

Second Submission to the Productivity Commission Inquiry into Housing Supply Regulation

Matthew Maltman

Theo Gibbons

September 2026

Summary

- It is well established that land use regulations reduce housing construction and raise prices. We know much less about which regulations bind, where, and by how much.
- The objective of land use policy should be to maximise net benefits to the Australian community. And yet, regulations have accumulated without a clear rationale grounded in market failure or assessment of trade-offs. The available evidence supports ambitious reform, particularly for medium-density housing, although more evidence is needed on both the costs and benefits of specific land use regulations.
- Feasibility modelling is currently our best tool for assessing reform options. It asks which projects a representative developer could profitably build under different regulatory settings. No model is perfect, but in our paper it predicted out-of-sample roughly which sites developers chose and approximately matched the aggregate response to reform. This is encouraging evidence that such models can inform reform design.
- Reforms need to permit sufficiently large changes in built form to be effective. Small increases in development capacity may not reach the part of the feasibility curve that makes new construction stack up.
- High construction costs and interest rates do not make zoning reform ineffective. The evidence is that substantial development would become feasible if high-demand areas were upzoned, even in a high-cost environment. In weaker markets, reform may enable construction to pick up quickly when conditions improve.
- Feasibility modelling holds output prices, construction costs, and land prices fixed. This is reasonable for local reforms over short horizons where additional development is unlikely to affect broader market conditions. State-wide or national reforms, and longer-run analysis, may need to account for changing macroeconomic conditions and general-equilibrium effects from persistently higher construction.
- To improve accountability for planning reform across levels of government, we prefer a combination of feasibility targets and direct reform in larger jurisdictions, with greater reliance on direct reform where modelling capacity is limited. Housing targets are usually ineffective, as they often just delay political action and are confounded by the business cycle.

About the authors

Matthew Maltman is a PhD student in economics at the University of Chicago. Theo Gibbons is a research economist at the e61 Institute. This submission is made in our personal capacities: the views expressed are our own and should not be attributed to any institution.

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”, documents the relationship between New Zealand’s zoning reforms and its construction productivity performance.
- Behrens, Gibbons and Maltman (2026), “What local land use reform does, and doesn’t, do”, studies Campbelltown’s (SA) minimum lot size reform and provides evidence on its effects on quantities, construction sector activity, externalities, and movement chains.
- Behrens, Gibbons and Maltman (2026), “Where do minimum lot sizes bind? Evidence from a policy experiment”, uses the same experiment with a feasibility model to think through where and how minimum lot sizes bind.

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

1 Introduction and scope

This submission follows our first (submission 41), which was lodged before the interim report. That submission was largely interested in explaining the quasi-experimental literature on zoning reform, and putting together a simple framework about how to think about regulations in this market. In summary, we argued that:

- land use restrictions lower the quantity of housing and raise its price;
- supply is a curve and quantity is an outcome. It is important not to confuse the two. For example, construction volumes alone say little about how restrictive planning settings are;
- knowing what a reform did requires a credible counterfactual, either from a model or a (quasi-)experiment. If you want to think through prices, you need to do so at a sufficient scale of spatial aggregation;
- regulatory “bans”, meaning rules that prohibit housing outright, matter much more than regulatory “burdens” such as approval times, fees and contributions. An implication of this is that approval rates and timelines contain very little information about constraints on supply;
- claims that supply has little effect on prices usually rest on flawed designs or measurement at the wrong spatial scale;
- given the large evidence base on how land use regulations affect the housing market, the burden of proof should sit with those who want to retain a restriction.

We were pleased to see much of this framing taken up in the interim report, including the distinction between bans and burdens.

This submission focuses on a more granular question: how can we think through which regulations bind, where, and how should trade-offs be managed? It draws heavily on our paper, *Where do minimum lot sizes bind? Evidence from a policy experiment* (Behrens et al., 2026b). This paper contains a feasibility model and is currently under review at an academic journal. It is a subset of our broader working paper, *What local land use reform does, and doesn't, do*, which contains a wider range of analysis, including on housing movement chains and externalities. We also draw on our experience writing papers across a variety of quasi-experimental settings and on early-stage work underway in other contexts.

The interim report says the Commission will “compare commercially feasible and observed housing density at the parcel level to identify where development is theoretically feasible in the absence of controls but not currently occurring”, and will “model the impact of existing land-use controls on potential housing supply, and how changing controls would affect zoned capacity and flow through into a supply response” (pp. 15, 115–116). From this, it sounds like the Commission is thinking through a feasibility model for housing. Much of this submission is therefore written based on our understanding of how such a model should be built and interpreted.

2 Does it matter which regulations bind?

The microeconomics of housing development is understudied because of historical constraints on data. To understand a market you need the “holy trinity” of prices, quantities and costs.

Without all three, you cannot fully determine what is driving what within a market.¹

Prices have historically not been publicly available at a granular level. Quantities were not available at a unit-record level until state governments began to standardise their records, and technology made it much easier to scrape and machine-read large numbers of PDFs. Costs are still hard to observe, particularly differences between firms, although there are growing ways to approximate them.

So although it has long been observed that zoning restrictions lower the quantity of housing built and therefore raise its price, we have not had granular information on which restrictions bind, where, and how.

In some respects this is not a problem. A simple argument is sufficient for directionally correct policy: cutting restrictions should² increase housing quantities and therefore lower prices relative to the counterfactual. If increasing housing supply were the only policy objective, a reduction in any arbitrary planning restriction would be directionally positive. We can take this logic to its extreme, and argue that removing all zoning and land use restrictions will maximise construction, although it is unclear by how much and where.

On that view it is unclear what more granular market analysis adds to the debate. If cutting any restriction increases quantity, why does it matter how much a particular restriction binds?

One reason is that there are trade-offs associated with increased housing construction and other public policy objectives. A simple example is to imagine a government policy that paid a construction subsidy of \$100,000 for every new dwelling. That would unambiguously increase housing quantities. We would build more. But would it be good policy? It would impose a large cost on public finances and redistribute resources poorly. Nor would it necessarily solve the fundamental spatial problem of a lack of construction in places where people want to live.

This is a mistake that often appears in debates about housing and housing supply. Commentators put moral weight on increased supply being “good” or “bad”. We think, as we articulated in our first submission, that an Australia that builds more housing would be a better place to live as it would sustainably lower housing costs, increase consumer choice, and allow for consumption and production agglomeration effects. Reducing land use regulations would get us all of these benefits. But how we get there matters. These benefits need to be weighed against costs:

1. Local externalities. More housing may result in costs to existing residents, particularly in relation to congestion. That an increase in population will put pressure on local infrastructure is almost mechanical. That is, put more people in an area and, holding everything else fixed, there will be more cars parked on the street, more use of sewerage and other infrastructure, and more people at the local park.

But there are also effects in the other direction. More people can create positive externalities. The literature on the relationship between density and production externalities, carbon emis-

¹An example of why this distinction matters comes from the landbanking literature. An observation sometimes made in greenfield markets is that, following an increase in market-level construction at time t , some developers may reduce or delay their own development at $t + 1$. One interpretation of this is strategic: developers withhold supply in order to avoid depressing prices, consistent with market power or landbanking. But an alternative explanation is that a construction boom may push developers up a steep short-run cost curve as labour, materials, or financing become scarce. Firms may then optimally postpone projects until these input costs ease. Without information on developers’ costs, or some other source of exogenous variation that separates these mechanisms, we cannot conclude either way.

²In some settings a reform may have no detectable effect, or a very small one, but it would be odd for allowing more housing to reduce the amount built. Plausibly there are some general equilibrium conditions that could achieve this (for example, removing some building regulations could produce high-profile building defects that collapse confidence in the sector).

sions, and infrastructure costs is large and well documented, and we will not repeat it here. But often understated are the externalities related to consumption: more people bring more customers, and so allow for a wider range of businesses and local services. Our broader Campbelltown work is a first attempt to put numbers on both sides of this ledger, including estimating the impact of zoning reform on business formation and a range of potential local costs.

