

Dr Cameron K. Murray
Chief Economist
Fresh Economic Thinking

Sue Higginson MLC
Committee Chair
NSW Legislative Council Portfolio Committee No. 7 – Planning and Environment
Committee

17th June 2024

Dear Chair,

Thank you for the opportunity to provide a written response to evidence provided on the 7th June 2024 to the committee by Dr Peter Tulip, Chief Economist, Centre for Independent Studies, and Mr Brendan Coates, Economic Policy Program Director, Grattan Institute.

The testimony of these witnesses involved claims about how upzoning increased housing supply in New Zealand, particularly in Auckland. However, the evidence from Auckland is not as strong as many believe.

In 2023, Director of Research at Prosper Australia, Dr Timothy Helm, and myself, published a methodological critique of one of the major studies of the effect of Auckland's 2016 upzoning on new housing production.¹

I am unaware of any reputable expert economist who has argued that our critiques are wrong and have attached both the original critique, and a follow up, to this letter.

Auckland's experience shows that housing markets adapt physically due to planning rules, but that the quantity of new homes produced is regulated by market forces.

Here I provide additional evidence (for publicly distribution) for the committee's consideration.

Summary of important economic context:

1. Planning controls regulate the location and density of housing types, not the quantity of new homes per period. The quantity of homes produced per year is the combination of the density per project *and* the number of projects.

$$\text{Quantity (homes/period)} = \text{Density (homes/project)} \times \text{Frequency (projects/period)}$$


¹ Greenaway-McGrevy, Ryan, and Peter CB Phillips. "The impact of upzoning on housing construction in Auckland." *Journal of Urban Economics* 136 (2023): 103555.

Project density is limited by both market prices (high density is economical in high value areas) and planning controls, but the frequency, or the number of housing projects in a period, is unregulated and can respond to market forces.

2. Across the housing market there is an equilibrium rate of housing production per period that trades-off the benefit of building today and delaying and building later. This rate is very responsive to market conditions, and indeed, Sydney saw during the 2010s period more than a tripling of building approvals from the 2006-09 lows, before falling 35% from their 2017 peak (below image).

Sydney new dwelling activity

Completions fall far below building approvals during booms



Sources: Greater Cities Commission, Australian Bureau of Statistics

fresheconomicthinking.com




3. Individual cities cannot reduce their housing rents relative to other cities without inducing inwards migration. A spatial equilibrium means that the relative rental premium of more desirable locations reflects the value of moving.
4. Housing concerns in Sydney are not unique. Cities across the world are having the same housing debates, all with their own planning and local policy settings.

For context, Sydney asking rents in mid-2022 were the as they were in mid-2016 (nominally), and Sydney rents in the consumer price index (CPI) have risen less than the overall CPI for about 15 years. Recent rapid rental price growth is mostly due to high household income growth associated with general inflation.

Summary of important Auckland context:

1. Housing markets are cyclical. Extrapolating a straight-line will not predict well future paths of housing approvals, commencements, or completions (see the above Sydney chart). However, the main study cited earlier does this, implying a rapid decline in their growth rate of dwelling consents (equivalent to building approvals) if not for the 2016 Auckland upzoning.
2. Dwelling consents are not completions. The compositional change in housing types created by upzoning meant more homes demolished per new home. The chart below shows that although consents spiked (as they did in Sydney's 2010s boom), only since 2023 have net additions to the housing stock grown rapidly. However, this recent spike in completed new dwellings has coincided with rents rising faster in Auckland than in Wellington and other urban areas.



Planning regulations can allow for housing development in some areas but not others for many valid reasons. It appears that this has been the case in Sydney for many years and supply has flexibly responded to cyclical demand.

The Auckland myth: There is no evidence that upzoning increased housing construction

The one study showing it did is methodologically flawed and contradicted by the data



CAMERON MURRAY AND TIM HELM

JUN 04, 2023



11



15



2

Share



Thank you to the 3,000 subscribers who read Fresh Economic Thinking.

I want to ask a small favour. If you enjoy these articles, the best thing you can do to support me is to share them on your social media and forward the emails to your friends.

And don't forget to subscribe. It's free after all.

Many housing analysts argue that large-scale upzoning policies create affordable housing. This was a justification for the major upzoning in the 2016 Auckland Unitary Plan (AUP). The Auckland experience is now of intense interest in housing debates.

Widely cited in media reporting is a paper by Ryan Greenaway-McGrevy and Peter Phillips (GMP) entitled [The Impact of Upzoning on Housing Construction in Auckland](#).

[Here's the Vancouver Sun](#) covering it. The study even made it to the [New York Times](#).

In areas that were upzoned, the total number of building permits granted (a way of estimating new construction) more than quadrupled from 2016 to 2021.

