and development contributions: the archetypal burden-reduction policy. Builders largely did not respond: just 294 applications used the scheme over six years (Maltman and Greenaway-McGrevy 2025; Tso 2019).

Then from 2017, the Council progressively rezoned residential land from a "General

[10]Donovan and Maltman (2025), published in *Land Use Policy*, review the critiques in detail and find them to have little to no merit: they misunderstand the methods, rely on inappropriate comparisons, and index data to certain time periods to hide policy effects.

Residential" zone (one dwelling per lot) to a "Medium Density Residential" zone (allowing up to three storeys and multiple units per lot, including townhouses and apartments). Maltman and Greenaway-McGrevy (2025) use a synthetic-control method, constructing a weighted combination of comparable New Zealand cities to estimate what would have happened without the reform. Figure 5 shows the result: actual consenting diverges from the synthetic counterfactual after 2017, and over 60 per cent of the dwellings consented over the following six years are attributable to the upzoning. (You probably do not need a synthetic control to see the effect).

Figure 5: Lower Hutt dwelling consents per 1,000 residents: actual versus synthetic control.

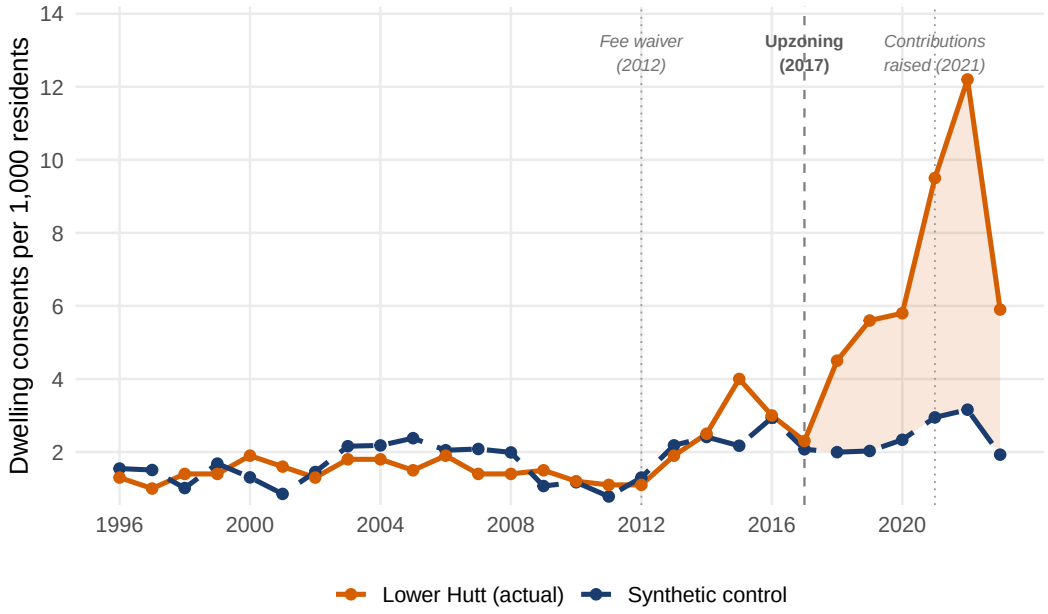


Source: Maltman and Greenaway−McGrevy (2025). Shaded area shows the estimated treatment effect.

Shaded area shows the estimated treatment effect. Source: Maltman and Greenaway-McGrevy (2025).

## Campbelltown, Adelaide: minimum lot size reform (2014)

This case study draws on forthcoming work by Behrens et al. (2026), which provides the first causal evidence on the effects of zoning reform on housing construction in an Australian city. This submission may be published before we release this research, so we describe the setting and findings in some detail here. We note however that this work is forthcoming and therefore is subject to potential revisions.[11]

Campbelltown is a middle-ring council in Adelaide's eastern suburbs, roughly 10 kilo-

[11]That caution applies to every figure and data point related to Campbelltown in this submission.

metres from the CBD. Before 2014, its Development Plan prescribed a blanket minimum lot size of 500 square metres (350 square metres near centres), with no dwelling-type-specific parameters for medium-density housing. In 2014, a Development Plan Amendment reduced minimum lot sizes to as low as 150 square metres for row dwellings and apartments within 400 metres of designated centres and transport corridors. This made it feasible to subdivide a standard 600–800 square metre detached-house lot into two to five smaller lots. The policy produced a large increase in infill development, which resulted in community opposition. This drove a policy change in 2019 which reversed the changes, raising the minimums from 150 to 250 square metres.

Panel (a) of Figure 6 shows the distribution of the land use changes across Campbelltown. Panel (b) shows that the supply response was concentrated in the highly treated areas, where lot sizes were cut to as low as 150sqm. Panel (c) shows that the minimum lot size was a binding constraint. Before the reform, output lot sizes in land divisions clustered around 350 square metres (the near-centre minimum; the blanket minimum was 500). During the reform, output lots shifted to the 150–250 square metre range made possible by the new rules. After the 2019 reversal, lots re-bunched at 250 square metres, the new binding floor. Panel (d) shows that there was a large and reasonably immediate increase in subdivisions after the policy changes, which returned to normal levels after the reform was reversed.

Figure 6: Campbelltown zoning reform and land division response.

