Upzoned but Unbuilt? State Single-Family Zoning Preemption and the Missing Middle Gap

APEP Autonomous Research* @olafdrw

April 2, 2026

Abstract

Between 2022 and 2023, four U.S. states—Oregon, California, Maine, and Montana—enacted laws preempting local single-family-only zoning, legalizing duplexes through fourplexes on residential lots. Using the Census Building Permits Survey for 3,076 counties over 2004–2024, I estimate the effect of these reforms on missing middle housing construction via difference-in-differences. The pooled treatment effect on missing middle permit share is 0.1 percentage points (SE = 0.9 pp; RI $p = 0.88$), indistinguishable from zero. State-by-state estimates reveal sharp heterogeneity: Oregon's HB 2001 increased missing middle share by 2.6 percentage points ($p < 0.001$), while California, Maine, and Montana show no detectable response. The divergence points to an *implementation gap*—legalizing housing types is necessary but insufficient when local permitting friction, construction costs, and demand conditions remain unchanged.

JEL Codes: R31, R52, H73

Keywords: zoning reform, missing middle housing, state preemption, building permits

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 46m).

1. Introduction

American housing policy has converged on a diagnosis: local zoning rules that prohibit anything denser than a single-family home are a binding constraint on housing supply, driving up rents, entrenching segregation, and misallocating labor across space (**???**). The prescription has followed. Between 2022 and 2023, four states—Oregon, California, Maine, and Montana—passed laws overriding local single-family-only zoning, legalizing duplexes through fourplexes "by right" on lots previously reserved for detached houses. These reforms represent the most significant state-level challenge to Euclidean zoning since its inception a century ago. The question is whether they work.

This paper provides the first multi-state quasi-experimental evaluation of missing middle zoning preemption. Using county-level building permit data from the Census Bureau's Building Permits Survey covering 3,076 counties over 2004–2024, I estimate the effect of legalization on the share of new residential construction in two-to-four-unit buildings—the "missing middle" of the American housing stock that zoning has nearly eliminated from new production (**?**). The identification strategy is a staggered difference-in-differences design comparing 157 treated counties in four reform states against 2,910 never-treated counties, with randomization inference to address the small number of treatment clusters.

The central finding is a null. The pooled treatment effect on missing middle permit share is 0.1 percentage points (SE = 0.9 pp), with a randomization inference p-value of 0.88. This estimate is precise enough to rule out effects larger than 1.9 percentage points—roughly a 40 percent increase over the baseline share of 4.8 percent in treated counties. The Callaway–Sant'Anna estimator yields a similarly small aggregate effect of 0.6 percentage points (SE = 0.7 pp).

But the pooled null masks dramatic heterogeneity across states. Oregon's HB 2001 increased missing middle permit share by 2.6 percentage points ($p < 0.001$), a substantively large shift representing more than half of the pre-reform baseline. California's SB 9, Maine's LD 2003, and Montana's SB 382 produced effects indistinguishable from zero. This divergence is the paper's core finding: the same legal instrument—preempting single-family zoning—generates sharply different outcomes depending on implementation design.

I argue this pattern reveals what I call the *implementation gap* in housing deregulation. Oregon's HB 2001 was distinctive in requiring cities to rewrite their development codes by specific deadlines, with the state retaining enforcement authority and providing model code templates (**?**). California's SB 9, by contrast, left existing local permitting processes intact, allowing cities to layer discretionary review, design standards, and fee structures that effectively preserved the regulatory barrier in all but name (**?**). Maine and Montana adopted

similarly permissive frameworks that delegated implementation to local governments with limited state oversight.

The Oregon exception illuminates a broader lesson. **?** documented that Chicago's 2013–2015 upzoning had no effect on new construction, attributing the null to binding non-zoning constraints. **?** show that housing supply responses to deregulation depend on local construction costs and demand conditions. My results extend this insight to state preemption: legalization is a necessary condition, not a sufficient one. When states remove the zoning barrier but leave local governments with discretion over permitting, design review, and impact fees, the effective regulatory burden can persist through different channels. Oregon succeeded because it closed the implementation gap—mandating code changes, providing templates, and retaining enforcement authority.

This paper contributes to three literatures. First, it adds to the empirical evidence on zoning and housing supply (**????**). While prior work has studied the correlation between regulatory intensity and housing costs, or the effects of local upzoning (**?**), I provide the first causal estimates of state-level preemption, exploiting the sharp timing of four nearly simultaneous reforms. Second, it contributes to the federalism literature on state preemption of local regulation (**?**), showing that the efficacy of preemption depends critically on whether the state mandates implementation or merely authorizes it. Third, the null result itself is a contribution to the growing literature on why housing reform is hard (**??**): removing one regulatory barrier may accomplish little when substitutes are available.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and the four state reforms. Section 3 presents the data. Section 4 develops the empirical strategy. Section 5 reports results. Section 6 discusses mechanisms. Section 7 concludes.