Closer to home, the [SMH](#), [The Conversation](#), the [AFR](#), and [housing_policy_websites](#) have covered the topic and this research.

Eyeballing the data from GMP's paper in the chart below, it looks like Auckland's upzoning policy was a huge success—there was an obvious boom in dwelling consents (building approvals) in the upzoned areas after implementation.

The conclusion of GMP's paper is that


21,808 additional dwellings were permitted over the five years following the zoning reform, corresponding to approximately 4.11% of the dwelling stock of the Auckland region. ¹

Astonishingly, they note that

permits issued per year have approximately doubled over the five subsequent years to the reforms.

But does it pass the sniff test?

Auckland dwelling consents ('000s) — Upzoned — Non-upzoned



An initial sniff test

Unpacking the upzoning effect is a tricky empirical problem. While the authors have piloted an innovative approach to solving it, there are issues with the analysis that we think seriously undermine the main result.


It is the nature of research to disagree on methods, assumptions, and interpretations. We welcome rigorous feedback on any of our papers; you may find errors, strange assumptions, or contestable interpretations that should be discussed.

We provided feedback to the authors, who have [now published](#) some extensions and sensitivity testing of their analysis.

The next graph provides the first clue to the problem.

It plots the annual dwelling consents in New Zealand's two largest urban areas, Auckland and Wellington. One of these cities had a home-building boom that doubled housing consents compared to the pre-2016 trend, thanks to upzoning. The other has a zoning policy some describe as "[a racket to stop housing being built.](#)"

Can you pick which is which? Is Auckland City A or City B?



Data coverage

The first issue with GMP's paper is the partial dataset used for the analysis.

Unfortunately, it's a biased sample.


Dwelling consents from a separate Special Housing Area (SpHA) program that ran prior to the 2016 upzoning are removed, as are certain locations in the city. SpHAs were an inclusionary zoning policy enacted in late 2013 that allowed developers across New Zealand to build at higher density in exchange for providing a proportion of new dwellings as affordable housing.

The authors explain that:

Care is taken to remove construction activity generated by other government policies during our sample period. In particular, consents issued under ‘Special Housing Area’ (SpHA) authority are omitted from our sample.

When the AUP began in 2016, these SpHA approvals became redundant in Auckland. These approval types dropped away fast in the consent data, showing they were true substitutes—substitute locations for buyers and substitute markets for developers and builders. They were part of the same city-wide housing market.

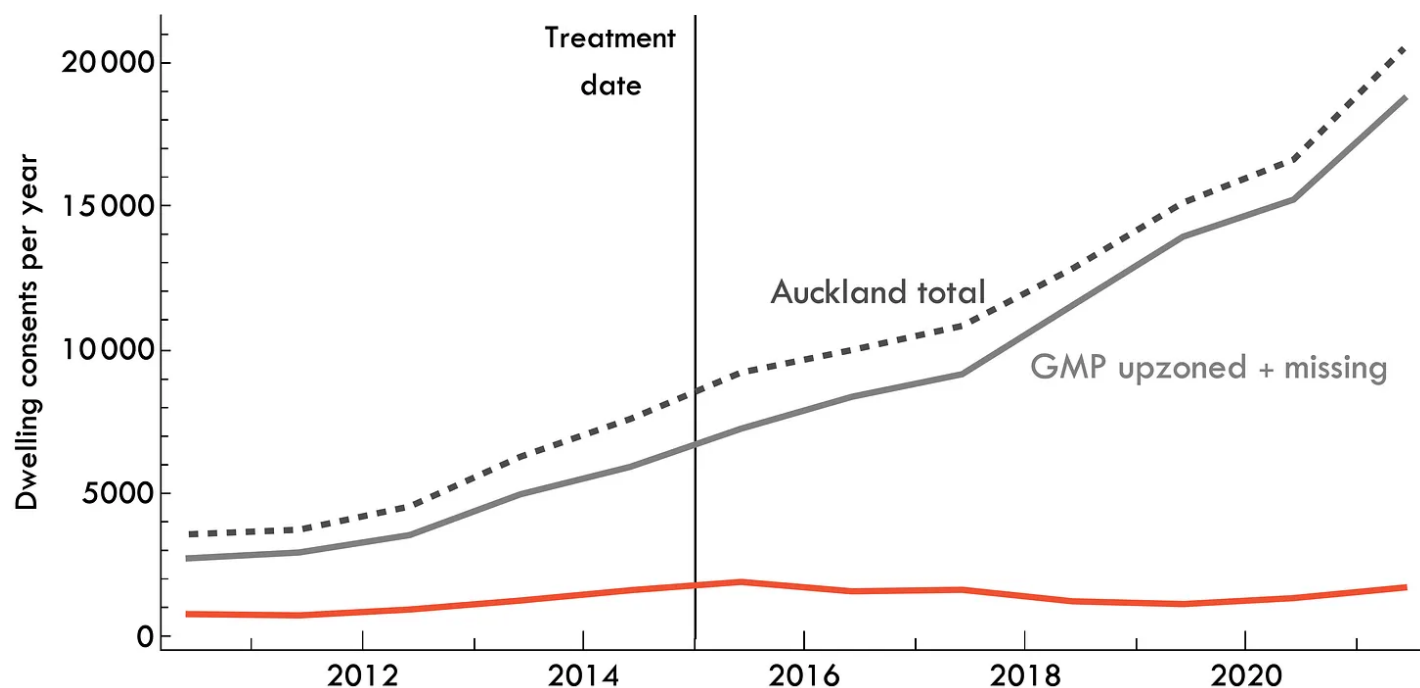
The light grey line in the chart below shows dwelling consents for the upzoned areas in GMP’s sample. The kink in the line suggests a clear change in trend in 2016.



