

The Missing Middle Stays Missing: Upzoning and Dwelling Type Composition in New Zealand

APEP Autonomous Research* @olafdrw

April 9, 2026

Abstract

New Zealand’s 2022 Medium Density Residential Standards forced five major cities to allow up to three dwellings per residential lot without planning consent. Using annual region-level building consent data and a difference-in-differences design comparing four treated regions against eleven never-treated controls, I find no evidence that this upzoning shifted the composition of new housing toward multi-unit dwellings. The multi-unit share point estimate is 0.7 percentage points (SE = 3.6 pp) — a precise null. A Callaway–Sant’Anna estimator confirms the finding (−1.8 pp, SE = 3.3 pp). Pre-treatment parallel trends hold. The result suggests that relaxing planning rules is insufficient to close the “missing middle” gap when binding constraints lie in construction costs, financing, and developer capacity.

JEL Codes: R31, R52, H70

Keywords: upzoning, housing supply, building consents, missing middle, New Zealand

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 18m).

1. Introduction

For every three houses New Zealand built in 2022, it built roughly one multi-unit dwelling — a townhouse, flat, or apartment. Despite decades of planning reform rhetoric, the “missing middle” of medium-density housing remains stubbornly under-supplied in most Anglophone cities (Parolek, 2020). In August 2022, New Zealand implemented the most aggressive upzoning in its history: the Medium Density Residential Standards (MDRS) mandated that Tier 1 cities allow up to three dwellings per residential lot, as of right, without resource consent. The reform eliminated a discretionary approval process that averaged 40–70 days and thousands of dollars in fees (New Zealand Productivity Commission, 2012). If permitting barriers matter for the housing mix, the MDRS should have produced a visible compositional shift toward multi-unit forms.

This paper tests that prediction. I exploit the staggered introduction of MDRS across New Zealand’s regions using a difference-in-differences design. Four regions containing Tier 1 cities — Waikato (Hamilton), Bay of Plenty (Tauranga), Wellington, and Canterbury (Christchurch) — received mandatory upzoning in August 2022, while eleven smaller regions remained subject to the less prescriptive National Policy Statement on Urban Development but not MDRS. Auckland, treated earlier under its 2016 Unitary Plan, is excluded. The outcome is the multi-unit share: the fraction of new dwelling consents classified as townhouses, flats, or apartments rather than standalone houses.

The main finding is a precise null. The TWFE estimate of the MDRS effect on multi-unit share is 0.7 percentage points (SE = 3.6 pp, $p = 0.84$) — economically trivial against a pre-treatment treated-region mean of 31.4%. A Callaway–Sant’Anna estimator designed for staggered adoption yields -1.8 percentage points (SE = 3.3 pp). The pre-test for parallel trends passes comfortably ($p = 0.39$). Event-study coefficients show no divergence in any pre-treatment year and no break in trend after 2022. The null survives exclusion of Canterbury (where earthquake-related construction may confound), removal of small regions, delayed treatment timing, and controls for the level of overall consenting activity.

The precise zero on composition sits alongside a suggestive negative effect on the *level* of multi-unit consenting. A specification using log multi-unit consents as the outcome yields a coefficient of -0.40 (SE = 0.17, $p = 0.03$), indicating that treated regions may have experienced a larger decline in absolute multi-unit building during the post-2022 construction downturn. This pattern is consistent with binding supply-side constraints — construction costs, developer financing, and builder capacity — that overwhelmed the demand-side relaxation of planning rules (Glaeser and Gyourko, 2018).

This paper contributes to a growing literature on the limits of upzoning as a housing

supply tool. [Greenaway-McGrevy and Phillips \(2023\)](#) find that Auckland’s Unitary Plan increased land values at upzoned parcels but produced a muted supply response. [Maltman and Greenaway-McGrevy \(2024\)](#) study a single Lower Hutt suburb and find limited compositional change. [Kulish et al. \(2012\)](#) model the supply elasticity of housing and show that planning constraints are only one of several binding margins. The most closely related paper is [Been et al. \(2023\)](#), who study Minneapolis’s 2018 upzoning and find similarly modest effects on total housing supply. My contribution is to test the *composition* hypothesis — whether upzoning shifts *what type* of housing gets built, not just how much — in a setting with unusually sharp policy variation and clean pre-trends.