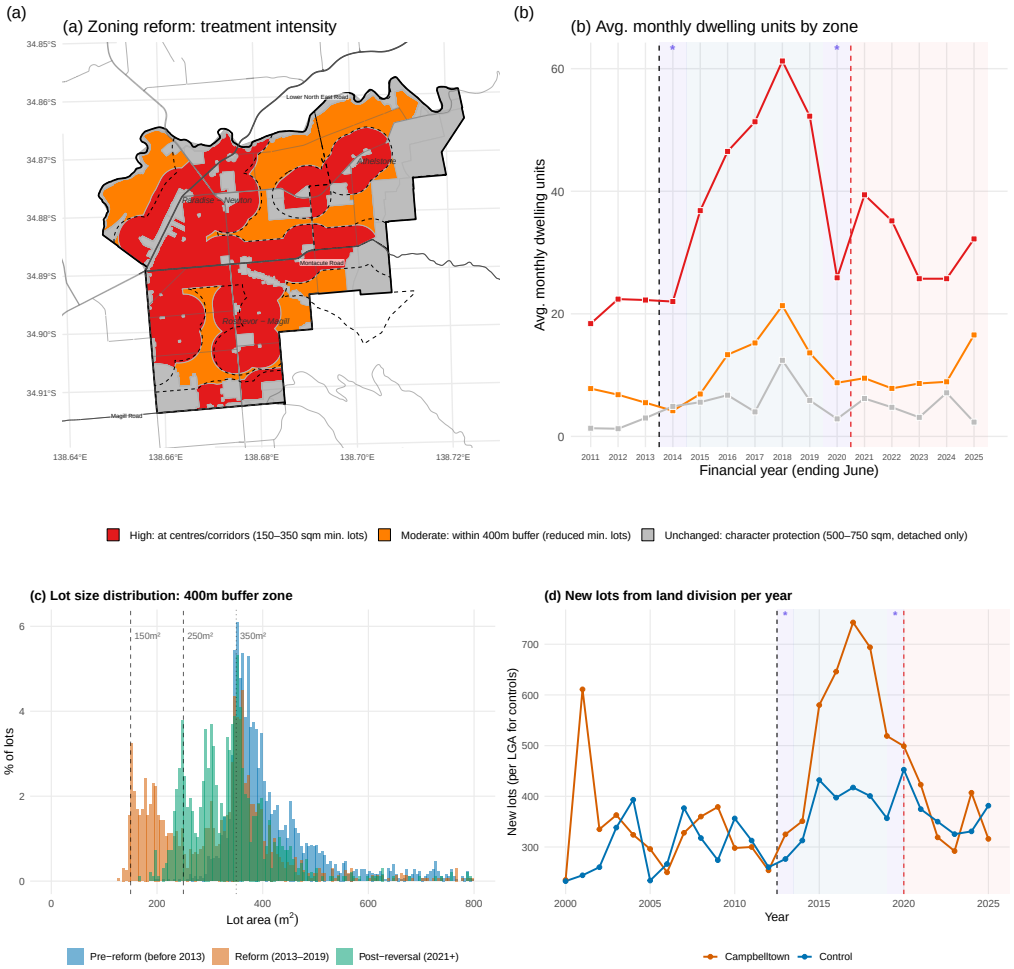


Panel (b): months with fewer than 15 total DAs excluded as incomplete. Panel (c): dashed lines mark the 150 and 250 $m^2$ thresholds; dotted line marks the pre-reform minimum of 350 $m^2$. Panel (d): control is the per-LGA average across six inner and middle Adelaide LGAs. Blue shading = reform-active (2013–2019); red = post-reversal (2021–2025); purple = announcement/grandfathering. Source: Behrens et al. (2026). Preliminary results from unreleased work, subject to change.

We estimate the causal effect of the reform using building-approval data across metropolitan Adelaide.[12] Figure 7 shows the event-study estimates: building approvals rise sharply during the reform period, and the effect fades after the 2019 reversal, providing a natural test that strengthens our finding that the construction boom was driven by policy and not some other factor.

Two further findings from the Campbelltown data bear on how the Commission should think about what reforms might have success at raising supply. **First, going broad**

[12]In the paper we also test for effects using DA data, lot divisions data, completions data, and population data, and find similar results.

Figure 7: Event study: effect of the Campbelltown minimum-lot-size reform on total building approvals (FY2008–2024).

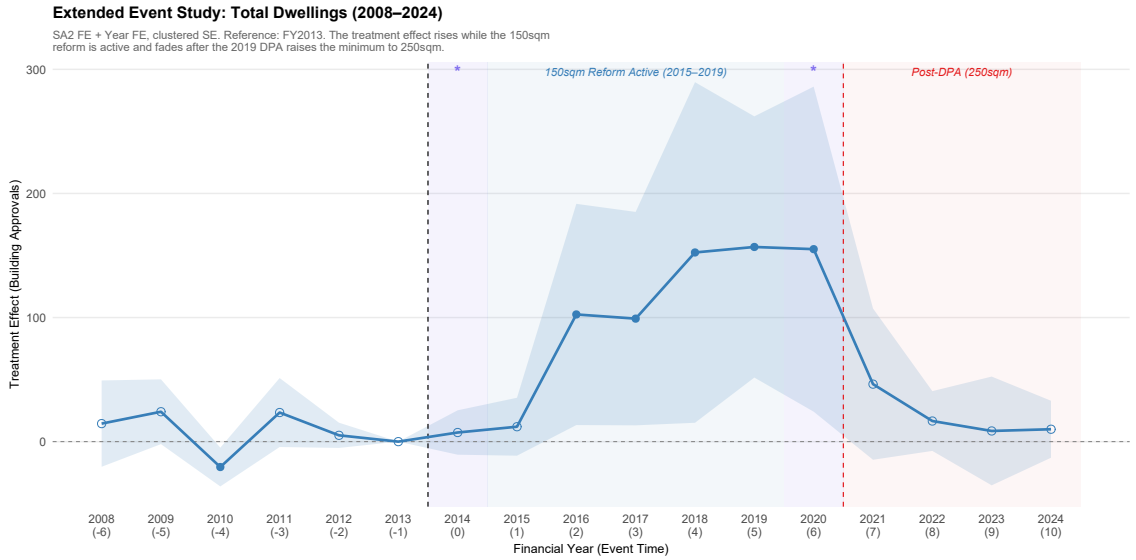


Source: Behrens et al. (2026). Preliminary results from unreleased work, subject to change.

**matters, because not all sites can or will be developed in the short run.** Using detailed microdata linked to DAs and land divisions, we are able to track every lot in the council through time. We find both extensive and intensive margin effects of the policy on development.