Weighing these two effects against each other, the literature that estimates how new housing affects local amenities is mixed-to-positive. New construction appears to improve local consumption amenities, with lower grocery prices, greater product variety, more restaurants, and more business turnover (Borraz et al., 2024; Li, 2022; Pennington, 2025; Blanco and Neri, 2026; Behrens et al., 2026a). There is little evidence that it raises overall crime, and some evidence that it lowers it (Freedman and Owens, 2011; Zhuo et al., 2025; González-Pampillón, 2022; Blanco and Neri, 2025; Behrens et al., 2026a), nor of adverse effects on schools, traffic or local government finances (Neri, 2024; Odayappan, 2026; Behrens et al., 2026a). The clearest costs are hyperlocal nuisances: increased parking citations, illegal dumping and complaints (Tanri-sever, 2025). Much of this evidence comes from subsidised housing or public housing rather than private medium-density infill, and should be read with that in mind.

Given that effects can be both positive and negative, more empirical work is needed to estimate their magnitudes so policymakers can weigh the trade-offs. A reason this is important is that it is unclear how much of the reported negative effect is perceived rather than real. For example, in our Campbelltown paper, we noted that locals were greatly concerned about parking, traffic, trees, character and infrastructure. These complaints generated opposition that eventually contributed to the reforms being reversed. It was not clear to us that these negative effects were sizeable, although hyperlocal effects are hard to detect with our current data.

The hyperlocal nature of externalities is highlighted by the fact that most objectors to new construction tend to live very close to the development in question (Figure 1). In addition, these local objectors were much more likely to raise privacy, overlooking, overshadowing and setback concerns than those further away, who were more concerned about general effects like traffic or safety.




Figure 1: Who objects, and what they object to in Campbelltown

Whereas the costs of housing construction are concentrated and salient, the benefits are broad and diffuse across the housing market. This means that even small negative effects, regardless

of whether they are perceived or genuine, can have a large effect on the social licence for reform. These negative effects are also often self-reinforcing: new housing is opposed because the infrastructure will be unable to cope, but when government proposes more infrastructure, that is opposed because either it is not needed, or due to concerns it will bring in more people.

Although it may seem obvious that land use restrictions exist to manage externalities, they are often not motivated or framed that way. Why, for example, was Campbelltown’s minimum lot size 350 square metres before the reform? One could say that the rule was designed in part to manage congestion or on-street parking through limiting construction and density. But why 350? Why not 300, or 450? The same question can be asked of planning rules all over the country. Why is a frontage minimum five metres rather than seven? Why is a height limit two storeys rather than three? Why is a particular setback three metres rather than four?

These parameters are almost never motivated by an estimated relationship between the regulation and the external cost. There is no established elasticity telling governments how a further 50-square-metre cut in minimum lot sizes would change street parking, traffic, sewerage use or pressure on local infrastructure. Very little empirical work has tried to answer those questions. Instead, governments tend to choose a broad kind of urban outcome, set the regulations somewhere in that area, and do so without a clear understanding of which rules bind, how much housing they prevent, or what their effects on local communities actually are. Much more empirical work is needed to put numbers on both sides of the trade-off.

2. Internalities. There are potential internalities associated with people moving into new construction, in particular related to apartment or lot sizes. The argument is usually motivated by normative ideas that certain types of housing are not good for individuals to live in.³ Generally speaking, individuals can trade off the conditions of their dwelling against its price and location, and therefore greater diversity in the housing stock is good for consumer choice.

There are, of course, regulations aimed at protecting consumers from harmful consumption, usually as a result of asymmetric information. A buyer or tenant may not know that a dwelling is poorly insulated against extremes of heat and cold, that its electrical wiring is unsafe, or that there is some other hidden defect. Building regulations that ensure apartments are safe, insulated, ventilated and otherwise fit to live in are, in principle, well-justified policy. The specifics of these regulations will differ: some will be net positive and others net negative, but they at least in principle address a recognisable market failure.

Land use regulations are different. People can observe whether an apartment is small, whether a house has a backyard, and how much private space a dwelling offers. They can trade those things off against price and location. These directly observable features should be treated differently from hidden safety or building-quality risks. This is the distinction the interim report draws in Principle 2 and Table 1.2, and we endorse it (information request 1.2).

Even if one accepts the premise that some forms of housing are undesirable for their occupants, regulating the characteristics of newly constructed dwellings is generally a very blunt way of improving the housing people actually consume. Housing markets are connected. Prohibiting small apartments does not ensure that households consume more space. Some may instead share larger apartments with more people, or remain in overcrowded existing dwellings. In addition, there may be other social harms, such as these households having to live farther from their preferred location, or pay more for the same amount of housing.

More fundamentally, the majority of the housing stock in any given year is existing rather than newly constructed. Sufficiently costly requirements on new construction can reduce the rate at which new housing is built, so slow the replacement and renovation of the existing stock.

³Variants of this argument are that it is “un-Australian” not to have a backyard, or that people “need space”.

Households may then remain for longer in older, lower-quality dwellings.

So, if the concern is a genuine hidden risk to consumers, directly regulating that risk may be appropriate. If the concern is instead that consumers are choosing dwellings that policymakers regard as too small, too dense or otherwise undesirable, restricting the supply of those dwellings is a weak and blunt response.

3. Policy resources and political licence. In theory, land use reform can be as simple as changing one number in legislation: a 350-square-metre minimum lot size becomes 250, or a three-storey height limit becomes six. In practice, however, even these reforms require public-servant time, policy analysis, community consultation, public attention, and votes. The direct fiscal cost may be small, but the policymaking process still has an opportunity cost. Governments can do *anything* but not *everything*.

One reason we prefer ‘ban’ reforms (e.g. land use and zoning reform) to ‘burden’ reforms (e.g. reforms to approval processes) is that the latter often carry direct fiscal costs, which means that their effects have to be sufficiently large to justify them. Waiving fees or reducing development charges requires a budget or shifts infrastructure costs elsewhere. Speeding up approvals may require more staff, new systems, or additional administrative capacity. Lower Hutt’s 2012 fee waiver, for example, generated little construction but put a large hole in the Council’s budget (Maltman and Greenaway-McGrevy, 2025; Tso, 2019).

A further cost is political licence. If reforms are poorly targeted—for example, a cut to non-binding regulations—the resulting lack of a supply response may erode public and academic trust in land use reform. Institutions may then look to other, less effective reforms.

These are all reasons to target reforms carefully. Governments should concentrate scarce resources and political capital on the controls that actually bind, rather than changing planning parameters indiscriminately and hoping that some of the changes produce housing.

2.1 Getting the reform ballpark right

The upshot is that zoning reform requires some assessment of trade-offs. There is some optimal policy: some mix of lot sizes, height limits, frontage rules, setbacks and other controls that manages the concerns above as well as it reasonably can while still allowing more housing. That is to say, optimal policy would be, either through a literal or implicit cost-benefit analysis, to maximise net benefits to the Australian community.

We are not necessarily implying that every change to land use policy should require a quantitative cost-benefit analysis, including for abstract concepts such as how much a community values the aesthetics of buildings. We are saying that optimal policy should justify regulations by what externality they manage, provide evidence that they are proportionally aligned with the costs they impose, and aim to maximise the welfare of the entire Australian community, not just local residents.

The problem with achieving this under the status quo is twofold. First, we do not currently have sufficient evidence on how specific land use regulations relate to local externalities in order to weigh up everything accurately. Second, a shift to this policy mix may generate a large backlash, as Campbelltown demonstrates.

Given this, a reasonable proxy is for governments to approximate externalities by retaining particular regulations to mitigate the harms communities care most about while demonstrating that they still have shifted out the supply curve materially.

For example, if a community really dislikes on-street parking, a government could raise floor

area ratios or cut lot sizes dramatically but keep parking minimums in place. This may be a politically durable transition package if the parameters are set at the right levels.⁴ If communities really value the aesthetics of buildings, they could retain strong architectural regulations while allowing buildings of all different shapes and sizes. If a community dislikes small lots, it could retain larger lot size minimums but relax height limits. If it dislikes overshadowing, it could keep height limits but reduce lot sizes, and so on. Provided the net effect of the package is to allow substantially more overall construction—directly and empirically measured, through feasibility modelling or otherwise—it can be designed so that quantities rise and prices fall while mitigating the harms the community minds most.