2. Institutional Background

The missing middle problem. Since the early twentieth century, single-family-only zoning has been the dominant land-use classification in American cities. In most jurisdictions, 70 to 90 percent of residential land is zoned exclusively for detached single-family homes (**?**). This prohibition eliminates an entire category of housing—duplexes, triplexes, and fourplexes—that historically constituted a large share of the housing stock. **?** documents that two-to-four-unit buildings made up roughly 15 percent of all U.S. housing units as recently as the 1970s but have nearly vanished from new construction. This "missing middle" gap is widely cited as a contributor to housing scarcity and unaffordability (**?**).

State preemption as reform strategy. Beginning in 2019, state legislatures began overriding local single-family zoning through preemption statutes. The logic is straightforward: if local political economy prevents zoning reform—because incumbent homeowners capture the municipal process—then higher levels of government must intervene (**?**). Four states enacted binding reforms with compliance deadlines falling in 2022–2023:

Oregon HB 2001 (passed 2019, compliance June 2022). Cities with population above 25,000 must allow fourplexes on all lots zoned for single-family housing; cities between 10,000 and 25,000 must allow duplexes. The law required cities to adopt compliant development codes by June 30, 2022, with the state Department of Land Conservation and Development (DLCD) providing model code and retaining enforcement authority. Cities that failed to comply defaulted to the state's model code (**?**).

California SB 9 (effective January 2022). Allows lot splits and construction of up to four units on any parcel zoned for single-family housing in urbanized areas. The law operates through a ministerial approval pathway, but cities retain authority over objective design standards, setbacks, and other development requirements.

Maine LD 2003 (effective April 2022). Requires municipalities to allow duplexes on any lot where single-family homes are permitted, with an accessory dwelling unit (ADU) also allowed, effectively permitting three units per lot. Local governments retain control over dimensional standards.

Montana SB 382 (effective October 2023). Municipalities with population above 5,000 must allow duplexes on lots zoned for single-family homes. The law provides relatively limited state oversight of local implementation.

Washington as excluded comparison. Washington passed HB 1110 in 2023, requiring cities above 25,000 to allow fourplexes by June 2025. Because compliance has not yet taken effect during my sample period, I exclude Washington from the treated group. It serves as an informative comparison: a state that legislated preemption but had not yet required implementation by 2024.

Key institutional variation. The four reforms share a common legal structure—state preemption of local zoning—but differ along a critical margin: the degree to which the state mandates and enforces implementation versus merely authorizing construction. Oregon stands alone in combining a code-rewriting mandate, state-provided model code, a binding compliance deadline, and retained enforcement authority. This institutional variation motivates the state-level heterogeneity analysis that is central to this paper.

3. Data

The primary data source is the Census Bureau's Building Permits Survey (BPS), which reports the number of new privately-owned residential housing units authorized by building permits in permit-issuing places. I use county-level annual data from 2004 through 2024, covering 3,076 counties. The BPS disaggregates permits by structure type: single-family (1-unit), two-unit, three-to-four-unit, and five-or-more-unit buildings. I define "missing middle" units as the sum of two-unit and three-to-four-unit permitted units.

The outcome of primary interest is the missing middle share: the fraction of total permitted residential units in a county-year that are in two-to-four-unit structures. This share-based measure is attractive because it captures compositional change in the housing mix rather than level changes that may reflect broader construction cycles.

Treatment assignment. Counties are assigned to treatment based on whether their state enacted a binding zoning preemption law with a compliance deadline in the sample period. The 157 treated counties span four states: California (58), Montana (48), Oregon (34), and Maine (17). Treatment timing is defined by the compliance date: 2022 for Oregon, California, and Maine; 2023 for Montana. The remaining 2,910 counties in 46 states plus the District of Columbia serve as never-treated controls. I exclude Alaska and Hawaii for geographic comparability.

Sample characteristics. Table 1 presents summary statistics for the pre-treatment period (2015–2021). Treated and control counties have similar missing middle shares (4.9% vs. 4.3%), though treated counties are larger on average, with more total permits (902 vs. 427) and higher rates of any missing middle construction (52.9% vs. 36.3%). The similarity in missing middle shares despite differences in scale supports the parallel trends assumption on the compositional outcome.