The finding speaks to a broader debate about whether the “missing middle” gap ([Parolek, 2020](#)) is primarily a regulatory artifact or reflects deeper structural constraints. If planning rules were the binding barrier, their removal should have produced visible compositional change within four years. The null suggests that the constraint lies elsewhere — in the gap between what is *permitted* and what is *profitable* to build at medium density ([New Zealand Productivity Commission, 2012](#); [Ministry for the Environment, 2020](#)). For policymakers, the implication is that upzoning is necessary but not sufficient: complementary interventions in construction productivity, financing, and infrastructure provision may be required to close the missing middle gap.

The remainder of the paper proceeds as follows. Section 2 describes the MDRS reform and its institutional setting. Section 3 presents the data. Section 4 outlines the empirical strategy. Section 5 reports results and robustness checks. Section 6 discusses mechanisms and implications.

2. Institutional Background

New Zealand’s resource consent system has long been identified as a binding constraint on housing supply ([New Zealand Productivity Commission, 2012](#)). Under the Resource Management Act 1991 (RMA), construction of more than one dwelling on a standard residential lot required “resource consent” — a discretionary process involving notification, neighbor objections, and council assessment. For medium-density proposals, the process averaged 40–70 working days and cost NZD 2,000–5,000 in direct fees, with substantially higher costs when objections triggered hearings ([Ministry for the Environment, 2020](#)).

The MDRS reform. The Resource Management (Enabling Housing Supply and Other Matters) Amendment Act 2021 inserted Schedule 3A into the RMA, establishing the Medium Density Residential Standards. From August 20, 2022, all Tier 1 territorial authorities were

required to permit, as of right and without resource consent, up to three residential units of up to three stories on any urban residential lot. The five Tier 1 urban areas are Auckland, Hamilton, Tauranga, the Wellington region (including Porirua, Lower Hutt, Upper Hutt, and Wellington City), and Christchurch.

Staggered treatment. Auckland had already implemented comparable medium-density provisions under the Auckland Unitary Plan (operative September 2016). The remaining Tier 1 cities received MDRS in August 2022, though Wellington-region councils had slightly different operative dates (Porirua: September 2022; others: 2023–2024). Tier 2 cities — including Napier, Hastings, Palmerston North, Nelson, Dunedin, and others — were subject to the National Policy Statement on Urban Development (NPS-UD), which required plan changes to enable density but did not mandate as-of-right construction.

The construction cycle. The MDRS was implemented during a severe construction downturn. National new dwelling consents fell from approximately 51,000 (year ended February 2022) to roughly 33,000 (year ended February 2025) — a 35% decline driven by rising interest rates, construction cost inflation, and post-COVID supply chain disruptions ([Reserve Bank of New Zealand, 2023](#)). This downturn is the primary threat to identifying a composition effect, as it may have differentially affected the multi-unit segment.

3. Data

I use Stats NZ Building Consents data, the official source for construction activity in New Zealand. The data report the number of new dwelling consents issued by region and dwelling type (houses versus multi-unit homes, where the latter includes townhouses, flats, apartments, and retirement village units) on an annual basis for years ended February 2016 through February 2026. The unit of observation is region \times year.

Sample construction. The analysis sample comprises 15 regions \times 11 years = 165 observations. Auckland is excluded because it received comparable upzoning six years before MDRS (Auckland Unitary Plan, September 2016). Four regions are treated (contain MDRS Tier 1 cities): Waikato (Hamilton), Bay of Plenty (Tauranga), Wellington, and Canterbury (Christchurch). Eleven regions serve as never-treated controls. The pre-treatment period spans seven years (2016–2022, where “2022” denotes the year ended February 2022, entirely before the August 2022 MDRS operative date). The post-treatment period comprises four years (2023–2026).

Key variable. The primary outcome is the *multi-unit share*: the fraction of new dwelling consents in a given region-year that are classified as multi-unit homes rather than standalone houses. In the pre-treatment period, the mean multi-unit share across all regions is 21.0% (SD = 12.4 pp), with treated regions averaging 31.4% and control regions 17.9%.

3.1 Summary Statistics

Table 1: Summary Statistics

| | Treated | Control | Full Sample |
|------------------|------------------|------------------|------------------|
| Multi-unit share | 0.376 (0.132) | 0.232 (0.148) | 0.270 (0.157) |
| Total consents | 3729 (1893) | 657 (605) | 1476 (1751) |
| Houses | 2337 | 481 | 976 |
| Multi-unit homes | 1393 | 175 | 500 |
| N (region-years) | 44 | 121 | 165 |

Notes: Standard deviations in parentheses. Treated regions contain MDRS Tier 1 cities (Waikato, Bay of Plenty, Wellington, Canterbury). Control regions are all other non-Auckland regions. Data: Stats NZ Building Consents, annual year ended February, 2016–2026. Multi-unit share is the fraction of new dwelling consents classified as multi-unit homes (townhouses, flats, apartments) rather than standalone houses. N = 165 region-years (15 regions × 11 years).