The share of lots developed each year in the highest treated area (the reformed 400m buffer zone) rose from 0.80 to 1.11 per cent after the reform, against an essentially flat 0.29 to 0.31 per cent in the untreated part of the council. On the intensive margin, the yield of each development also increased from 1.45 to 2.23 in these highly treated areas. This suggests that even a major liberalisation of land use rules might see only around 1 per cent of lots developed each year, but can still generate the large quantity increase seen in Figure 7.

Second, redevelopment concentrates in older housing on large, cheap lots. The lots most likely to be redeveloped are those that are cheap per square metre of land and carry little existing structure relative to their land (a low floor-to-lot ratio). Figure 8 shows the pattern on the full universe of Campbelltown lots: the share subdivided each year rises with the projected development margin and with lot size, and falls with the lot's price per square metre. This is directly useful for the Commission's interest in where "commercially feasible capacity" actually lives: the redevelopment pipeline runs through older, under-capitalised housing stock on large and relatively inexpensive lots.[13]

[13]We have also heard anecdotes from Lower Hutt that this was the case there too. Developers would

Figure 8: Share of all Campbelltown lots developed per year, by characteristic quintile.

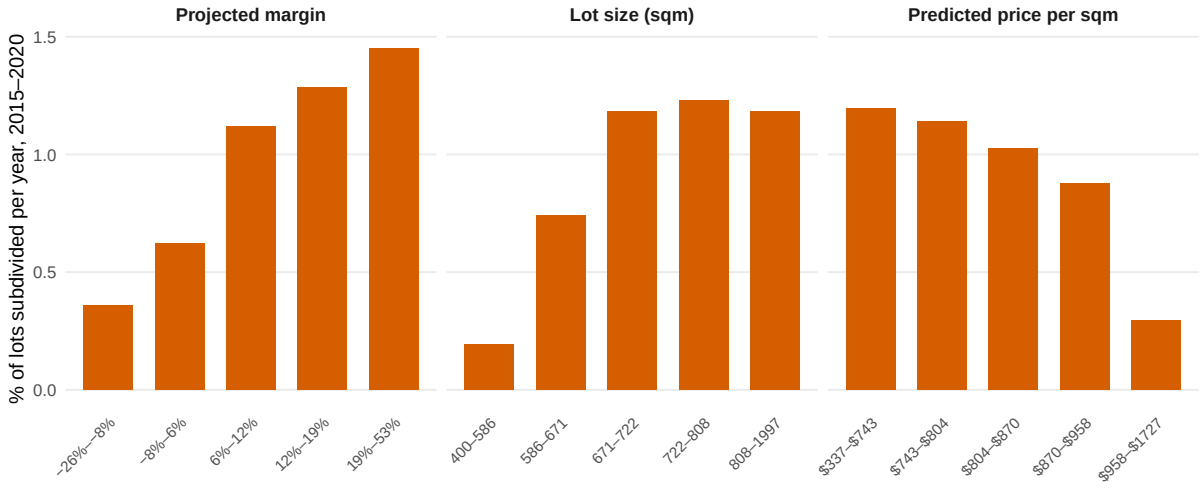


The denominator is every lot of 400–2,000 square metres in Campbelltown. The outcome is whether the lot was developed by a land division between May 2015 (where our data begin) and the end of 2020, expressed per year. The projected margin is the developer feasibility margin under the post-reform minimums. Source: Behrens et al. (2026). Preliminary results from unreleased work, subject to change.

## International evidence

These Australasian findings are reinforced by studies in very different institutional settings:

- **Zurich:** Büchler and Lutz (2024) exploit the staggered timing of multiple rezonings across Zurich municipalities and find that upzoning raises floorspace and dwelling numbers, with no local rent effects because it dissipates across the city.
- **São Paulo:** Anagol et al. (2025) use block-level variation from an upzoning in São Paulo and find more multi-family permitting and lower prices in blocks granted higher allowable density.
- **New York:** Rollet (2025) builds a dynamic general-equilibrium model and finds that removing density restrictions could raise total floorspace by around 71 per cent over 40 years, against 13 per cent under current rules. The same model finds that a 20 per cent reduction in construction costs lifts growth only from 13 to 15 per cent. This is the bans-versus-burdens point quantified inside a single model: relaxing the ban does far more than reducing marginal costs (e.g. a "burden").

buy large amounts of land on a single street where land was cheap in the middle/outer suburbs and would work down the street post-upzoning.

## Why removing restrictions increases the quantity of housing built

A developer can build when a project's revenue clears its costs by enough to meet a hurdle margin. Take a suburban lot that yields one detached house under a binding minimum lot size. Relax the minimum and the developer can instead build several townhouses. Three things move at once:

- The sale price *per dwelling* falls (each townhouse comes with less land than the detached house, and is probably smaller).
- The *number* of dwellings in the project rises.
- *Costs* rise (there are more dwellings to build, and land costs often increase due to option value premiums).

Whether the project becomes feasible depends on a horse race between these forces: the project goes ahead if the gain from selling more dwellings outruns the fall in price per dwelling and the rise in costs.

Two features of this race do most of the work, and both are convexities. The first is mechanical: the number of dwellings that fit on a parcel rises quickly as the minimum falls.[14] The second is in prices: a dwelling's price per square metre of land rises as lots get smaller, because buyers are paying for a dwelling in a location rather than for raw land.

Figure 9 shows the price convexity directly, plotting price per square metre against lot size for development activity in Campbelltown across the three regulatory periods, with developer purchases (large parent lots, in blue) and developer sales (newly created small lots, in orange).[15] Per-square-metre prices rise steeply as lots shrink, and the steepness accelerates below roughly 200 square metres: two 150 square metre lots are worth substantially more, in total, than one 300 square metre lot. The average dwelling on a 150 square metre lot sold for around $500,000, so two of them are worth roughly $1 million, against around $580,000 for one dwelling on a 300 square metre lot.