To conclude, it is striking how little attention historically has been given to land-use regulation relative to other areas of public policy. Over the past 25 years, economists have lamented the lack of regulatory and economic reform. And yet when we talk about economic reform—through media attention or substantial analysis—we mention only areas like tax and transfer policy, industrial relations, government procurement, and business regulation.

A more mature approach to land-use policy requires treating it like any other major area of public policy: identifying the market failures being addressed, measuring the effects of existing rules, and comparing alternative policy designs. The evidence base remains much thinner than it should be given its importance.

3 Feasibility modelling as a way of measuring market conditions

3.1 What feasibility modelling does

The explosion of data access and availability has allowed much more detailed analysis of the microeconomics of housing construction because data on the “holy trinity” of prices, quantities and costs has become available to some researchers. The best example is Rollet (2025), which builds a parcel-level dataset of all of New York, and is able to work through both the demand and supply sides of what gets built, where, and how this plays out dynamically over the long run under different regulatory conditions. There are still gaps in this literature, discussed below, but broadly speaking it explains much of what we need to understand.

In Australia the long-standing form of analysis has been the “zoning tax” (Kendall and Tulip, 2018; Jenner and Tulip, 2020; Tulip, 2022; Coates et al., 2025). More recently, feasibility analysis has taken off: by Grattan (Coates et al., 2025), by the Centre for International Economics (Rajaratnam et al., 2024), and by us in our Campbelltown paper. We believe feasibility analysis to be the more powerful tool of the two, but they are conceptually similar and both useful. Feasibility analysis lets us horse-race which regulatory changes unlock development, and where that development occurs, under different assumptions about prices and costs. Like any model, it simplifies reality to construct counterfactuals, but forces us to be explicit about what regulations, prices, and costs actually determine what gets built.

The conceptual basis of feasibility modelling is straightforward. Take a representative developer facing a given cost structure, a set of possible development opportunities, and the output prices those projects could receive. Land use regulations restrict the set of projects the developer may choose, including projects that would be profitable.

The formal mathematical setup we use can be found in Behrens et al. (2026b). In that model we estimate the feasibility of every site in Campbelltown by:

⁴Although the permanent policy test should remain whether parking minimums have quantified net benefits and are more effective than a more direct instrument such as residential parking permits or pricing on-street spaces.

1. Imputing the fixed input costs for each site, beginning with the price a developer would pay for the parent lot and adding financing, taxes, and other fixed costs.
2. Estimating the maximum number of medium-density dwellings that could legally be built under each regulatory setting. For minimum lot sizes, this means dividing developable site area by the minimum and rounding down, since one cannot build 0.2 of a dwelling.
3. Once the legal yield is known, multiplying marginal per-dwelling costs (construction costs) by the number of dwellings, and adding fixed costs to obtain total development costs.
4. Imputing revenue from the sale of those dwellings at their resulting lot sizes, including the new-dwelling premium, using observed sales in the area.
5. Comparing revenue with cost and classifying a site as feasible when the developer earns at least a 20 per cent margin on development cost.⁵

Our model was bespoke, built to answer a simple question in a single council area: where do minimum lot sizes bind? We made many simplifying choices to get at that question. Such choices get harder as the question gets larger. In practice, feasibility modelling becomes more difficult as regulations are compounded and simplifying assumptions are relaxed. Some examples are:

How should construction costs be treated? We estimate construction costs as fixed per dwelling because we study one form of medium-density housing, although we test alternative assumptions. A broader application will need to allow costs to vary by a wide range of dwelling types, for example, structure type and floor area. Modelling results are highly sensitive to choices about construction costs, so these details matter.

Should construction costs be endogenous? Construction costs can differ structurally across regions because of labour markets and logistics. But measured differences can also be produced by regulation: construction in a constrained, high-value market may be expensive because developers respond to scarce sites by building a higher-end product. So if a modeller wants to estimate feasibility differences across a broad region, they would probably want to retain the first effect, but not the second. Therefore, they need to decide which regional differences are likely to persist after reform and which converge as the mix of projects changes.

How important this issue is depends on how sale prices are modelled. If sale prices are estimated with less regional specificity, or can be adequately adjusted for differences in dwelling quality, there is more scope to allow construction costs to vary regionally. Alternatively, using regional averages for both sale prices and construction costs corrects for this at the cost of local precision. A problematic case is using highly localised sale prices which cannot be fully quality-adjusted, but construction costs at a regional level. In that case, differences in project quality can make high-value locations appear more feasible than they really are.

How many regulations should the model include? We were broadly looking at one regulation. As more regulations are added, the number of interactions grows exponentially, each requiring more assumptions. If lot sizes, height limits, floor area ratios and setbacks were all allowed to vary, the model would need to consider a variety of scenarios where these changes interact. There may be little data on the costs or prices of some types of development under these different policy mixes, or even on whether they are architecturally feasible. For instance,

⁵This is close to the mean 19.3 per cent hurdle reported by Moorhead et al. (2024); we test lower and higher hurdles in sensitivity analysis.

it would be hard to say what a fifteen-storey apartment building would sell for in suburban parts of Sydney because the price imputation would be weak.

Should the model be partial or general equilibrium? Our model is partial equilibrium. In a small, local setting this is a reasonable approximation, particularly because we were able to benchmark this against actual empirical evidence on what happened. Large-scale feasibility modelling must think harder about prices and costs moving endogenously, as Section 4 discusses.

How much heterogeneity? We assumed one firm with a constant cost structure and one price schedule. Richer models could distinguish high- and low-cost firms or high- and low-end products. The value of that additional complexity will differ depending on the question being answered.

3.2 How well feasibility modelling performs

In our experience, feasibility modelling performs quite well both cross-sectionally and dynamically.

Cross-sectional performance. Are parcels the model ranks as most profitable subsequently developed more often than those it ranks as least profitable? Figure 2 shows a steep gradient across projected-margin quintiles in Campbelltown. Among transacted lots, around 7 per cent in the least profitable quintile were developed over the reform period, compared with around one quarter in the most profitable quintile. The same gradient appears when we extend to all lots—not just those transacted.




Figure 2: Share of Campbelltown lots developed by lot-characteristic quintile

There is still a large residual: some highly profitable sites are not developed, while some apparently unprofitable sites are. This is consistent with there being heterogeneity the model

cannot capture. A representative developer misses genuine differences in costs, products, and sale prices across firms. In addition, there are market shocks developers may not anticipate or may miscalculate.

These are real limitations. The relevant comparison, however, is not with a perfect model, but with no model.⁶ The objective of these models is to explain the market well on average, not to explain every site exactly.

Dynamic performance. Can the model predict how construction changes after reform? Our Campbelltown paper contains, to our knowledge, the first feasibility model tested against a quasi-experiment: minimum lot sizes were relaxed and then partially reversed, allowing us to compare model-implied construction with reduced-form estimates. Figure 3 shows this.

This was not a full out-of-sample exercise from our perspective, because we knew the quasi-experimental result and built the model partly to understand it. However, the model uses a standard feasibility structure and was not calibrated to aggregate construction flows.⁷ Regardless, we think this is encouraging evidence that feasibility modelling can approximate the magnitude of construction responses from reform.




Figure 3: Model-implied and reduced-form additional dwelling flows in Campbelltown

Most critiques of feasibility modelling that we have heard cherry-pick a particular cost or regulation that was not perfectly measured. Our conclusion from working on the back end of these models is that these things usually do not matter that much, unless they relate to construction costs themselves or sale prices. You could dial up or down things like demolition costs, infrastructure or water connection costs, etc., and not see much movement in the aggregate conclusions or predictions of the model.

⁶Sometimes, of course, no model is better than an actively bad one.

⁷The conversion from feasible stock to annual flow does, however, use observed redevelopment rates. The results do not change materially when we relax that assumption.

3.3 Regulatory interactions

Modelling also allows a researcher to identify how regulations interact. For example, Campbelltown permitted row dwellings with a 5-metre frontage, while comparison-council standards ranged from 8 to 15 metres. This matters because frontage and setback rules become geometrically binding on small lots and cap the gain from reducing lot-size minimums. The modelled gain from cutting a uniform minimum lot size from 350 to 150 square metres varies substantially across Adelaide councils (Figure 4). Frontage requirements are the largest or second-largest source of the gap in every case.