4. Empirical Strategy

4.1 Identification

I estimate the effect of MDRS on dwelling type composition using a two-way fixed effects (TWFE) difference-in-differences design:

$$Y_{it} = \alpha_i + \lambda_t + \delta \cdot (\text{Treated}_i \times \text{Post}_t) + \varepsilon_{it} \quad (1)$$

where Y_{it} is the multi-unit share in region i in year t , α_i are region fixed effects, λ_t are year fixed effects, and $\text{Treated}_i \times \text{Post}_t$ is an indicator equal to one for treated regions after August 2022. Standard errors are clustered at the region level.

The identifying assumption is that treated and control regions would have followed parallel trends in multi-unit share absent the MDRS. Because all treated regions received MDRS simultaneously (August 2022), there is no staggered-adoption heterogeneity concern in the TWFE specification. I nonetheless report Callaway–Sant’Anna (2021) estimates as a robustness check, using never-treated regions as the comparison group.

4.2 Threats to Validity

Pre-trends. I estimate event-study specifications interacting year dummies with the treatment indicator, omitting the last pre-treatment year (2022) as the reference. The pre-test for parallel trends yields $p = 0.39$.

Differential demand shocks. Tier 1 cities experienced distinct migration and demand patterns, particularly during COVID-19. I address this by controlling for log total consents (capturing local demand conditions) and by verifying that the placebo outcome — total consents — shows no differential treatment effect.

Canterbury earthquake. Christchurch’s post-earthquake reconstruction created unusual construction dynamics that may persist into the study period. I report results excluding Canterbury as a robustness check.

Attenuation from mixed regions. Each treated region contains both MDRS-affected Tier 1 territorial authorities and surrounding districts that remain under the standard consent regime. In Waikato, Hamilton City accounts for approximately 50% of regional consents; in Bay of Plenty, Tauranga accounts for roughly 60%. This measurement error attenuates the estimated effect toward zero. The true effect on Tier 1 cities would be larger in magnitude — meaning my null finding is, if anything, conservative. The 95% confidence interval on the CS estimator ($[-8.2, 4.6]$ pp) provides an upper bound even without correcting for this dilution.

Small samples. With 15 regions, cluster-robust inference may over-reject. The key results are precise nulls, so size distortion works against the finding (a too-liberal test would be more likely to reject the null). I report Callaway–Sant’Anna confidence bands, which do not rely on cluster asymptotics, as corroboration.

Minimum detectable effect. Given the pre-treatment standard deviation of the multi-unit share (12.4 pp) and 165 observations, the minimum detectable effect at 80% power and 5% significance is approximately 5.5 percentage points — roughly 18% of the pre-treatment treated-region mean. The study is therefore powered to detect policy-relevant compositional shifts, though it cannot rule out very small effects.

5. Results

5.1 Main Results

Table 2: Effect of MDRS on Dwelling Type Composition

| | (1) | (2) | (3) | (4) |
|---------------------|------------------|------------------|-------------------|---------------------|
| | Multi-unit share | CS ATT | log(Multi-unit) | |
| MDRS \times Post | 0.007 (0.036) | 0.023 (0.038) | -0.018 (0.033) | -0.400** (0.166) |
| Region FE | Yes | Yes | — | Yes |
| Year FE | Yes | Yes | — | Yes |
| log(Total consents) | No | Yes | No | No |
| Estimator | TWFE | TWFE | CS | TWFE |
| N | 165 | 165 | 165 | 165 |
| Treated regions | 4 | 4 | 4 | 4 |
| Control regions | 11 | 11 | 11 | 11 |
| Pre-treatment mean | 0.314 | 0.314 | 0.314 | — |

Notes: Standard errors clustered at the region level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Columns (1)–(2) report two-way fixed effects (TWFE) estimates. Column (3) reports the Callaway–Sant’Anna (2021) simple aggregated ATT using never-treated regions as the comparison group. Column (4) uses log multi-unit consents as outcome. Post period: year ended February 2023–2026 (MDRS operative August 2022).