Costs, by contrast, contain much less of this offsetting convexity: construction costs more or less scale with floor area rather than land area, and the fees and connections attributable to each extra dwelling are comparatively small.

Where the schedule turns steep is where the regulation binds. Campbelltown's pre-reform 350 square metre minimum confined subdivision to the flat region of the curve,

[14] For example, on an 800 square metre lot, a 250 square metre minimum allows 3 dwellings whereas a 150 square metre minimum allows 5.

[15] Preliminary results from unreleased work, subject to change (Behrens et al. 2026).

where splitting a lot gains little and few projects clear the hurdle. Cutting the minimum to 150 square metres let developers sell output on the steep part, opening a large gap between what they pay for one big lot and what they receive for several small ones. And the 2019 reversal to 250 square metres, although it looks like only a partial unwind on paper, pushed output back off the steep part of the curve, which is why it extinguished most of the supply response (Figure 7).

Figure 9: Developer transactions in Campbelltown, by reform period.

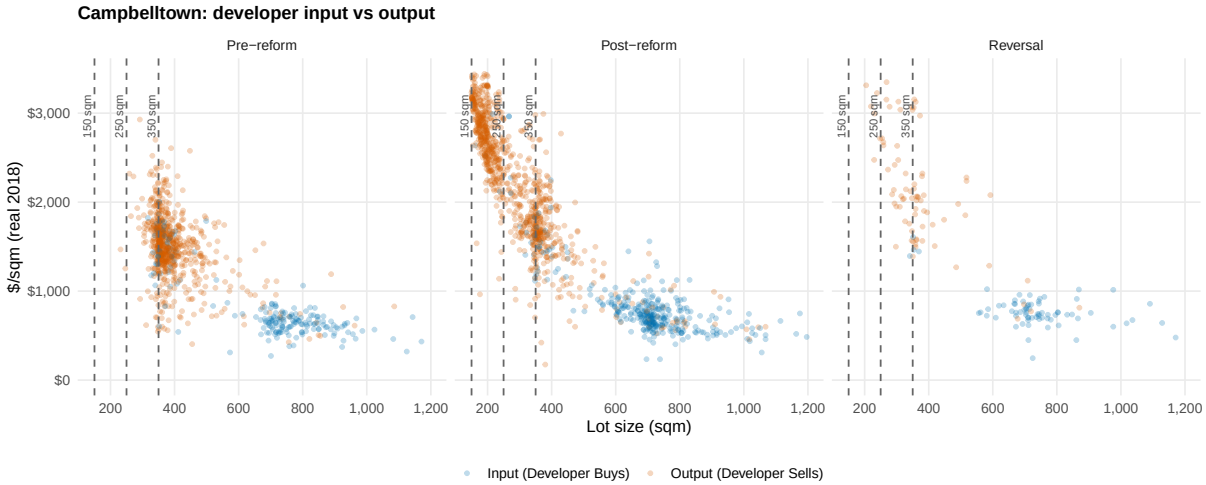


Each panel plots per-square-metre transaction prices against lot size for one regulatory period, with developer input purchases (large parent lots) and output sales (newly created child lots). Source: Behrens et al. (2026). Preliminary results from unreleased work, subject to change.

Minimum lot size reform works by dividing land. A second lever works by stacking floorspace: allowing two and three storey townhouses and low-rise apartments on the same parcel. It is worth being clear about why going up matters, because the logic is the same horse race viewed from the other side.

When a parcel that held one dwelling is allowed to hold a three-storey building with six units, the cost of the land is spread across six homes instead of one. Cooper et al. (2026) show this in Auckland. In the zones the 2016 reform liberalised most, the price per square metre of *land* rose, yet the land cost per square metre of *floorspace* fell, because each dwelling now needed far less land. In the zones left restrictive, land cost per unit of floorspace kept climbing. Even where upzoning raises the price of land, it lowers the land cost per home, which is the policy relevant fact.[16]

[16] A common worry is that upzoning undoes this benefit by raising land prices. At the level of an individual parcel this is well documented, but not enough to reduce the development premium. The net effect of upzoning an entire city is less settled, because two forces push against each other: option value raises the price of each developable site, while the expanded supply of developable floorspace pushes land prices down across the market.

## 5. Supply shocks have large effects on prices

The previous sections establish that removing zoning restrictions shifts the supply curve and increases the quantity of housing built. This section addresses the price side: supply shocks substantially reduce prices and rents, when measured correctly. As Section 1.2 argued, you cannot regress quantity on prices, and you cannot take a sub-metro supply shock and make general equilibrium claims about the prices. Given those constraints, there are two - and only two - credible research designs for recovering the effect of supply on prices.

**The first is to write down a model of the housing market, estimate it, and use it to run the counterfactual.** Rollet (2025) is the best example. He builds a dynamic general equilibrium model of floorspace supply and demand, estimated on a parcel-level panel of buildings and zoning in New York City, and finds that removing all density restrictions would reduce citywide rents by around 18 per cent in the long run, with welfare gains equivalent to roughly a 15 per cent income increase, disproportionately benefiting lower-income households. Greenaway-McGrevy (2025) does a simpler exercise for Auckland, estimating that the Unitary Plan will push dwelling prices down by between about 15 and 27 per cent relative to a counterfactual without the upzoning.

Saunders and Tulip (2019) is the canonical Australian model, finding that a 1 per cent increase in national supply pushes prices down by around 2.5 per cent. Given that the quantity effects documented above can exceed 5 per cent of the stock within a few years, this implies upzoning can quickly put large downward pressure on prices.

**The second is an experiment: find a policy change that delivers a supply shock to an entire market**, and compare that market against a credible counterfactual. City-wide shocks are the cleanest such instance, because although people can and do move between cities, the elasticity of movement is much lower than at a local level. In other words, the inner-west and inner-east of Sydney are highly substitutable, and while Sydney and Melbourne still are, they are much less so.