But the total dwelling consent data for Auckland, in dark grey, shows no such change.

The dashed black line shows the data missing from GMP's sample, which by its absence generates the kink in their upzoned area data.

When the missing consents are included in the upzoned area total, the chart looks like this, and the kink disappears.



Excluding this data effectively converts a growth trend for Auckland indistinguishable from that in Wellington, which was not upzoned, into a series with a distinct structural break.

In their extension paper, the authors include these missing dwelling consents, noting:

Total permits no longer exhibit a substantial break in trend in 2016, when the AUP became operational.

Linear assumption and identical pre-AUP trends

The second issue is that the counterfactual from which GMP estimate the upzoning effect assumes (a) linear growth, and (b) identical trends in upzoned and non-upzoned areas prior to the AUP.

This introduces significant biases, as we will show.


The chart below explains GMP’s method of estimating the net additional dwelling consents caused by upzoning.

First, the straight line of best fit for non-upzoned areas from 2010 to 2015 is extrapolated out to 2021 to define a counterfactual for non-upzoned areas. That’s the lower of the two straight grey lines.

Then, the pre-2015 levels difference between upzoned and non-upzoned areas is added to this to define a counterfactual for the upzoned areas. That’s the top grey straight line.


Next, the difference between actual and counterfactual non-upzoned area consents is calculated in order to estimate the spillover from these areas to the upzoned areas. That’s the grey shading. The idea is to ensure the policy effect isn’t over-cooked by counting pure substitution of activity from one place to another as an increase.

Finally, this spillover is subtracted from the deviation of upzoned area consents above its counterfactual, in order to estimate the additional dwelling consents net of spillovers (the red shaded area).



To our point (a), there is no reason the counterfactual trend should be linear. Not many

economic trends are. Fitting a *curve* to the pre-treatment trend fits that data better, and extrapolating it produces a bigger spillover and smaller net (spillover-robust) treatment effect. By our estimate, the net effect falls from 22,000 extra dwelling consents to about 4,500 with this one small change.



To our point (b), assuming upzoned areas have the same pre-treatment trend as non-upzoned areas implies a fixed difference between the two counterfactuals in all future years (about 2,000 consents, in GMP's sample). But what if the upzoned areas were already growing at a faster rate pre-treatment?

The authors note this concern and explain it is in principle possible to deal with it:


...our approach to modelling counterfactual scenarios based on pre-treatment trends in the control group could easily be extended to incorporate trends in *both* the treatment and the control group.

However, in practice, they assume the non-upzoned trend also applies to the upzoned areas (and hence to the market as a whole).

For their partial sample, this is fine. But for the full data, the assumption of linear and parallel trends is clearly inappropriate. Total consents in Auckland were growing non-

linearly, so there is no reason to expect the levels difference between sub-areas to be fixed over time.


When we replicate their method with the full data (counting all missing consents in upzoned areas) and with both (a) non-linear and (b) independent pre-treatment trends, the estimate of a spillover-robust treatment effect turns *negative*.



We're not suggesting the AUP actually *reduced* consents. Extrapolating growth this far forward is [unrealistic](#). And this is part of our point: whether using a linear or non-linear trend, extrapolating a short and highly-cyclical series a long way into the future is an inherently unreliable way of defining a counterfactual.

The chart below from the extension paper shows the enormous cyclical variation in dwelling contents. On it, we've marked in blue the linear trend the authors extend into the future for the non-upzoned areas. The authors assume an identical linear trend and fixed levels difference for upzoned areas. So in pink, we've shown what this implies as their counterfactual assumption of total dwelling consents without the AUP upzoning.