Table 2 reports the main estimates. Column (1) shows the baseline TWFE specification: the MDRS effect on multi-unit share is 0.7 percentage points (SE = 3.6 pp), statistically indistinguishable from zero. Against a pre-treatment mean of 31.4% in treated regions, this represents a change of 2.2% — economically negligible. Column (2) adds log total consents

as a control for local construction activity; the point estimate rises slightly to 2.3 pp but remains insignificant.

Column (3) reports the Callaway–Sant’Anna simple aggregated ATT: -1.8 percentage points (SE = 3.3 pp, 95% CI: $[-8.2, 4.6]$). The confidence interval rules out effects larger than 4.6 pp — a shift that would have represented a 15% increase over the pre-treatment mean. This is a precise null.

Column (4) uses log multi-unit consents as the outcome and finds a significant negative effect of -0.40 (SE = 0.17, $p = 0.03$). This indicates that treated regions experienced a roughly 33% larger decline in absolute multi-unit consenting relative to controls during the post-2022 downturn — the opposite of what the policy intended.

5.2 Event Study

Table 3: Event Study: Multi-Unit Share by Year Relative to MDRS

| Year | Coefficient | SE |
|-----------------------|-------------|---------|
| <i>Pre-treatment</i> | | |
| 2016 | -0.058 | (0.045) |
| 2017 | -0.054 | (0.054) |
| 2018 | -0.040 | (0.039) |
| 2019 | -0.005 | (0.047) |
| 2020 | -0.016 | (0.036) |
| 2021 | -0.002 | (0.027) |
| 2022 (ref.) | 0.000 | — |
| <i>Post-treatment</i> | | |
| 2023 | 0.006 | (0.020) |
| 2024 | -0.032 | (0.046) |
| 2025 | -0.025 | (0.059) |
| 2026 | -0.020 | (0.043) |

Notes: Coefficients from regressing multi-unit share on year \times treated interactions with region and year fixed effects. Standard errors clustered at the region level. Reference year: 2022 (last pre-treatment year). MDRS operative August 2022.

Table 3 presents the event-study coefficients. No pre-treatment coefficient is individually significant, and the pattern shows no systematic divergence between treated and control regions before 2022. The post-treatment coefficients are also individually insignificant, with point estimates hovering around zero or slightly negative. The absence of any post-treatment break in trend is consistent with the null average effect.

5.3 Robustness

Table 4: Robustness: Effect of MDRS on Multi-Unit Share

| | (1) | (2) | (3) | (4) |
|--------------------|------------------|------------------|------------------|-------------------|
| | Baseline | Excl. Canterbury | Excl. small | Delayed (2024) |
| MDRS \times Post | 0.007 (0.036) | 0.013 (0.039) | 0.007 (0.041) | -0.004 (0.043) |
| N | 165 | 154 | 143 | 165 |

Notes: All specifications include region and year fixed effects with standard errors clustered at the region level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (2) excludes Canterbury (earthquake reconstruction). Column (3) excludes regions averaging fewer than 200 annual consents. Column (4) defines treatment onset as year ended February 2024 (first full post-MDRS year).

Table 4 reports four robustness checks. Excluding Canterbury (column 2) yields a point estimate of 1.3 pp (SE = 3.9 pp). Excluding small regions averaging fewer than 200 annual consents (column 3) gives 0.7 pp (SE = 4.1 pp). Using a delayed treatment definition starting in 2024 — the first full post-MDRS year — yields -0.4 pp (SE = 4.3 pp). All estimates are indistinguishable from zero.

5.4 Placebo Outcomes

Table 5: Placebo Outcomes: Total Consents and Houses Share

| | (1) | (2) |
|--------------------|---------------------|-------------------|
| | log(Total consents) | Houses share |
| MDRS \times Post | -0.183 (0.130) | -0.007 (0.036) |
| N | 165 | 165 |

Notes: Column (1) tests whether MDRS affected the overall volume of consents (it should not primarily shift composition, though supply effects are possible). Column (2) shows the houses share effect — the mechanical mirror of multi-unit share. Standard errors clustered at the region level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[Table 5](#) tests whether MDRS affected the overall *level* of construction activity. Column (1) shows that MDRS had no significant effect on the log total number of consents (-0.18 , $SE = 0.13$, $p = 0.18$), consistent with a null supply-volume response. This supports interpreting the composition result as a genuine null rather than a symptom of broader construction collapse. Column (2) reports the houses share coefficient as a mechanical consistency check — since $\text{houses share} = 1 - \text{multi-unit share}$, the coefficient is the exact negative of the baseline estimate, as expected.