A sufficiently detailed model lets governments see which control becomes binding under each package of reforms. If lot sizes and frontages are relaxed, for example, site coverage, setbacks or parking may become the next constraint.




Figure 4: Feasible capacity and the peak-response minimum lot size across Adelaide councils

In addition, modelling suggests the effects of relaxing development regulations can be highly nonlinear. In Campbelltown, we find that a 100-square-metre reduction in the minimum lot size from 350 to 250 square metres generates far fewer feasible development opportunities than an equally sized reduction from 250 to 150 square metres. Feasibility increases particularly rapidly once minimum lot sizes fall below roughly 250 square metres, with the marginal gain in

feasibility peaking at around 175 square metres.

This means that a government may reduce minimum lot sizes, frontage requirements, or height limits, yet generate little additional construction if the reform does not relax the constraint that is actually binding, or does not relax it far enough. Effective reform therefore requires policymakers first to identify which regulation constrains development at the margin and then to relax that constraint sufficiently to move developments into the region where additional capacity becomes economically feasible.⁸

This distinction matters for how unsuccessful reforms are interpreted. A modest relaxation that produces little new construction need not imply that planning regulations are unimportant or that further reform would also be ineffective. It may instead indicate that the reform did not reach the relevant threshold. Governments can therefore undertake reforms that appear meaningful while creating relatively few new feasible development opportunities.

3.4 More data will improve feasibility modelling

Despite improvements in data access and availability, good analysis of the housing market is held back by a lack of linkable unit-record data on transactions, approvals, and completions, together with parcel geometry and construction costs. State and territory governments generate much of this administrative information, but public access is uneven. New South Wales sets the gold standard. Valuation NSW provides free bulk sale-level files from 1990 onward, while the NSW Planning Portal provides a daily public feed of development applications lodged through the portal since 2019 (Valuation NSW, 2026; NSW Department of Planning, Housing and Infrastructure, 2026a).

South Australia is by far the worst jurisdiction we have seen. In 2017, the state appointed Land Services SA as the exclusive private provider of specified land-titling and valuation services for 40 years and granted it rights to commercialise state land-information assets, while retaining ownership of the underlying data (Auditor-General’s Department of South Australia, 2018). Unit-record property-sales data are consequently supplied through fee-paying licences rather than a free public bulk release (Land Services SA, 2026). Putting a core input to housing-market research behind a commercial paywall makes independent analysis unnecessarily difficult. Anecdotally, even SA government departments struggle to get access to basic property data.

Completion data are also poor: a complete national unit-record series is effectively unavailable. New South Wales comes closest by maintaining a feed of post-consent certificates, including occupation certificates, although coverage is incomplete before July 2021 (NSW Department of Planning, Housing and Infrastructure, 2026b). The ABS’s public small-area Building Approvals release stops at SA2 and local-government-area level (Australian Bureau of Statistics, 2026). The ABS should publish much more granular approvals data, at least at Mesh Block or SA1 level and preferably as geocoded unit records. Privacy risks seem manageable as state systems already publish planning applications at individual locations.

The Commission should recommend that State, Territory and Commonwealth governments make unit-record data available to researchers. This would be an extremely important piece of research infrastructure.

⁸Of course, as discussed above, one approach is just to reduce *everything* substantially.

4 General equilibrium and macroeconomic dynamics

Feasibility models provide microeconomic explanations for what gets developed, where, and why. Bridging the gap from those parcel-level choices to market-wide prices, quantities, and land values remains an open problem in the literature, particularly once business-cycle dynamics are allowed to matter.

4.1 Partial versus general equilibrium

Feasibility modelling is usually partial equilibrium: prices and costs are held constant. Sensitivity analysis can vary them, as we do in our paper, but does not make them move endogenously with demand and supply.

In some contexts (e.g. for us in Campbelltown) this is a second-order concern. If a reform is local, the price and cost responses should be limited. But imagine instead reform across all of New South Wales, or all of Australia. The expected construction response should be large enough for the partial-equilibrium assumption to break down, through three main channels.

1. **Construction inputs.** A national expansion would increase demand for labour and materials, raising costs and narrowing the set of feasible projects. This is intuitively true, but we think these concerns are somewhat overstated for a few reasons:
 - Many materials are traded in global markets in which Australia is a small buyer. Their prices likely do not rise strongly with Australian demand, although this is an empirical question.
 - Labour markets are good at reallocating workers after structural shocks. Employment moved into mining during the mining boom and has more recently moved into care. Adjustment takes time and wages can rise, but usually these pressures settle in the medium run. In other words, we should not just assume the current construction workforce would be fixed if we wanted to build substantially more homes.
 - Empirical evidence supports the construction sector’s ability to scale up with zoning reform. New Zealand is the closest available empirical experience. Its construction workforce almost doubled between 2013 and 2023. Wage pressure did emerge, but sector productivity also rose over much of the decade (New Zealand Ministry of Business, Innovation and Employment, 2024; Maltman, 2025). While this is not Australian evidence, it shows that sector capacity is not fixed.
2. **Output prices.** Prices should fall as construction rises, which makes some modelled feasible projects infeasible. In Saunders and Tulip (2019), a permanent 1 per cent increase in the dwelling stock lowers real housing prices by around 2.5 per cent in the long run.
3. **Land prices.** Development rights may capitalise into parcel values. The policy-relevant question is whether that response is large enough to erase most of the construction gain. No inner-urban study we know of has shown that result. Section 4.2 explains this in detail.

The question is whether modelling is sensitive to these effects. The existing evidence suggests that it is not. In Campbelltown, output prices would have to fall by about 23 per cent, or parent-land prices rise by about 81 per cent, to return modelled feasibility to its pre-reform level. In New York, Rollet (2025) allows rents and floorspace prices to respond endogenously and estimates that removing density limits raises long-run floorspace by about 71 per cent.

In addition, the transition path from one equilibrium to another matters, not just the end point. Reform can generate an initial construction surge; as output prices soften and the regulatory




Figure 5: Price and land-cost changes required to reverse the Campbelltown feasibility gain

wedge closes, activity may settle at, or even below, its pre-reform rate. The area under the curve matters here, not the amount of housing built at a given time. In other words, this can represent success if cumulative housing is higher and prices are lower than they otherwise would have been.

Our model also illustrates how much the macroeconomic environment can change measured capacity without overturning the regulatory result. Figure 6 applies the post-pandemic increase in real construction costs and development finance rates across a range of counterfactual minimum lot sizes in Campbelltown. Higher construction costs and interest rates decrease feasibility, yet it remains well above the pre-upzoning counterfactual (350-square-metre minimum lot size).

Governments regularly tell us that developers say projects do not stack up at current construction costs and interest rates, and then sometimes infer that zoning reform therefore cannot work. The first clause is unsurprising, but the inference that therefore zoning reform will not work is illogical. Although construction costs and interest rates matter a great deal, a lot of development still stacks up if restrictions are loose enough in high-demand areas. Empirically, this is consistent with high rates of development in New Zealand over the recent business cycle.




Figure 6: Modelled feasible capacity in Campbelltown under construction-cost and finance-rate shocks

4.2 Land prices

Land prices deserve specific attention here. A challenge to feasibility modelling is that it is run on land prices before reform, so it does not account for the “change of use premium”, or the general equilibrium effect of upzoning on land prices.

The intuition is clearest at the urban fringe. Agricultural land can be much cheaper than adjacent residential land, due to the fact that residential land is more scarce and more productive. If we ran a model estimating the feasibility of residential development using the agricultural price of land, the concern is that we will overstate post-rezoning feasibility because we are not accounting for the rise in land prices after rezoning.

That mechanism is real. But it does not follow that land prices rise *enough* to extinguish construction. It also does not follow that the above example is representative of what happens in inner-urban, medium-density land markets. In fact, no study we know of has shown that result empirically.

Indeed, it is not even clear that upzoning will increase land prices at all in many cases. When we upzone suburban land, we see two forces racing against one another. First, upzoning allows for greater use of the land, increasing developer demand, resulting in its price being bid up. But second, a broad reform also expands the supply of potential sites, forcing landowners to compete to sell to developers, which pushes down prices. Which force dominates is an empirical question.