Auckland is a good example because the Unitary Plan was a city-wide shock. The cleanest price evidence comes from rents, for two reasons: rents are not contaminated by any redevelopment option value into land prices, and they are less affected by macroeconomic conditions such as interest rates. Greenaway-McGrevy and So (2025) compare Auckland's rents to other New Zealand cities that did not upzone. Their central estimate is that, after quality adjustment, Auckland rents ended up around 28 per cent lower than they would have been without the reform.

The more important finding from Greenaway-McGrevy and So (2025) is that Auckland's

rents grew less than every other metropolitan region of New Zealand since zoning reform. In other words, whichever city you compare Auckland to, rents have risen by less since upzoning. Lower Hutt is a second case: Maltman and Greenaway-McGrevy (2025) find rents fell by about 21 per cent relative to the synthetic counterfactual.

Two features of these estimates deserve note. First, the effect builds over time. The effect of a supply shift accumulates as the flow of new construction continues to be high relative to its counterfactual. Second, these city-wide estimates are, if anything, conservative as they miss spillovers to other cities. Insofar as people moved to Auckland due to cheaper housing, this will push up prices in Auckland, and down prices in other cities.

The affordability effects are so strong that they are visible in raw data. As Figure 10 shows, rent-to-income ratios in Auckland have *fallen* since the Unitary Plan, while in the rest of New Zealand they have risen (Donovan and Maltman 2025). In real terms, median house prices in Auckland are around 13 per cent lower than in late 2016, while real prices in the rest of the country rose roughly 27 per cent over the same period (Donovan and Maltman 2025).

These figures likely understate the true affordability gain. New housing tends to sell at a premium because it is higher quality than the existing stock, which pushes up the median price even as the reform makes housing cheaper on a quality-adjusted basis. This compositional effect is documented across multiple settings, including Lower Hutt and Campbelltown (Behrens et al. 2026; Maltman and Greenaway-McGrevy 2025), so the raw price declines in Figure 10 are conservative.

Figure 10: Housing affordability outcomes in Auckland versus the rest of New Zealand.

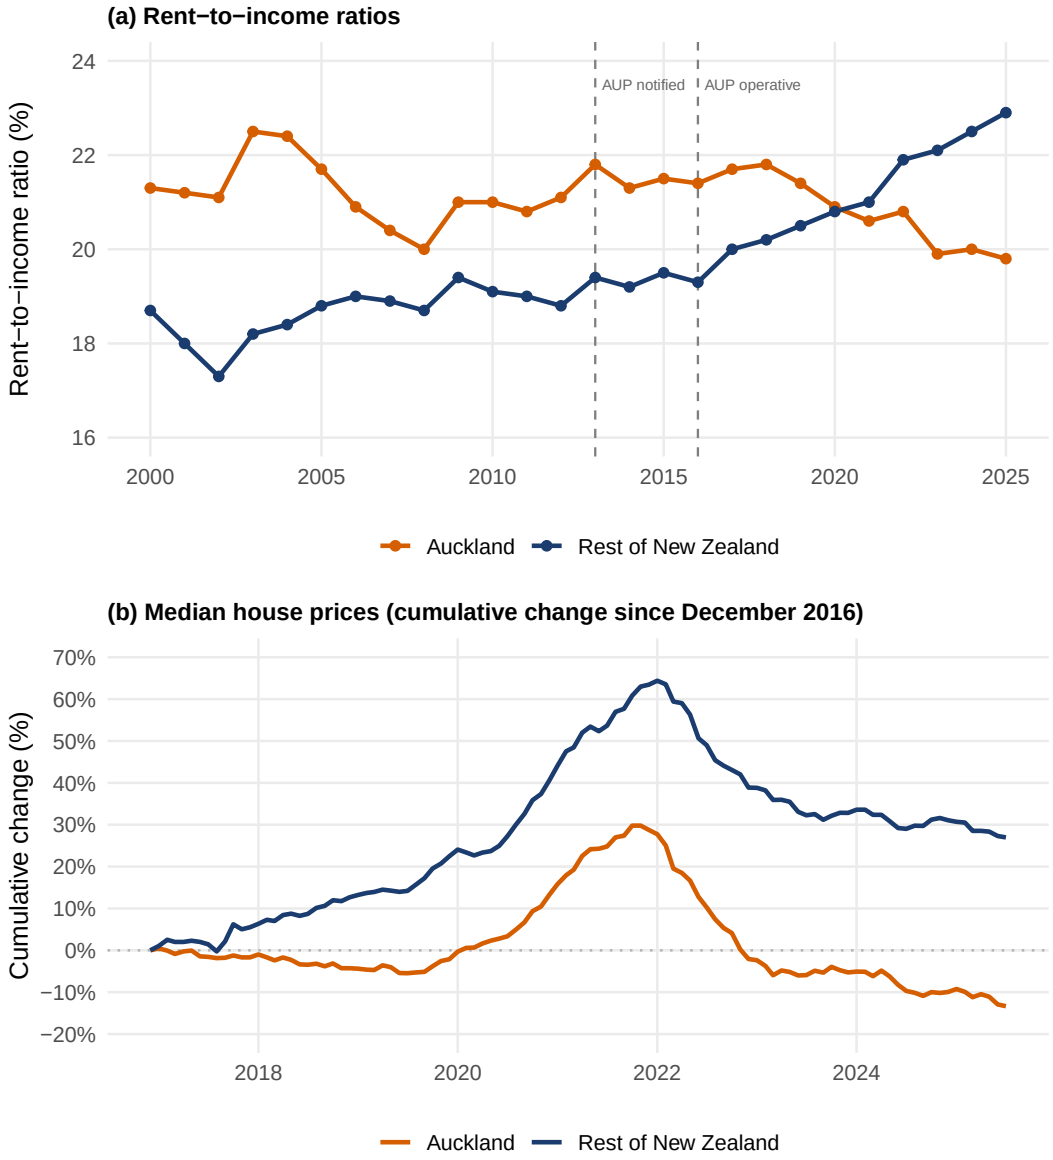


Source: Stats NZ; REINZ; Donovan and Maltman (2025).