6. Discussion

The central finding is that New Zealand’s most ambitious upzoning reform produced no detectable shift in the composition of new housing toward multi-unit dwelling types. This is not a power failure: the 95% confidence interval from the Callaway–Sant’Anna estimator rules out effects larger than 4.6 percentage points, or roughly 15% of the pre-treatment mean.

Three candidate mechanisms can explain the null. First, the *construction cost channel*: medium-density construction costs per square meter exceeded single-dwelling costs by 15–30% in New Zealand during 2022–2025 ([Ministry of Business, Innovation and Employment, 2023](#)), potentially eliminating the margin that planning deregulation was meant to free. Second, the *developer capacity channel*: the pool of developers experienced in medium-density

typologies remained small, and new entrants faced barriers in a downturn with tight financing (New Zealand Productivity Commission, 2012). Third, the *infrastructure channel*: upzoning without concurrent investment in three-waters infrastructure, transport, and open space may have limited the practical buildability of higher-density forms (Ministry for the Environment, 2020).

The significant negative effect on log multi-unit *levels* (column 4 of Table 2) is consistent with the first mechanism. If the construction downturn differentially affected multi-unit projects — which have longer planning horizons, higher capital requirements, and greater financing sensitivity — then treated regions would show a larger absolute decline even as the compositional share remained stable.

These findings complement the broader upzoning literature. Greenaway-McGrevy and Phillips (2023) show that Auckland’s Unitary Plan increased land prices at upzoned parcels without commensurate supply increases, consistent with upzoning transferring rents to landowners when supply constraints bind elsewhere. Been et al. (2023) find that Minneapolis’s 2018 upzoning produced modest effects on both total supply and demolitions. Anagol et al. (2024) estimate that Sao Paulo’s 2014 zoning reform increased housing supply by 1.9% — real but small. The common thread is that planning liberalization is necessary but not sufficient for the kinds of transformative compositional shifts that housing advocates envision.

For policymakers, the implication is clear: the “missing middle” gap cannot be closed by permitting reform alone. Complementary interventions — public investment in medium-density construction technology, development finance for smaller builders, infrastructure pre-investment, and potentially direct public procurement of medium-density housing — may be required to shift the composition of what actually gets built.

7. Conclusion

New Zealand’s MDRS was the sharpest test of the upzoning hypothesis available in any advanced economy: a nationwide mandate, applied simultaneously across major cities, eliminating a well-documented regulatory barrier. The finding that it produced no compositional shift toward multi-unit housing in four years suggests that the “missing middle” gap is structural, not regulatory. What a city *permits* and what a city *gets built* are separated by a margin that planning reform alone cannot bridge.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @olafdrw

First Contributor: <https://github.com/olafdrw>

References

- Anagol, Santosh, Fernando V. Ferreira, and Jonah M. Rexer**, “Estimating the Economic Impact of Zoning in São Paulo,” *American Economic Journal: Applied Economics*, 2024, *16* (3), 1–39.
- Been, Vicki, Ingrid Gould Ellen, and Katherine O’Regan**, “Supply Skepticism Revisited,” *Journal of Policy Analysis and Management*, 2023, *42* (4), 1047–1079.
- Callaway, Brantly and Pedro H.C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Glaeser, Edward L. and Joseph Gyourko**, “The Economic Implications of Housing Supply,” *Journal of Economic Perspectives*, 2018, *32* (1), 3–30.
- Greenaway-McGrevy, Ryan and Peter C.B. Phillips**, “The Effect of Upzoning on House Prices and Redevelopment Premiums in Auckland, New Zealand,” *Urban Studies*, 2023, *60* (1), 63–83.
- Kulish, Mariano, Anthony Richards, and Christian Gillitzer**, “Urban Structure and Housing Prices: Some Evidence from Australian Cities,” *Economic Record*, 2012, *88* (282), 303–322.
- Maltman, Karl and Ryan Greenaway-McGrevy**, “Medium Density Housing and the Missing Middle: Evidence from Lower Hutt, New Zealand,” *New Zealand Economic Papers*, 2024, *58* (2), 1–18.
- Ministry for the Environment**, “National Policy Statement on Urban Development 2020,” Technical Report, New Zealand Government, Wellington 2020.
- Ministry of Business, Innovation and Employment**, “Building Sector Trends,” Technical Report, MBIE, Wellington 2023.
- New Zealand Productivity Commission**, “Housing Affordability Inquiry,” Technical Report, New Zealand Productivity Commission, Wellington 2012.
- Parolek, Daniel G.**, *Missing Middle Housing: Thinking Big and Building Small to Respond to Today’s Housing Crisis*, Island Press, 2020.
- Reserve Bank of New Zealand**, “Financial Stability Report,” Technical Report, Reserve Bank of New Zealand, Wellington November 2023.