Consistent with this, the empirical evidence is mixed. In Houston, Shortell (2022) estimates positive capitalisation after reform, while Hamilton (2024) finds no positive effect, and negative estimates in some specifications. Greenaway-McGrevy et al. (2021) find that Auckland upzoning increased the value of undeveloped upzoned parcels *relative* to more-developed lots and non-upzoned properties. That result identifies a relative redevelopment premium; it is not the citywide general-equilibrium effect.⁹ In addition, total land prices are not the object of interest. Cooper et al. (2026), for example, find that land prices per square metre rose in Auckland’s most liberalised zones while land cost per square metre of permitted floorspace fell, because land could be used more efficiently.

In our Campbelltown study, land values rose only modestly after the reform. Across several specifications attempting to pin down the effect on house prices and land values, our estimates range from zero to 7.5 per cent (Table 1). The upper estimate compares prices developers paid for parent lots in Campbelltown with developer purchases in other councils. The assumption is that a developer buying a lot to redevelop it is paying for the land alone, since the existing structure will be demolished and is therefore worth nothing to them (or less than nothing, since demolition is costly). By contrast, our feasibility model implies that parent-land prices would have to rise by around 81 per cent to eliminate the gains from the reform (Figure 5).

We also test the change-of-use theory directly. We run the model under pre-reform market conditions and then under post-reform land prices (Table 2). Around 90 per cent of lots receive the same feasibility classification under the two price schedules. Modelled capacity is 25,658 dwellings at observed post-reform prices and 30,530 using the pre-reform price schedule. Higher land prices therefore remove some feasibility, but nowhere near enough to alter the substantive conclusion.

We also have forthcoming work on zoning changes in other contexts and have again been sur-

⁹To be specific, this is the “missing intercept problem” of applied microeconometrics. Causal studies identify the effect of a treated group relative to a control group (in general). A policy could cause land prices to fall across a whole metropolitan area, but more so in control regions. This would be estimated as a positive effect on land prices in the treated group. This is a useful finding, but it is not what we are interested in here: whether land prices rise enough to offset feasibility.

Table 1: Land and house prices in Campbelltown relative to control councils

Specification	Campbelltown $\times$ reform		Campbelltown $\times$ reversal	
	Estimate	SE	Estimate	SE
<i>Repeat sales</i>				
All pairs	0.010	(0.016)	−0.046	(0.010)
High subdivision probability	0.048	(0.021)	−0.066	(0.017)
<i>Level prices</i>				
All houses	−0.000	(0.010)	−0.019	(0.006)
Developer inputs (2+ lots)	0.075	(0.017)	−0.026	(0.015)
Hedonic, high probability	0.024	(0.009)	−0.004	(0.005)
<i>By development uplift</i>				
High uplift, hedonic	0.023	(0.007)	−0.037	(0.009)
Low uplift, hedonic	0.002	(0.009)	−0.025	(0.008)
High uplift, repeat sales	0.000	(0.018)	−0.033	(0.015)
Low uplift, repeat sales	0.000	(0.028)	−0.067	(0.016)

Notes: Pre-reform is the omitted period. All specifications include council and year fixed effects. Randomisation-inference p -values are 1.00 for all houses and 0.29 for developer inputs. Coefficients are log points. Source: Behrens et al. (2026b).

Table 2: Feasibility with observed post-reform prices and the pre-reform price schedule

Zone	Feasible dwellings		Lots classified identically
	Observed prices	Pre-reform prices	
400-metre transit buffer	23,908	28,087	85%
Policy area	1,248	1,921	95%
Rest of Campbelltown	502	522	98%
Campbelltown total	25,658	30,530	90%

Notes: The same lots are evaluated using observed post-reform prices and prices imputed from the pre-reform schedule. Source: Behrens et al. (2026b).

prised by how little land prices appear to move. Other researchers have broader work underway that may identify the citywide effect more directly, but we cannot speak to those results here.

To be clear, above we are talking about medium-density reform spread across a city. Highly concentrated apartment rezonings around transit can look different. Scarce designated sites could see a stronger change-of-use premium, particularly as such projects often require neighbouring parcels to be amalgamated to reach a viable site area. Neighbours coordinate rather than compete, which gives existing landowners more power relative to the developer. But it is important to keep in mind that there is a reporting bias here. People report and remember large land-price increases after rezoning and forget, or never hear about, cases where prices did not move. A null result is less interesting than a spike, in both media coverage and the academic literature.

The evidence is therefore mixed on whether zoning reform will push up land prices in general, and does not support the idea that land-price capitalisation extinguishes the construction gain from broad medium-density reform. Multiple applied microeconomic studies find large construction responses after zoning reform. By definition, land prices did not rise enough in those settings to extinguish the quantity response.

Finally, even if we assume full capitalisation, this does not imply that planning reform does not increase housing supply. Mechanically, land value capitalisation can only occur where a site's development potential increases. Consider a site in a desirable inner-city location currently zoned for low-density detached housing, that is rezoned for high-density apartments. If we assume full capitalisation, its land price will increase to the point that development is only marginally feasible. However, despite this land value having increased to "offset" feasibility, expected future development on this site is clearly higher. Previously, one dwelling was the most that could be provided; developers are now paying more because they can now build more. Developers will not pay a premium for an option they have no intention of exercising.

Regardless, these findings raise questions about how to model land markets that we do not yet know how to answer. Most models of land markets seem inconsistent with emerging empirical evidence.

The real-options literature models a landowner choosing the optimal time to develop or sell in response to expected rents, prices and costs (Clarke and Reed, 1988; Guthrie, 2010). This seems incomplete in established suburbs, where households sell in response to lifecycle shocks—for example, because they move for work, downsize, separate, or experience a death in the family. They want the best available price when that event occurs and are unlikely to want to wait for an optimal bid from a developer.

Consistent with this, in Campbelltown even in the most profitable quintile of sites that changed hands, only one in four was bought by a developer and then subdivided. The majority of buyers were owner-occupiers or investors who valued the existing dwelling and retained it. A developer in this land market would therefore largely be competing against non-developers, and a landowner would largely be selling to buyers who want to retain the structure.

This may partly explain the lack of capitalisation, but it raises another puzzle: if these sites are so profitable to redevelop, why are they not snapped up and built out? Weak explanations invoke a desire not to "flood the market". The more likely explanation is market structure and liquidity constraints.

Much infill development is bespoke and idiosyncratic. Sites arrive irregularly, differ in shape and support different products, which makes the work difficult to standardise and scale. Campbelltown's response was correspondingly unconcentrated. This is a market of many different actors building different projects, not a small number of large developers exercising market power.

If this model broadly holds, these developers may not want to expose themselves to balance sheet risk, or may have difficulty obtaining the financing to expand to many sites at once. They may hold only a few sites in case one fails, or lack the liquidity or capital that a bank would require. As a result, they may not be competing against one another at auction.

This affects how we think about developer contributions (information request 4.4). The interim report argues that contributions are eventually capitalised into lower land bids. That mechanism may describe greenfield markets, but it is less obvious for infill: our evidence is that the average buyer, including on highly feasible sites, is not a developer. Where a developer must bid against an occupier or investor, that bid comes out of their margin. This does not matter on an inframarginal project, but in a market with many marginally feasible sites, it could kill development. The same caution applies to rezoning-triggered windfall-gains taxes, which in our view meaningfully risk reducing housing supply, particularly in marginally feasible markets. Broad-based land taxation is more principled.

4.3 Macroeconomics: the business cycle and segmented markets

Feasibility models describe the microeconomics of development in a particular place. The literature has not yet built those parcel-level choices into a satisfactory account of the macroeconomy, including how feasibility moves over the business cycle.

This micro-to-macro problem matters for a few reasons. One is identifying which reforms work and which do not.

Consider a meaningful zoning reform in a marginal market when construction and finance costs are unusually high. A researcher wishes to understand the impacts of this reform, and so runs a difference-in-differences analysis and finds no construction response. They conclude either that regulations were not binding, or that land use reform does not work. Three to five years later, costs or interest rates fall and construction takes off—potentially in absolute terms, but at least relative to the no-reform counterfactual—because the greater zoned capacity enabled by the reform has now become feasible. The researcher missed this because the difference-in-differences relies upon a policy effect *at or near* the timing of the reform for identification.