It doesn't match the data (the black line) at all. [2](#)



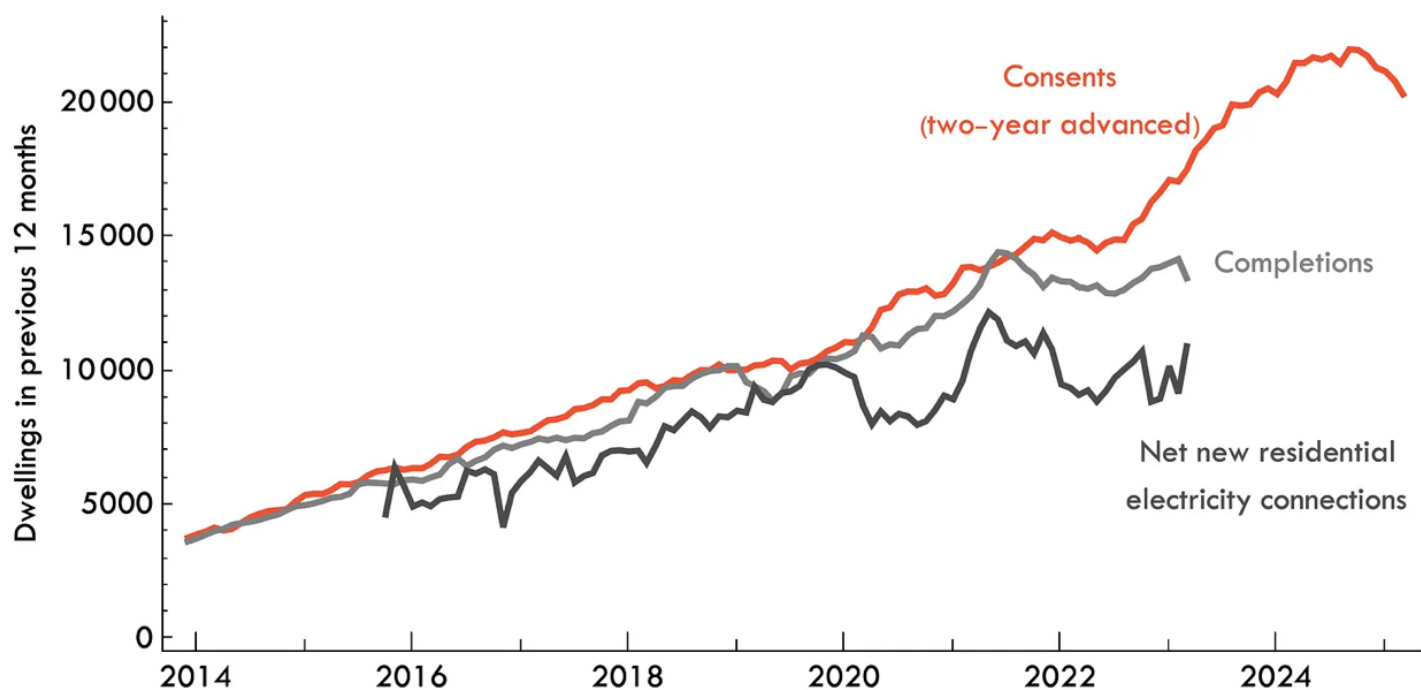
The authors re-estimate the policy effect in their extension paper, this time from a treatment date of 2013 (when the SpHAs began). As in the original, actual permits over the treatment period—in this case the nine years from 2013 to 2022—are compared to a counterfactual linear trend.

This kind of linear extrapolation is a highly unreliable guide to the counterfactual outcome.

Interpretation of consents as extra dwellings

A final note of caution concerns the interpretation of dwelling consents as extra dwellings. Historically, about 90% of consents become completed dwellings after two years, as the authors note.

Recently, however, net additional dwellings, as measured by the change in the number of residential electricity connections, have not grown as fast as completions would suggest. Net additional dwellings two years after approvals fell from 77% prior to 2018 to 69% since 2020. This implies that more existing homes are being demolished for each new home.



The total number of residential electricity connections increased by 49,726 in Auckland between January 2016 and January 2022. It seems generous to imagine that almost half of those would not have occurred but for the AUP—especially given similar trends in Wellington we noted earlier.

We agree that more homes are better than fewer, but good planning matters too. Organising the location and types of dwellings can reduce total infrastructure costs and avoid social costs that cause a [net loss in utility](#).

The disorderly nature of where post-AUP development occurred is [hated by even the most strident supporters of the policy](#) now [living through it](#).

The single road in and out of our suburb is a shitshow in the morning. And in the afternoon. And on weekend mornings. Yep, we have traffic jams trying to leave our suburb at 10am on a Saturday. The tacit agreement we made when accepting density

out west was that it would come with improvements to public transport and infrastructure; or at the very least that dense inner suburbs would reduce the load on our outer transport links. The best we've seen is a half-assed bus lane and [ridiculous bar leaners instead of seats in bus stops](#).