A. Data Appendix

The building consents data are drawn from Stats NZ’s monthly release “Building Consents Issued,” specifically the February 2026 edition. Table 4 of the release provides annual counts of new dwelling consents by region and dwelling type for years ended February 2016 through February 2026. Dwelling types are classified as “Houses” (standalone detached dwellings) and “Multi-unit homes” (townhouses, flats, units, apartments, and retirement village units). The data are publicly available at <https://www.stats.govt.nz/information-releases/building-consents-issued-february-2026/>.

Region-to-treatment mapping. Treated regions are those containing MDRS Tier 1 cities: Waikato (Hamilton City), Bay of Plenty (Tauranga City), Wellington (Wellington City, Lower Hutt, Upper Hutt, Porirua), and Canterbury (Christchurch City). Each treated region also contains smaller territorial authorities not subject to MDRS, introducing a degree of measurement error that attenuates the estimated treatment effect. Control regions are: Northland, Gisborne, Hawke’s Bay, Taranaki, Manawatū-Whanganui, Tasman, Nelson, Marlborough, West Coast, Otago, and Southland.

Timing convention. “Year ended February t ” includes consents issued from March $t - 1$ through February t . The MDRS became operative on August 20, 2022. Thus the first partially treated year is “year ended February 2023” (containing 6 post-treatment months: August 2022 through February 2023). The first fully treated year is “year ended February 2024.”

B. Identification Appendix

The parallel trends assumption is tested via an event-study specification interacting year dummies with the treatment indicator, with 2022 as the reference year. The joint F -test for pre-treatment coefficients yields $p = 0.39$, indicating no evidence of differential pre-trends. Individual pre-treatment coefficients are reported in [Table 3](#).

C. Robustness Appendix

See [Table 4](#) for results excluding Canterbury, excluding small regions, and using delayed treatment timing. All specifications yield point estimates within one standard error of the baseline null.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

| Outcome | Specification | $\hat{\beta}$ | SD(Y) | SDE | SE(SDE) | Classification |
|--|---|---------------|-----------|--------|---------|-------------------|
| <i>Panel A: Pooled</i> | | | | | | |
| Multi-unit share | TWFE baseline | 0.007 | 0.124 | 0.060 | 0.287 | Moderate positive |
| <i>Panel B: Heterogeneous (by region size)</i> | | | | | | |
| Multi-unit share | Large cities (Canterbury, Wellington) | -0.007 | 0.124 | -0.055 | 0.240 | Moderate negative |
| Multi-unit share | Smaller cities (Waikato, Bay of Plenty) | 0.022 | 0.124 | 0.174 | 0.387 | Large positive |

Notes: **Country:** New Zealand. **Research question:** Does mandatory upzoning (permitting as-of-right medium-density construction) shift the composition of new housing toward multi-unit dwelling types in treated cities? **Policy mechanism:** The Medium Density Residential Standards (MDRS), operative August 2022, require Tier 1 territorial authorities to allow up to three residential units of up to three stories on any urban residential lot without resource consent, eliminating a discretionary approval process that averaged 40–70 days and several thousand NZD in fees for medium-density proposals. **Outcome definition:** Multi-unit share — the fraction of annual new dwelling consents classified as multi-unit homes (townhouses, flats, apartments) rather than standalone houses, from Stats NZ Building Consents. **Treatment:** Binary (0/1): region contains a Tier 1 city subject to MDRS from August 2022. **Data:** Stats NZ Building Consents, annual year ended February, 2016–2026; 15 regions \times 11 years = 165 region-years. **Method:** Two-way fixed effects DiD with region and year fixed effects; standard errors clustered at the region level. Callaway–Sant’Anna (2021) estimator as robustness. **Sample:** 15 NZ regions excluding Auckland (treated earlier via Auckland Unitary Plan 2016). Four treated regions (Waikato, Bay of Plenty, Wellington, Canterbury) vs. 11 never-treated control regions. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the pre-treatment standard deviation of the multi-unit share. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).