In this example, the applied microeconomic design was not sufficient for policy analysis. It correctly concluded that the policy had no short-run effect, but this is probably not the object that policymakers are interested in. Governments likely care about long-run outcomes, which might vary across the business cycle.

The issue is sharpest in weak-demand regional markets. Output prices may be low relative to the national distribution yet housing can remain expensive relative to local incomes. Construction costs are set in national and international markets, and may be higher locally because of logistics. In favourable parts of the cycle, lower costs and interest rates may allow some projects to stack up; in unfavourable periods, they may not. Demand can continue to shift out while supply remains unresponsive, so prices rise without triggering construction until they clear a large cost hurdle.

In other words, if zoning restrictions are tight, supply is inelastic in the “bad” macroeconomic periods due to construction and financing costs, and in the “good” periods due to zoning regulations. We need to build an evidence base to identify *if* and *where* this is true.

Feasibility modelling in these markets should therefore ideally be dynamic, or find ways to consider different business-cycle scenarios. It should show which housing types become feasible under different cost and finance conditions, so governments can set rules that permit a strong construction response when conditions improve. In these markets, we may have to wait out unfavourable parts of the cycle, and then ensure we enter the favourable parts with large amounts

of zoned capacity to build quickly.

Macroeconomic conditions can influence construction, but construction also influences the macroeconomy. Construction is one of Australia’s most cyclical sectors and dwelling investment is an important channel of monetary policy (Reserve Bank of Australia, 2024). If zoning keeps projects close to the feasibility margin, small changes in rates and costs can produce large swings in activity. Reform may raise the floor, leaving projects highly feasible in good conditions and marginally feasible in bad ones rather than moving between marginal and impossible. The upshot of this is that zoning reform may stabilise the construction sector, at least over the next few decades while the “zoning tax” is eroded, at higher levels of construction. This would have broader implications for monetary and fiscal policy, and macroeconomic stabilisation, as it would reduce the cyclicity of one of Australia’s most cyclical sectors.

In summary, for local, short-run reform, a partial-equilibrium model remains a useful way to identify which controls bind. Any national estimate of benefits, however, must take an explicit position on how prices, land values, construction costs and productivity respond, and on how those responses unfold over the cycle.

5 Housing targets and intergovernmental relations

It is broadly accepted that Australia needs to build more housing. The submission thus far has focused on the specific, granular regulations that need to change to allow more housing to be built. A separate question is how the political economy and structure of government can create incentives for land-use policy to move in that direction. On this issue, there is considerably more disagreement.

State governments largely manage land-use policy through their statewide planning systems and oversight of local government. Half a decade ago, a common argument was that State governments already have incentives to get land-use policy “right” for fiscal and competitive-federalism reasons: for example, additional housing may generate stamp-duty revenue or attract residents from other states. On this view, there was little reason for the Commonwealth Government to become involved in the political process of housing construction, and little need for State governments to be pushed to reform their processes, as they should be close to optimal.

We have moved away from this view over the past few years, in large part because the status quo has not delivered sufficient housing construction and because reasonable arguments have been made for Commonwealth involvement in housing.¹⁰ It has also become clear that local governments can exercise substantial political influence, particularly at the state level.

Housing policy therefore requires some degree of political accountability and coordination across levels of government. There are several tools which aim to get governments to want to build more housing, discussed below.

5.1 Housing targets

The simplest version of this is “housing targets”. The NSW Government has used versions of this approach, and the Commonwealth Government’s New Homes Bonus provides another example. The logic is straightforward: each area is assigned an amount of housing it is expected

¹⁰Versions of these arguments include that some state governments can export the effects of their housing shortages, or that the Commonwealth bears political responsibility for housing irrespective of whether it controls the relevant policy levers. The latter has important intergovernmental implications: if state governments do not reform their planning systems, political pressure may fall on the Commonwealth to pursue less effective reforms within its own policy domain to improve housing affordability.

to build over some period, according to a specified methodology, and its performance is assessed against that target.

A major attraction of housing targets is that they preserve the principle of subsidiarity. Lower-tier jurisdictions retain discretion over which regulations to change and can concentrate development in the locations and typologies their constituents prefer. If a community wants substantial apartment development along major roads while limiting development within existing suburbs, it can pursue that approach. It can also choose the reverse.

That said, we generally do not think housing targets are a good way to manage the political-economy problems surrounding housing construction, for several reasons.

First, targets require some form of remedy or accountability when jurisdictions do not meet them. This is particularly important when targets are imposed by a higher-tier government on lower-tier jurisdictions that are opposed to new construction. Those jurisdictions may have little political incentive to meet their targets, particularly where their constituents are also opposed to development.

A common remedy is to remove or override some planning powers if a lower-tier jurisdiction fails to meet its target. In other words, there is a “stick” of losing control of your planning system altogether if you don’t reform it enough to generate construction. This can be complemented with a “carrot” of some kind of fiscal reward from meeting or exceeding your target.

We see this as delaying the underlying political problem: at some point, the higher-tier government must absorb the political costs associated with more development in these areas. Either the higher-tier government is willing to require planning reform now, or it will need to intervene later when the lower-tier jurisdiction fails to meet its target. If it is unwilling to do so for political reasons now, it is probably unwilling to do so later. At best, it simply delays reform for several years and may ultimately produce a watered-down response.

The second problem is that while governments can increase construction through planning reform relative to the counterfactual, they cannot always do so in absolute terms. A recession, an increase in construction costs, or higher interest rates can result in missed targets even where substantial good-faith effort has occurred. This is likely relevant to the Commonwealth Government’s current New Homes Bonus, where high interest rates and construction costs have made targets more difficult, though importantly not impossible, to achieve.

The interaction of these two problems makes housing targets difficult to enforce. Some jurisdictions may have little desire to meet them and ultimately only partly control the absolute amount of construction. It can therefore be ambiguous whether a jurisdiction that falls short has failed to reform adequately or has simply faced unfavourable market conditions.

For the above reasons, we think housing targets are dominated by direct reform of planning rules.

5.2 Feasibility targets

A newer approach seeks to improve on housing targets by instead setting feasibility targets. Instead of promising to build a given amount of housing, it requires a jurisdiction to maintain some specified amount of feasible zoned capacity. This shifts accountability from realised construction in the future towards the planning settings that determine feasibility today, forcing relevant policy questions to be resolved earlier.

Feasibility targets retain some of the subsidiarity benefit of housing targets. Lower-tier jurisdictions can choose which regulations to reform and can attempt to concentrate development in the locations and typologies they prefer within their communities.

They also allow for an independent assessor, such as a government commission or department, to apply a common feasibility model across an entire region. This can move disputes away from a vague assessment of whether an area is “doing enough” and towards quantifiable measures that can be assessed and compared consistently.

In this sense, we think feasibility targets dominate housing targets, as they retain many of their benefits while removing the ambiguity over whether a shortfall reflects inadequate reform or unfavourable market conditions. They do not, however, solve the broader political problem. If a lower-tier jurisdiction does not want to meet its target and resists the higher-tier government, the higher-tier government still requires some remedy with which to enforce compliance. The political problem therefore remains, although feasibility targets at least bring the problem forward in time.

Feasibility targets also raise several new, largely methodological issues. The first is determining how much feasibility is “enough”. This is an empirical question, and early attempts have been made to estimate the relationship between feasible capacity and construction, so it is potentially solvable.

The second issue is that feasibility varies over the business cycle and with broader macroeconomic conditions. Areas will naturally have greater feasibility when conditions are favourable and less when they are weak. The higher-tier government therefore needs to specify how feasibility will be measured over the cycle. Should feasibility be measured under a time-invariant, “neutral” economic environment, or against the current economic environment? If it is the latter, should lower-tier jurisdictions be required to upzone during downturns when feasibility falls? (And, what about the converse, that they could be allowed to *downzone* during market upswings?) This problem is also solvable, but it requires a strong framework for how we want the business cycle to interact with development.

Relatedly, in some areas feasibility will be low through much of the cycle, not because of planning controls but because prices are too low to cover construction costs. In these markets zoning binds only weakly, and upzoning would add little feasible capacity. How should these areas be treated? Should they be exempt from feasibility targets during downturns if no policy settings can generate substantial feasibility? And how should those exemptions operate politically if marginal or low-feasibility markets receive exemptions while wealthier and more central areas do not?