Source: Stats NZ; REINZ; Donovan and Maltman (2025).

The takeaway from these research designs is simple. To say something credible about the effect of supply on the level of prices, you need either a model that imposes the equilibrium relationship, or a shock that moves the whole market, so that the market-level price is the thing actually being treated. Any other attempt to say anything about supply on prices is econometrically invalid.

For example, Murray and Limb (2023) regress dwelling prices and quantities on lagged zoned capacity across 19 Brisbane activity centres and conclude that increasing permitted density has no effect on prices. The design cannot support that conclusion: with

contemporaneous price and quantity entering on both sides and no source of exogenous variation, the regressions recover a mix of supply and demand elasticities rather than a causal effect.

## 6. Infrastructure and developer contributions: get the prices right, but keep them in perspective

*Addresses Terms of Reference matter 3 (infrastructure and developer contributions).*

The Terms of Reference rightly ask about infrastructure delivery and the developer-contributions model. Efficient, predictable infrastructure funding matters, and poorly designed contributions can distort the location and timing of development. But the Commission should be careful not to over-weight contributions reform relative to the bans listed above.

As discussed in Section 4, Lower Hutt waived fees and contributions from 2012 with little effect, then upzoned from 2017 with a large construction response. In 2021, the council *raised* development contributions substantially, increasing estimated revenue from $10 million to $37.5 million over ten years. Developers protested, warning costs would be passed on to buyers (Tso 2021). The following year, Lower Hutt approved the most dwellings on record, around six times its pre-upzoning average (Maltman and Greenaway-McGrevy 2025). The upzoning had created enough profitable development opportunities that construction continued even as contributions rose.

The binding constraint was the ban, not the burden.

## 7. High construction costs are a reason for more reform, not less

A common objection is that construction costs are now so high that building more is simply not feasible, rendering regulatory reform beside the point. This belies the fact that costs are only one half of the profitability equation. What matters is *profitable opportunities to build*.

Over the full global cost cycle of the late 2010s and early 2020s, New Zealand built more housing per capita than Australia (Figure 11; Maltman 2025), despite a weaker macroeconomic position and despite never having outbuilt Australia previously this

century.

Figure 11: Dwelling approvals/consents per 1,000 residents: Australia versus New Zealand.

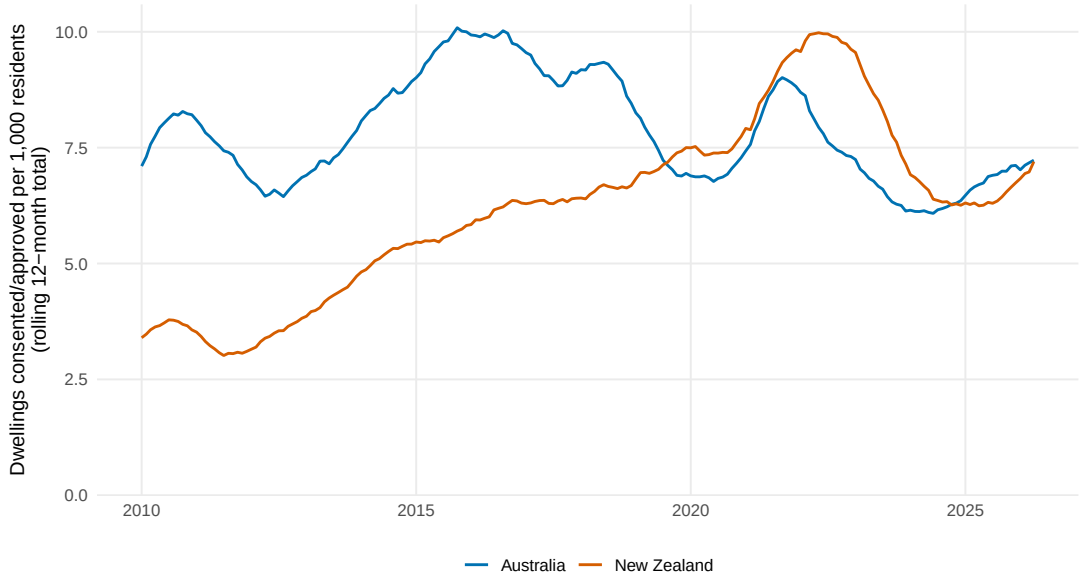


Source: ABS Building Approvals; Stats NZ Building Consents; World Bank.

When more intensification is legal in high-value locations, more projects clear the hurdle rate even at elevated costs. Rollet (2025)'s model makes the same point directly: relaxing the density ban does several times more for the quantity of housing built than even a sizeable fall in construction costs (Section 4).

This objection also conflates the building cycle with the structural floor beneath it. Reform does not promise to override a downturn. What it does is raise the level of construction at every point in the cycle: more in the upswing, and a shallower trough when conditions deteriorate. Auckland shows this: even during the 2022–23 downturn, with elevated material costs and rising interest rates, construction never fell below its pre-upzoning average (Donovan and Maltman 2025). High costs and elevated interest rates are therefore an argument for doing *more* reform, not less: when the hurdle rate is high, expanding the set of profitable opportunities matters most.

In fact, the New Zealand experience suggests zoning reform does not just increase the quantity of housing built. It might also make the construction sector more productive.

New Zealand's construction productivity rose over 15 per cent during the 2010s, even as it fell in Australia, the United Kingdom, Canada, and the United States (Figure 12). Maltman (2025) shows that these gains were concentrated in the subsectors and cities affected by zoning reform, and aligned to the timing of each reform (Figure 13).

Figure 12: Residential construction productivity: New Zealand versus Australia.