If upzoning could deliver a benefit in the form of more homes and lower prices, even with the social costs of disorderly and haphazard development locations, it would be a trade-off worth making. But where are these benefits?

This is why housing supply analysis is hard

Housing and property analysis struggles with the fact that markets are cyclical, locations are substitutes, consumption responds to price, and individuals migrate to improve their quality of life. Housing analysis faces many of the same problems as empirical macroeconomics—a lack of controls, idiosyncratic local factors, and more.

It may simply be that GMP's approach of estimating a policy effect on a housing market without reference to trends in comparable markets, by choosing a control group that is essentially *within* the true intervention group—the Auckland housing market—just cannot be made to work.

Auckland's consent trend, like Wellington's, looks a lot like a growth cycle spurred on by an uptick in migration around 2014, and a normal cyclical boom, one also seen in Australia's major cities following the post-GFC recovery period.

Judgement on whether the AUP delivered the goods will need to await a fine-grained comparison between Auckland as a whole and comparable cities. On the basic data thus far, we've seen nothing to justify the headlines.

Thanks for reading Fresh Economic Thinking!
Subscribe for free to receive new posts and
support my work.

- 1 An earlier version of the paper [dated May 2022](#) concluded that 26,903 additional dwellings were consented as a result of the policy (equal to 5% of the pre-AUP housing stock). The numbers cited above are from the [May 2023 version](#), in which a calculation error in the original has been corrected.
- 2 Thanks to Matthew Maltman, author of the [onefinaleffort.com](#) blog, for pointing out an [error](#) in the original version of this chart.



11 Likes · 2 Restacks



A guest post by

Tim Helm

Dr Tim Helm is a freelance economic consultant and the Director of Research and Policy at think tank Prosper Australia.



The Auckland upzoning myth: Response to comments

We identified three major flaws in a landmark paper about the construction effects of upzoning. No-one disagreed with these flaws. Now we respond to three new questions.



CAMERON MURRAY AND TIM HELM

AUG 27, 2023



10



12



1

Share



A widely cited paper in the zoning and housing supply debate is *The impact of upzoning on housing construction in Auckland* by Ryan Greenaway-McGrevy and Peter Phillips (GMP), [published](#) in the Journal of Urban Economics (free version [here](#)).

It's a popular reference for claims that city-wide upzoning boosts new housing supply.

Our [last post](#) in June laid out three major methodological flaws with this paper. There has been no disagreement with these three main points.

This post addresses [responses](#) from our readers to that post, explains in even simpler terms how GMP's estimates were derived, and provides a spreadsheet replicating the results, so you can see this for yourself.

We didn't claim to prove that upzoning had *no* effect on new construction. We claimed that GMP failed to use methods capable of identifying one.


No disagreement on the major flaws

No one disputed our three core points.

1. Biased sample

Did you see the [AFR](#) chart showing a clear uptick in upzoned areas? It's an artefact of using a biased data sample. There was no structural break in total consents in the full sample.

Auckland dwelling consents ('000s) — Upzoned — Non-upzoned



No one disputed this.

The only question was whether Ryan Greenaway-McGrevy's May 2023 [extension paper](#), which applied the same method as in the published paper to the full data sample, had adequately addressed the problem. We deal with that below.

2. A linear counterfactual in cyclical data

Ever climbed Rangitoto? It's a steady gradient – until it's not. There was no disagreement with our claim that applying a straight-line assumption to forecasting future growth in a cyclical data series is a practice fraught with danger.

The paper relies on an unrealistic counterfactual in which a linear assumption implies that the growth in building consents would fall to half the pre-upzoning rate.



3. Consents are not new homes


The paper didn't do what it said on the tin: it measured building consents (approvals) rather than dwelling construction.

Stockpiling pieces of paper with approval to build (consents) does not grow the stock, nor does replacing one building with another.

Net additions, as you can see below, have diverged from consents. No one disputed that the point under debate was the effect of zoning on *approvals/consents*, not on construction or net additions.

Auckland new dwelling activity

Consents are diverging from net increases in the dwelling stock



Sources: Auckland Council, StatsNZ, EMI of the Electricity Authority

Fresh Economic Thinking (.com)

Readers' responses

The responses to our post boiled down to three key questions:

1. Is that *really* the counterfactual? Surely not.
2. Doesn't the extension paper resolve this?
3. If upzoning didn't change supply, why did rents fall?

Let's take each in turn.

1. Is that *really* the counterfactual?

A friend of ours reckons that only science fiction writers and economists are truly happy working with 'counterfactuals'.