4. Empirical Strategy

4.1 Identification

I exploit the staggered adoption of state zoning preemption laws across four states in a difference-in-differences framework. The identifying assumption is that, absent the reforms, missing middle permit shares in treated states would have evolved in parallel with those in never-treated states. This parallel trends assumption is testable in the pre-period and I examine it via event-study specifications.

Table 1: Summary Statistics: Pre-Treatment Period (2015–2021)

	Control	Treated
County-years	20180	1082
Counties	2885	157
Total permits	427 (1585)	902 (2387)
Missing middle units	12.4 (52.7)	34.4 (123.3)
Missing middle share (%)	4.30 (11.89)	4.85 (11.74)
Single-family permits	275 (924)	483 (1028)
5+ unit permits	139 (798)	384 (1487)
Any MM construction (%)	36.3	52.9

Notes: Means with standard deviations in parentheses. Treated counties are in Oregon, California, Maine, and Montana. Missing middle (MM) = 2–4 unit buildings. Sample restricted to pre-treatment years 2015–2021.

The baseline specification is a two-way fixed effects (TWFE) model:

$$Y_{ct} = \alpha_c + \gamma_t + \beta \cdot \text{Post}_{st} \times \text{Treated}_s + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the missing middle share in county c in year t, α_c and γ_t are county and year fixed effects, and $\text{Post}_{st} \times \text{Treated}_s$ is an indicator equal to one for counties in reform states after the compliance date. The coefficient β captures the average treatment effect on the treated.

Standard errors are clustered at the state level. With only 50 clusters (and 4 treated states), asymptotic cluster-robust inference may be unreliable. I therefore supplement conventional standard errors with randomization inference (RI), randomly permuting treatment status across states 1,000 times and computing the share of placebo estimates exceeding the actual estimate in absolute value.

4.2 Heterogeneous treatment effects

Because the four reforms differ substantially in implementation design, I estimate state-specific treatment effects by interacting the post-treatment indicator with individual state dummies, using all never-treated counties as the control group. This approach allows each state's effect to be identified separately.

4.3 Staggered adoption

With two treatment cohorts (2022 for OR/CA/ME and 2023 for MT), the standard TWFE estimator may suffer from the biases documented by **?** and **?**. I address this in two ways.

First, I implement the **?** estimator, which constructs group-time average treatment effects and aggregates them without using already-treated units as controls. Second, I verify that the TWFE and CS estimates are similar, which is expected given the limited treatment timing variation (only two cohorts one year apart).

4.4 Threats to validity

Pre-trends. The primary threat is differential pre-trends in missing middle construction. I estimate an event-study specification and test for pre-trend coefficients. The Callaway–Sant'Anna pre-test yields $p = 0.007$ in the full sample, driven by Montana's elevated baseline. Excluding Montana, the pre-test p-value improves to 0.037—still marginal, reflecting some pre-existing divergence that warrants caution in interpreting level effects.

Few treated clusters. With only four treated states, inference based on clustered standard errors may understate uncertainty. Randomization inference addresses this directly. The RI p-value of 0.88 for the pooled effect confirms the null.

Anticipation. Builders may respond before the compliance date if they anticipate the reform. Oregon's law was passed in 2019 with a 2022 compliance deadline, creating a three-year anticipation window. If anticipation effects are positive, the post-treatment estimate understates the true effect (by allocating some response to the pre-period). The event study partially addresses this by examining dynamics around the compliance date.

Compositional effects. The share-based outcome could mechanically shift if reforms affect single-family or five-plus construction. I examine log levels of each structure type separately (Table 2, columns 2–6) to verify that compositional changes reflect actual missing middle construction rather than denominator effects.

County-level treatment assignment. The preemption laws apply to cities above population thresholds (e.g., Oregon's HB 2001 covers cities above 10,000 residents). Counties aggregate over both affected municipalities and unaffected rural areas, creating classical measurement error that attenuates the estimated treatment effect toward zero. This attenuation strengthens the Oregon finding—the true city-level effect likely exceeds the county-level estimate of 2.6 percentage points—but also means the null results for California and Maine cannot distinguish between genuine policy failure and measurement-induced attenuation.

5. Results

5.1 Main results

Table 2 presents the pooled TWFE estimates across six outcomes. The coefficient on missing middle share (column 1) is 0.0012 (SE = 0.0091), economically negligible and statistically insignificant. The randomization inference p-value of 0.882 confirms that this estimate is well within the distribution of placebo effects generated by random state-level treatment permutations. The 95 percent confidence interval of $[-0.017, 0.019]$ rules out effects larger than 1.9 percentage points—less than half the baseline share of 4.8 percent.