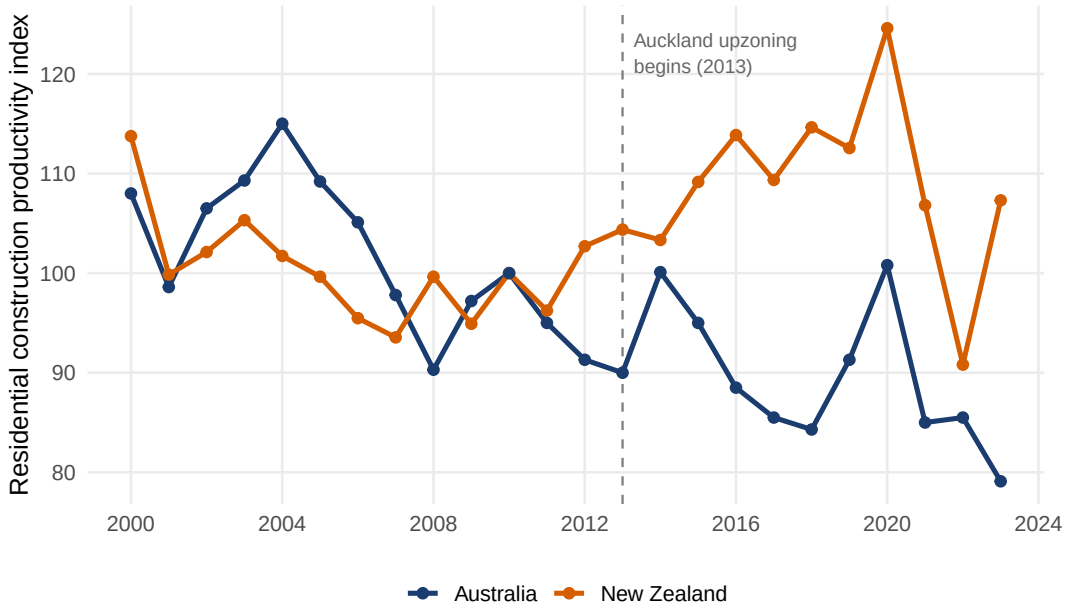


Source: Maltman (2025). Real output per hour worked in residential construction.

Source: Maltman (2025).

Figure 13: Residential construction productivity: Auckland versus the rest of New Zealand.

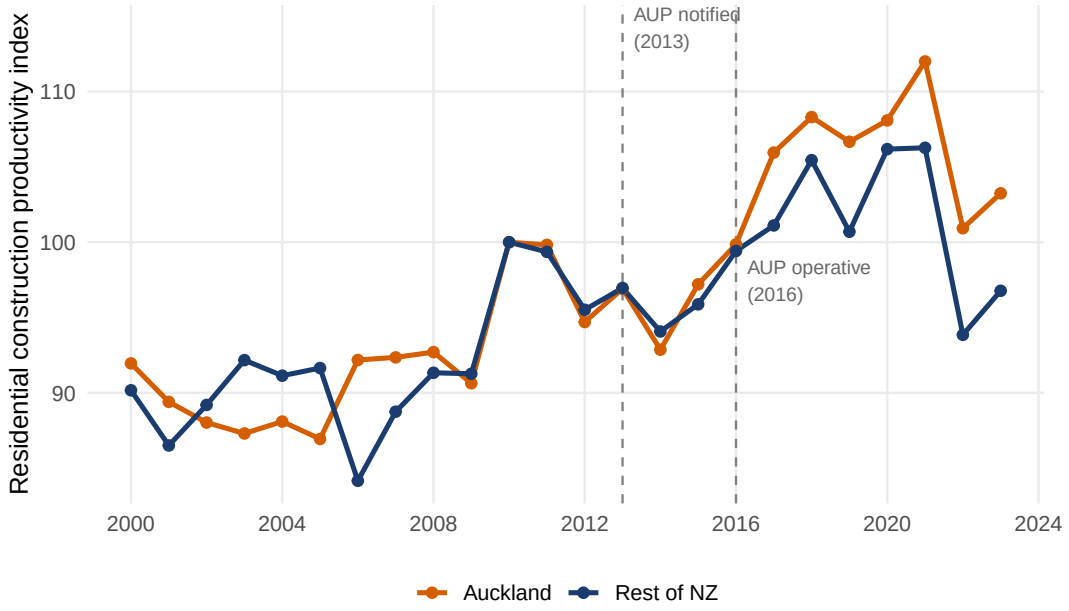


Source: Maltman (2025). Real output per hour worked in residential construction.

Source: Maltman (2025).

# 8. The limits of economics: this is partly a question of values

As Section 1.1 framed it, the relevant questions for public policy are: how much do land-use restrictions raise prices and reduce the quantity of housing built, and are these restrictions worth it? Economics can answer the first question with confidence: substantially, as the evidence reviewed above demonstrates.

The second question it cannot fully resolve. Defenders of land-use regulations argue they generate offsetting benefits, through the amenity value of low-density neighbourhoods, the management of congestion, heritage, or local externalities. There are genuine externalities that warrant regulation, and no serious reform proposal advocates eliminating land-use controls altogether. There is also a substantial political-economy literature documenting why incumbent residents resist new housing and why consultation processes systematically over-represent their preferences relative to those of prospective residents (Fischel 2001; Hankinson 2018). These preferences are real, and economics does not offer a basis for dismissing them.

To be specific about this, we can measure whether claimed externalities actually occur. For instance, we find in Campbelltown that on-street parking increased after the lot-size reforms (Behrens et al. 2026). Research can tell how much parking congestion a reform creates. What it cannot tell is how much weight that congestion should carry against more and cheaper housing. That is a question of values.

But, the burden of proof should shift. Given how large and how robustly established the costs of these regulations now are, the onus should sit with those who wish to retain a restriction to demonstrate the specific externality it addresses and that its benefit exceeds its cost.