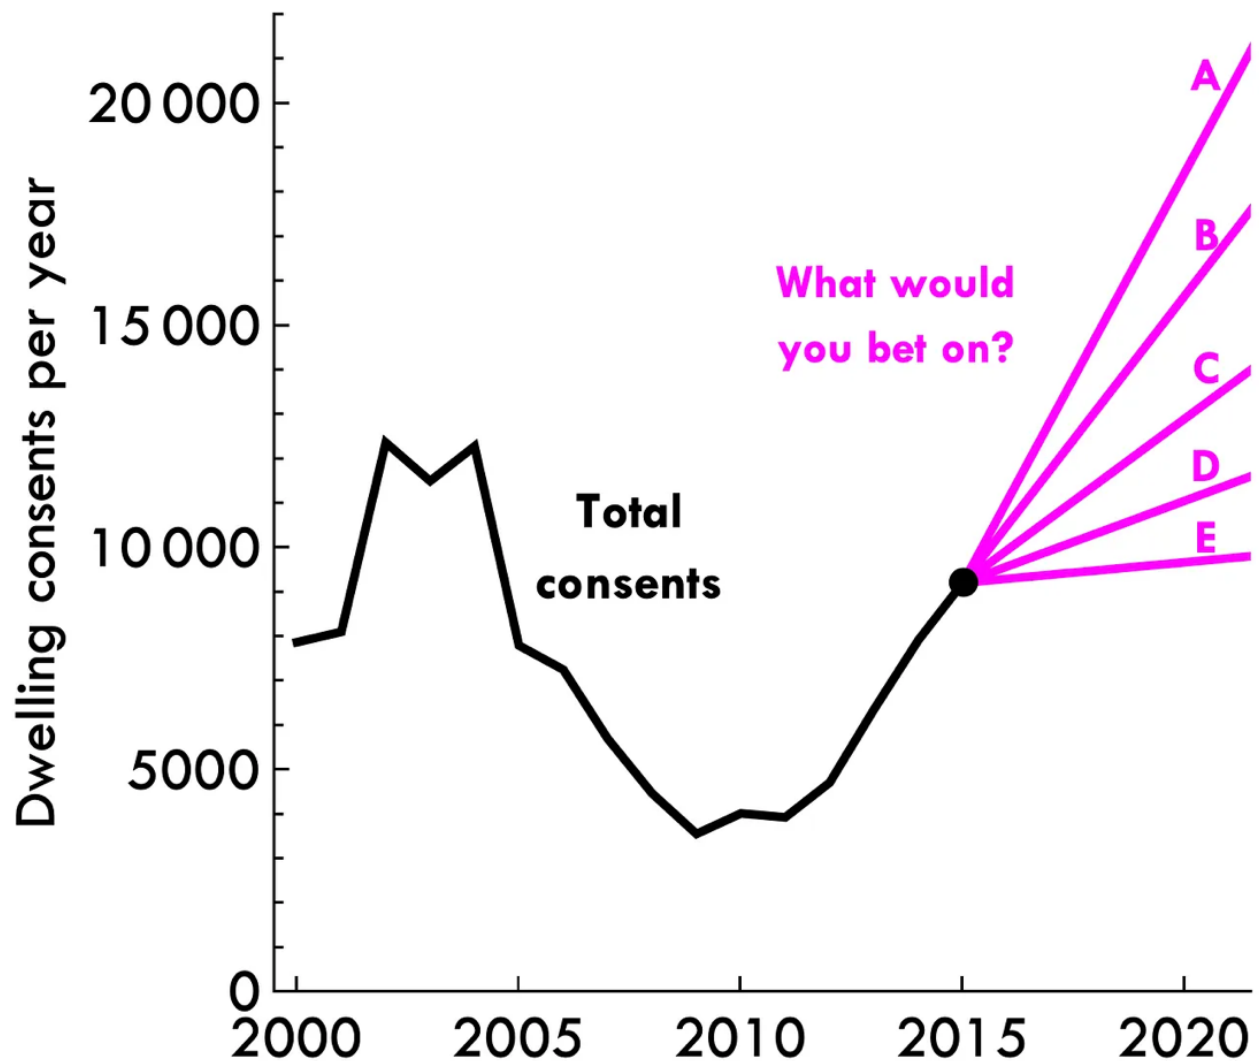
Was the counterfactual GMP used to estimate growth in consents due to upzoning realistic? Here's a test for you.

It's the end of 2015.

Building is booming. Consents have grown by 12% per annum since their post-GFC low to hit 9,000 per annum in 2015 (still a few thousand short of the pre-GFC high). And after a decade of zero net migration, immigration has kicked up rapidly. NZ's population is now growing 40% faster than Australia's in relative terms.