Governments also need the capacity to construct, maintain, and validate these feasibility models. As discussed elsewhere in this submission, feasibility models are complicated, and many planning regulations are difficult to quantify, particularly at the site level. Larger jurisdictions can plausibly develop the required modelling capacity. It is less clear that the same is worthwhile for jurisdictions such as the Northern Territory, or for small regional councils. In some parts of the country, detailed feasibility modelling may simply not be cost-effective, although some form of planning accountability is still required. Governments therefore need to assess these trade-offs.

5.3 Direct reform

The third option is to bypass decision-making by the lower-tier jurisdiction and target planning regulations directly, rather than attempting to induce reform through carrots, sticks, or performance targets. The simplest version would be for state governments to amend their planning systems to permit three-storey medium-density development throughout most of their residential land. At the Commonwealth level, a similar approach could involve payments to state governments conditional on implementing reforms, rather than for meeting targets or feasibility targets.

The Commission appears to have recommended something broadly in this direction in its interim report. Critics have argued that such an approach is blunt and could lead to inefficient use of infrastructure, additional congestion, or development in neighbourhoods that are poorly equipped to accommodate it.

We agree that the approach is blunt, but its bluntness is probably a virtue rather than a defect. A broad rule reduces the scope for politically powerful communities or concentrated interests to obtain carve-outs. It also allows market forces to play a larger role in determining where development occurs. Development should therefore tend to occur where underlying demand to live is strongest, helping to produce housing in places where people actually want to live.¹¹

Broad reform can also create large amounts of feasible zoned capacity quickly. This type of policy has been implemented successfully in places such as Auckland and Lower Hutt, as discussed in our first submission. This approach is particularly useful for jurisdictions without the capacity to construct and maintain sophisticated feasibility models, including smaller jurisdictions and regional communities.

One weakness of direct reform is that the complementary planning settings still need to be calibrated appropriately. The state planning system must ensure that controls such as floor area ratios, height limits, setbacks, and related standards allow the intended development to be financially and physically feasible. As discussed earlier in this submission, a reasonable rough approach for townhouse development may simply be to loosen all these constraints substantially.

A second weakness is that land-use reform can generate genuine local externalities that broad reforms will not target precisely, as we have discussed at length above.

As a first-order approximation, however, if the benefits from additional housing are large and the externalities associated with medium-density development are relatively small, broad direct reform is likely to generate substantial benefits. It therefore represents a plausible way of increasing housing supply where more targeted systems are difficult to design or enforce.

To summarise, housing targets, feasibility targets, and direct reform are all likely to somewhat improve on a system with little meaningful accountability for housing supply. They differ primarily in how they trade off subsidiarity, political accountability, management of local externalities, administrative capacity, and the likelihood of producing additional construction.

Our preference is for some combination of feasibility targets and direct reform in larger jurisdictions, with greater reliance on direct reform in smaller jurisdictions that lack the capacity to develop sophisticated feasibility models. Further evidence on which land-use restrictions actually bind would help both to calibrate feasibility models and to identify the most effective forms of direct reform. The externalities associated with local land-use reform should also continue to be studied so that these trade-offs can be assessed more precisely.

6 Responses to selected critiques

As the evidence that zoning regulations raise prices and lower quantities becomes more and more overwhelming, and as tools like feasibility modelling explain more and more of what is going on in these markets, critics of zoning reform are changing their arguments to explain the new evidence. We discuss two of these below.

¹¹One critique of blunt upzoning is that some areas have clearly inadequate infrastructure or are constrained by factors such as natural-disaster risk. Insofar as such areas exist, they can be explicitly exempted. What matters is that exemptions follow objective criteria rather than being negotiated area by area. Otherwise they become exactly the carve-outs that broad reform is designed to prevent.

“Reform does not add housing, it just moves it around.” An emerging critique is that if a city would have built 100,000 homes without reform, it still builds 100,000 afterwards, only in different locations.

This argument is highly useful for sceptics, as the counterfactual is never observed and so can be whatever they want. Critics of the Auckland studies, for example, attribute the construction boom to the business cycle or high migration rather than zoning reform (Donovan and Maltman, 2025). As we do not observe the true counterfactual, due to the fundamental problem of causal inference, this claim is impossible to definitively refute.¹²

However, the argument that “zoning just moves housing around” implies that it creates displacement from comparable areas. Applied microeconomic work is getting better at developing methods to identify these displacement effects. Where treated and untreated areas are spatially distinct, researchers can compare nearby areas with those further away, on the premise that nearby places are closer substitutes and should experience more displacement. In every context we are aware of, local quantity spillovers measured this way are at most modest and often absent (Behrens et al., 2026b; Büchler and Lutz, 2024; Anagol et al., 2025). Conceptually similar city-level comparisons likewise find displacement nowhere near large enough to offset the treatment effect (Greenaway-McGrevy, 2026; Maltman and Greenaway-McGrevy, 2025).

This design is easiest when a reform is geographically concentrated and its boundary is clear. Auckland’s patchwork of upzoned and non-upzoned areas makes the exercise harder. One approach was to bound the spillover required to overturn its estimate. Empirical work suggests it would need to be implausibly large (Greenaway-McGrevy and Phillips, 2023).

In addition to the empirical work, feasibility modelling shows that zoning reform shifts the supply curve out, and so should increase overall construction at all price points. Figure 7 traces the model-implied supply curves under Campbelltown’s three regimes. Holding total quantity fixed after that shift requires demand to be perfectly inelastic. If demand really were perfectly inelastic, we would instead observe sharp price falls without higher quantities, which would itself be a good thing and would achieve the central policy objective of reform.

¹²The beauty of the analysis of Auckland is that any reasonable attempt to build a counterfactual, from raw data comparisons within Auckland, between cities, or against Australian cities, finds extremely large effects. Other areas of New Zealand also had a strong business cycle over the 2010s, and importantly, Auckland built a lot during the recessionary early 2020s when it also had low or negative population growth. Other areas of New Zealand had higher rates of population growth, and yet built less. As such, the only way to dispute the Auckland research is by claiming that record-breaking quantities would have been reached in the absence of reform, which did not occur elsewhere in non-upzoned New Zealand, across multiple market conditions, while real prices have fallen.




Figure 7: Model-implied child-dwelling supply schedules under the three Campbelltown regimes

The raw data are also inconsistent with complete displacement. Construction in Auckland and Lower Hutt shifted above their historical averages and stayed there across different market conditions. If those gains merely displaced construction elsewhere in New Zealand, national construction should have been flat (as a rough approximation). Instead, as Figure 11 of our first submission reports, New Zealand built more housing per capita than Australia over the late 2010s and early 2020s, despite a weaker macroeconomic position and never having done so earlier this century (Statistics New Zealand, 2026; Australian Bureau of Statistics, 2026).

“Developers won’t build if prices fall”, or “if we build the market will be over-supplied.” Another claim we increasingly hear is that if housing reform succeeds in lowering prices, construction will slow or stop. Prices will therefore rise again, supposedly undoing the original fall.

This argument is circular: if we lower prices, prices will rise again, so apparently we never really lowered prices at all. This circularity comes from the mistake called “reasoning from a price change”: trying to infer what is happening in a market from a change in price without first understanding why the price changed.

Prices are determined jointly by demand and supply. They are an endogenous outcome of the market, not something that wanders around on its own. Prices only fall because demand shifts inward, supply shifts outward, or some combination of the two.

So when someone says, “if prices fell, development would no longer be feasible, so construction would fall”, the obvious question is: well, why did prices fall? If prices fell because housing demand collapsed, then yes: both prices and construction would fall. But that is not the mechanism being proposed when economists argue that upzoning can lower housing prices.

Upzoning is intended to shift the housing supply curve outward. An outward supply shift produces both a lower equilibrium price and a higher quantity. This is exactly what has happened in Auckland and Lower Hutt (Greenaway-McGrevy and Phillips, 2023; Maltman and Greenaway-McGrevy, 2025).

Saying that this cannot happen because lower prices would discourage construction gets the causal reasoning backwards. It is roughly equivalent to claiming that improvements in automobile manufacturing could not reduce car prices because, if cars became cheaper, manufacturers would no longer want to produce them.