## References

Anagol, Santosh, Fernando V. Ferreira, and Jonah M. Rexer (2025). "Estimating the Economic Value of Zoning Reform". In: *American Economic Journal: Economic Policy*.

Asquith, Brian J., Evan Mast, and Davin Reed (2023). "Local Effects of Large New Apartment Buildings in Low-Income Areas". In: *Review of Economics and Statistics* 105.2, pp. 359–375.

Behrens, Dominic, Theo Gibbons, and Matthew Maltman (2026). "The Effects of Minimum Lot Size Reform in Campbelltown, Adelaide". Forthcoming.

Büchler, Simon and Elena Lutz (2024). "Making Housing Affordable? The Local Effects of Relaxing Land-Use Regulation". In: *Journal of Urban Economics* 143, p. 103689.

Centre for International Economics (2013). *Reform of the NSW Planning System*. Tech. rep. Sydney: Centre for International Economics.

Coates, Brendan, Joey Moloney, and Matthew Bowes (2025). *More Homes, Better Cities: Letting More People Live Where They Want*. Tech. rep. Melbourne: Grattan Institute.

Cooper, Geoff, Ryan Greenaway-McGrevy, and James Allan Jones (2026). "Measuring the Cost of Land Inputs to Housing Construction". In: *New Zealand Economic Papers* 60.1, pp. 1–23.

Deloitte Access Economics (2016). *A Federal Incentives Model for Housing Supply*. Tech. rep. Prepared for the Property Council of Australia.

Donovan, Stuart and Matthew Maltman (2025). "Dispelling myths: Reviewing the evidence on zoning reforms in Auckland". In: *Land Use Policy* 151, p. 107498.

Fischel, William A. (2001). *The Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies*. Cambridge, MA: Harvard University Press.

Gabriel, Stuart A. and Edward Kung (2023). *Development Approval Timelines, Approval Uncertainty, and New Housing Supply: Evidence from Los Angeles*. Tech. rep. University of California, Los Angeles.

Greenaway-McGrevy, Ryan (2025). "Evaluating the Long-Run Effects of Zoning Reform on Urban Development". In: *Regional Science and Urban Economics* 111, p. 104062.

Greenaway-McGrevy, Ryan (2026). "Can Zoning Reform Increase Housing Construction? Evidence from Auckland". In: *Economic Modelling* 160, p. 107592.

Greenaway-McGrevy, Ryan and Peter C. B. Phillips (2023). "The Impact of Upzoning on Housing Construction in Auckland". In: *Journal of Urban Economics* 136, p. 103555.

Greenaway-McGrevy, Ryan and Yun So (2025). "Can Zoning Reform Reduce Housing Costs? Evidence from Rents in Auckland". In: Auckland EPC Working Paper 016.

Hankinson, Michael (2018). "When Do Renters Behave Like Homeowners? High Rent, Price Anxiety, and NIMBYism". In: *American Political Science Review* 112.3, pp. 473–493.

Housing Industry Association (Aug. 2025). *Planning Approvals: How Much and How Long?* HIA Economics (survey of 92 builders). National-average planning approval cost of $14,100 per dwelling.

Jenner, Keaton and Peter Tulip (2020). *The Apartment Shortage*. Research Discussion Paper 2020-04. Reserve Bank of Australia.

Kendall, Ross and Peter Tulip (2018). *The Effect of Zoning on Housing Prices*. Research Discussion Paper 2018-03. Reserve Bank of Australia.

Li, Xiaodi (2022). "Do New Housing Units in Your Backyard Raise Your Rents?" In: *Journal of Economic Geography* 22.6, pp. 1309–1352.

Local Government NSW (Oct. 2024). *Councils Already Playing Their Part in Solving the Housing Crisis*. LGNSW Media Release.

Maltman, Matthew (2025). "Zoning Reform and Construction Sector Productivity". Working paper.

Maltman, Matthew and Ryan Greenaway-McGrevy (2025). "Going it alone: The impact of upzoning on housing construction in Lower Hutt". In: *Journal of Housing Economics*. Also Auckland EPC Working Paper 018.

Mense, Andreas (2025). "The Impact of New Housing Supply on the Distribution of Rents". In: *Journal of Political Economy Macroeconomics* 3.1, pp. 1–42.

Murray, Cameron K. and Mark Limb (2023). "We Zoned for Density and Got Higher House Prices: Supply and Price Effects of Upzoning over 20 Years". In: *Urban Policy and Research* 41.2, pp. 129–147.

NSW Department of Planning, Housing and Infrastructure (2024). *Council League Table*. NSW Planning Portal.

Pennington, Kate (2021). "Does Building New Housing Cause Displacement? The Supply and Demand Effects of Construction in San Francisco". Working paper.

Productivity Commission (2022). *In Need of Repair: The National Housing and Homelessness Agreement*. Study Report. Canberra: Australian Government Productivity Commission.

Productivity Commission (Dec. 2025). *Creating a More Dynamic and Resilient Economy*. Inquiry Report 109. Canberra: Australian Government Productivity Commission.

Rollet, Vincent (2025). "Zoning and the Dynamics of Urban Redevelopment". Job Market Paper, MIT.

Saunders, Trent and Peter Tulip (2019). *A Model of the Australian Housing Market*. Research Discussion Paper 2019-01. Reserve Bank of Australia.

Soltas, Evan and Jonathan Gruber (2026). "How Costly Is Permitting in Housing Development?" Working paper.

Tso, Matthew (2019). *Lower Hutt Ratepayers Gift $29m to Developers Through Waived Fees*. Stuff.

Tso, Matthew (June 2021). *Developers to Pay for Bigger Slice of Infrastructure Costs in Hutt City*. Stuff.

Tulip, Peter (2022). *The Need for Reform of Land Use Policy: Submission to the Productivity Commission*. Centre for Independent Studies.

Wesselbaum, Dennis and Patricia Dela Cruz (Feb. 2024). *Land-Use Restrictions and Housing Affordability*. NZAE Member Survey #9, Asymmetric Information.