After a marathon debate, the proposed Auckland Unitary Plan (AUP) is rejected. Upzoning advocates hoped to see zoned capacity for new dwellings triple from 300,000 to 900,000. Instead, they are disappointed. Zoning rules stay as they are.

If you lived in 2015 in this alternative no-AUP world, which path would you bet on in the image below for dwelling consents?



Would you have picked D?

We wouldn't have either. But D is the counterfactual used by GMP to conclude that anything above this is the effect of the AUP on new dwelling consents.

Readers expressed disbelief at this.

It's hard to see it in the paper as the details are buried within the econometrics. We explained it in a simple, transparent way so that you can be the judge.


Are we sure? Yes. We read this paper cover to cover and replicated the method to successfully reproduce the results. Then we checked our understanding and replication with the authors.

GMP effectively assume that, without the AUP upzoning, growth in consents would

suddenly have slowed down. Over the five years prior to the AUP, annual growth in consents in GMP's sample averaged 12.1%. But GMP's counterfactual for the six years following involves an average annual growth of just 5.7%. So it's no surprise they found upzoning doubled growth.

Whether the authors appreciated this fact, or lost their intuition amid the econometrics, is beside the point.


You can now replicate GMP's results yourself. We've provided a spreadsheet to do just that. It uses just nine data points and three columns to replicate the paper's main results of 22,000 extra consents from upzoning. [1](#)



This is our simplest explanation of the method.

1. Take the 2015 non-upzoned data point, subtract the 2010 point, and divide by 5. That gives you the slope used to generate a counterfactual.
2. Begin from the 2015 total. Add twice the slope previously calculated (once for each of the two groups the total is partitioned into) and add that to the 2015 total to get a counterfactual 2016 figure. Then add twice the slope again for each year until 2021.
3. Subtract the actual total consents for 2016 to 2021 from those counterfactuals.


As you can see from the charts in the spreadsheet, this looks reasonable with the sample data, but quite odd when applied to the city-wide total data, as in the chart below.



The problem is that all locations and consent types are substitutes, and by removing certain data, GMP are simply tracking location substitution and misidentifying it as a net change in total consents.

What might have been better?

Comparison with other cities, for one. Or something like the “eyeball counterfactual” below. We’re not claiming this is the “right” counterfactual. But this is what similar data series look like in cities that didn’t have zoning changes.



2. Doesn't the extension paper resolve this?

In short, no.

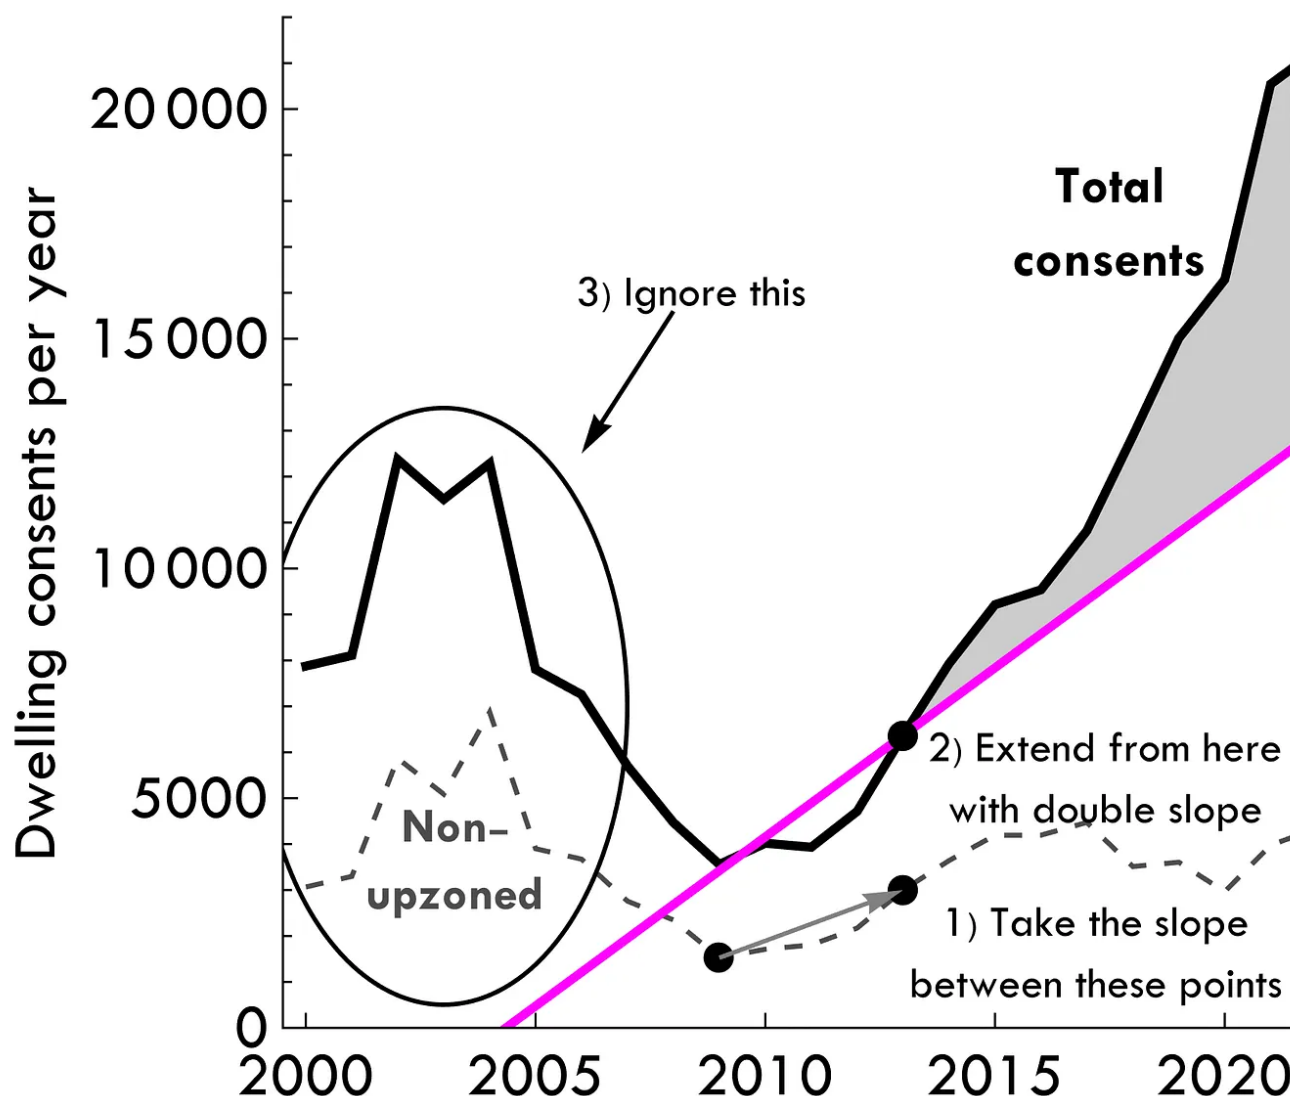
One of the authors released an [Update and Extended Results paper](#) in May 2023. It expands the sample to include data missing from the published paper, brings forward the starting point of the upzoning, and adds data for 2022.