In recent debates over tax and planning reform, a version of this argument has been that “we need high house prices to build housing”. By this logic, policies that might lower house prices are said to threaten new supply.

Yet much of the case for increasing housing supply is precisely that doing so should make housing cheaper. To then argue that we should avoid lower prices, because lower prices might reduce construction, risks turning the intermediate objective—building more housing—into the final objective itself. There are many reasons to value additional construction independently of its effect on prices: greater consumer choice, improved housing quality, or benefits associated with construction activity itself, for example. As discussed above, in our papers, and in our previous submission, we don’t think these are trivial. But these are rarely the primary objectives articulated by governments or advocates of housing reform. If the objective is housing affordability, then preserving high prices in order to have high supply to then have lower prices is self-defeating.

These arguments are often motivated by anecdotal claims such as, “I spoke to developers and they said that if prices fell, their projects would no longer be feasible.” What matters more is whether the opportunities that feasibility modelling predicts upzoning would create, and would be highly profitable, remain profitable at lower prices. The answer in our Campbelltown work is yes, and by a wide margin.

There is, however, a more sensible version of this argument. If additional construction lowers prices far enough, some projects close to the feasibility margin will no longer proceed. Developers do respond to prices, and construction will not expand without limit. Eventually, as more housing is supplied, the marginal project becomes unprofitable.

But this is not a rebuttal to the claim that supply reform lowers prices. It is simply how the market reaches its new equilibrium.

The substantive empirical question is therefore not whether sufficiently low prices can eventually discourage marginal development—they obviously can. It is how far the supply curve shifts following reform, how large the existing regulatory wedge is, and how much construction remains feasible as prices fall. In high-demand parts of major cities where that wedge is large, a substantial outward shift in supply can support considerably more construction even at materially lower prices.

References

- Anagol, S., Ferreira, F. V., and Rexer, J. M. (2025). Estimating the economic value of zoning reform. *American Economic Journal: Economic Policy*.
- Auditor-General’s Department of South Australia (2018). Land services commercialisation project. Report 12 of 2018.
- Australian Bureau of Statistics (2026). Building approvals, Australia. Australian Bureau of Statistics.
- Behrens, D., Gibbons, T., and Maltman, M. (2026a). What local land use reform does and does not do. Working paper.
- Behrens, D., Gibbons, T., and Maltman, M. (2026b). Where do minimum lot sizes bind? Evidence from a policy experiment.
- Blanco, H. and Neri, L. (2025). Knocking it down and mixing it up: The impact of public housing regenerations. *Review of Economics and Statistics*. Advance online publication.
- Blanco, H. and Neri, L. (2026). Neighborhood change and local economic activity. CESifo Working Paper 12891, CESifo.
- Borraz, F., Carozzi, F., González-Pampillón, N., and Zipitúa, L. (2024). Local retail prices, product variety, and neighborhood change. *American Economic Journal: Economic Policy*, 16(1):1–33.
- Büchler, S. and Lutz, E. (2024). Making housing affordable? The local effects of relaxing land-use regulation. *Journal of Urban Economics*, 143:103689.
- Clarke, H. R. and Reed, W. J. (1988). A stochastic analysis of land development timing and property valuation. *Regional Science and Urban Economics*, 18(3):357–381.
- Coates, B., Moloney, J., and Bowes, M. (2025). More homes, better cities: Letting more people live where they want. Technical report, Grattan Institute, Melbourne.
- Cooper, G., Greenaway-McGrevy, R., and Jones, J. A. (2026). Measuring the cost of land inputs to housing construction. *New Zealand Economic Papers*, 60(1):1–23.
- Donovan, S. and Maltman, M. (2025). Dispelling myths: Reviewing the evidence on zoning reforms in Auckland. *Land Use Policy*, 151:107498.
- Freedman, M. and Owens, E. G. (2011). Low-income housing development and crime. *Journal of Urban Economics*, 70(2–3):115–131.
- González-Pampillón, N. (2022). Spillover effects from new housing supply. *Regional Science and Urban Economics*, 92:103759.
- Greenaway-McGrevy, R. (2026). Can zoning reform increase housing construction? Evidence from Auckland. *Economic Modelling*, 160:107592.
- Greenaway-McGrevy, R., Pacheco, G., and Sørensen, K. (2021). The effect of upzoning on house prices and redevelopment premiums in Auckland, New Zealand. *Urban Studies*, 58(5):959–976.
- Greenaway-McGrevy, R. and Phillips, P. C. B. (2023). The impact of upzoning on housing construction in Auckland. *Journal of Urban Economics*, 136:103555.
- Guthrie, G. (2010). House prices, development costs, and the value of waiting. *Journal of Urban Economics*, 68(1):56–71.

- Hamilton, E. (2024). The effects of minimum-lot-size reform on Houston land values. *Cityscape*, 26(3).
- Jenner, K. and Tulip, P. (2020). The apartment shortage. Research Discussion Paper 2020-04, Reserve Bank of Australia.
- Kendall, R. and Tulip, P. (2018). The effect of zoning on housing prices. Research Discussion Paper 2018-03, Reserve Bank of Australia.
- Land Services SA (2026). Value added reseller licence. Land Services SA.
- Li, X. (2022). Do new housing units in your backyard raise your rents? *Journal of Economic Geography*, 22(6):1309–1352.
- Maltman, M. (2025). Zoning reform and construction sector productivity. Working paper.
- Maltman, M. and Greenaway-McGrevy, R. (2025). Going it alone: The impact of upzoning on housing construction in Lower Hutt. *Journal of Housing Economics*. Also Auckland EPC Working Paper 018.
- Moorhead, M., Armitage, L., and Skitmore, M. (2024). Real estate development feasibility and hurdle rate selection. *Buildings*, 14(4):1045.
- Neri, L. (2024). Moving opportunities: The impact of mixed-income public housing regenerations on student achievement. *Journal of Public Economics*, 230:105053.
- New Zealand Ministry of Business, Innovation and Employment (2024). Building and construction sector trends annual report 2023. Wellington.
- NSW Department of Planning, Housing and Infrastructure (2026a). Online DA data API. NSW Planning Portal.
- NSW Department of Planning, Housing and Infrastructure (2026b). Online PCC data API. Data NSW.
- Odayappan, S. (2026). Knock, knock: The impact of large subsidized housing developments on the surrounding community. *Journal of Housing Economics*, 73:102143.
- Pennington, K. (2025). Rents, movers, and stayers: The spillover effects of construction in San Francisco. Working paper, November 2025. Earlier version: “Does Building New Housing Cause Displacement? The Supply and Demand Effects of Construction in San Francisco” (2021), SSRN 3867764.
- Rajaratnam, N., McCarthy, D., Berry, T., and Shead, B. (2024). Cost and feasibility estimates for supplying new residential dwellings in NSW. Technical report, Centre for International Economics, Sydney. Supplementary modelling released with the NSW Productivity and Equality Commission’s *Building More Housing: Where It Stacks Up*.
- Reserve Bank of Australia (2024). Housing market cycles and fundamentals. Speech, 16 May 2024.
- Rollet, V. (2025). Zoning and the dynamics of urban redevelopment. Job Market Paper, MIT.
- Saunders, T. and Tulip, P. (2019). A model of the Australian housing market. Research Discussion Paper 2019-01, Reserve Bank of Australia.
- Shortell, J. (2022). The effect of a minimum lot size reduction on residential property values: The case of Houston. Master’s thesis, Universitat de Barcelona.

- Statistics New Zealand (2026). Building consents issued. Stats NZ.
- Tanrisever, I. (2025). Spillover effects of accessory dwelling unit development. *Regional Science and Urban Economics*, 114:104136.
- Tso, M. (2019). Lower Hutt ratepayers gift \$29m to developers through waived fees. Stuff.
- Tulip, P. (2022). The need for reform of land use policy: Submission to the Productivity Commission. Technical report, Centre for Independent Studies.
- Valuation NSW (2026). Bulk property sales information. NSW Government.
- Zhuo, Y., Owens, E., and Chen, M. K. (2025). Does neighborhood investment actually affect crime? New evidence from LIHTC and smartphone-based measures of policing. Working paper, January 2025. Earlier version: UCLA Ziman Center Working Paper 2024-04.