Unfortunately, it applies the same flawed counterfactual construction.

This time, instead of extrapolating a five-year straight-line trend for six years from the date of the AUP, it extrapolates a four-year trend for nine years from the date when an earlier upzoning policy began.

Here is the extension paper result, annotated to show the method. Please apply the same

test as above. Is the pink line a realistic counterfactual for the black line?



The extension paper fixes one problem by making another problem worse. It fixes the biased data sample but worsens the problem of using linear extrapolation on a clearly non-linear series. The counterfactual literally extends a straight line from the 2009 minimum, to a point four years ahead in 2013, on this clearly cyclical data.

Using this method, *any* deviation of total consents from a straight line fitted to a few years prior to the upzoning will be attributed to the upzoning. Had immigration stalled or markets crashed in 2016, this approach would be telling us that the AUP had *reduced* construction!

3. If upzoning didn't change supply, why did rents fall?

This one is straightforward. Rents changed in accordance with population, income, preferences, and housing stock.

The stock of homes did increase quickly during the building boom. That's plain in the data. But it's the attribution of this to upzoning rather than to ordinary construction cycles we're contesting. There's plenty of evidence that while zoning effectively regulates locations and housing types, it is market forces that regulate the overall rate of new dwelling construction.

Rents subsided in Sydney for many years after a large construction cycle in the 2010s too. Maybe, as in Sydney, Auckland rents will rise again soon. When consents fall away and rents rise, will upzoning get the blame?

We don't know how many extra dwellings the AUP is responsible for. We don't think anyone else does either.

But the story that upzoning produced a huge building boom is becoming an urban myth.

Cherry-picking figures, uncritically citing a paper with known methodological issues, and writing fairy tales about a small and plucky city far away is well and good when pushing a policy agenda.

YIMBY blogger Matt Yglesias is [frank](#) (and tongue-in-cheek) at times about the role these tales serve:

Like all self-respecting Americans, I mostly care about foreign politics in order to shadow box about American issues.

But if that's your game with Auckland and upzoning, please be honest enough to admit you're playing politics, not doing economic science.

Thanks for reading Fresh Economic Thinking!
Subscribe for free to receive new posts and
support my work.

- 1 The rest of the paper just establishes the statistical significance of treatment versus counterfactual control.
-



10 Likes · 1 Restack