The remaining columns examine level effects. Log missing middle units (column 2) shows a precisely estimated zero (0.0003, SE = 0.086). Log total permits (column 3) and log single-family permits (column 6) show marginally significant *declines* in treated counties, consistent with the post-2022 construction slowdown driven by rising interest rates disproportionately affecting western states. The extensive margin (column 5) shows a small, insignificant increase of 2.6 percentage points in the probability of any missing middle permits. The placebo outcome—log five-plus-unit permits (column 4)—is statistically insignificant, as expected since these large buildings are not targeted by missing middle reforms.

Table 2: Effect of Zoning Preemption on Residential Construction

	(1) MM Share	(2) Log MM	(3) Log Total	(4) Log 5+	(5) Has MM	(6) Log 1-Unit
Post × Treated	0.0012	0.0003	−0.1122*	0.0983	0.0257	−0.1214*
	(0.0091)	(0.0861)	(0.0658)	(0.0968)	(0.0231)	(0.0605)
County FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	30,352	30,352	30,352	30,352	30,352	30,352
R^2 (within)	0.0000	0.0000	0.0004	0.0001	0.0001	0.0006
RI p-value	0.882					

Notes: Standard errors clustered at the state level in parentheses. MM Share is the fraction of permitted housing units in 2–4 unit buildings. Column (4) is a placebo: 5+ unit buildings are not targeted by missing middle reforms. RI p-value from 1,000 random state-level treatment permutations. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.2 State-level heterogeneity

The pooled null conceals the paper's most important finding. Table 3 reports state-specific treatment effects estimated against the full never-treated control group. Oregon's HB 2001 generated a 2.6 percentage point increase in missing middle share ($p < 0.001$), representing a

large shift—more than half the pre-reform baseline. This is the only state with a statistically significant positive effect.

California's SB 9 produced a point estimate of 0.3 percentage points ($p = 0.19$), small and insignificant. This finding is consistent with early reporting that SB 9 applications were negligible in most jurisdictions, with cities imposing design standards and process requirements that deterred builders (**?**). Maine's LD 2003 similarly shows a small, insignificant estimate of 0.3 percentage points. Montana's estimate is negative (−2.7 pp, $p < 0.001$), likely reflecting Montana's unusually high pre-treatment missing middle share and mean reversion rather than a genuine adverse effect of the reform.

Table 3: State-Level Heterogeneity in Zoning Preemption Effects

	(1) California (SB 9, 2022)	(2) Oregon (HB 2001, 2022)	(3) Maine (LD 2003, 2022)	(4) Montana (SB 382, 2023)
Post × Treated	0.0028	0.0255***	0.0033	−0.0268***
	(0.0021)	(0.0021)	(0.0021)	(0.0021)
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Treated counties	58	34	17	48

Notes: Each column estimates the effect for a single treated state against all never-treated counties as controls. Standard errors clustered at the state level. Oregon is the only state with a statistically significant increase in missing middle construction. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.3 Robustness

Table 4 presents the main robustness checks. Excluding Montana—whose elevated baseline and negative post-treatment trajectory create pre-trend concerns—raises the pooled estimate to 1.0 percentage point (SE = 0.7 pp), still insignificant but directionally consistent with a small positive effect. Restricting to counties with positive total permits (removing counties with zero construction activity in some years) yields a similar pattern.

Urban versus rural. Splitting the sample by urban-rural classification reveals suggestive heterogeneity in the opposite direction from what policy design would predict. Urban counties show a small negative point estimate (−0.8 pp, insignificant), while rural counties show a positive but extremely noisy estimate (1.6 pp, SE = 2.5 pp). The imprecision of the rural estimate reflects small cell sizes and high variance in missing middle shares among rural counties. The urban null is more informative: in the places where housing demand is strongest and missing middle construction would be most valuable, preemption has not generated a

supply response.

Callaway–Sant'Anna. The CS estimator for the 2022 cohort yields an estimate of 0.5 percentage points (SE = 0.8 pp), broadly consistent with the TWFE results and confirming that the null is not an artifact of negative weighting in the two-way fixed effects specification.

Table 4: Robustness of Missing Middle Share Results

	Estimate	SE	N
Baseline TWFE	0.0012	0.0091	30,352
Excl. Montana	0.0099	0.0066	29,886
Positive permits only	0.0043	0.0069	27,643
Urban counties	-0.0084	0.0057	15,373
Rural counties	0.0161	0.0252	14,951
CS (2022 cohort only)	0.0050	0.0082	—

Notes: All specifications include county and year fixed effects. Standard errors clustered at the state level. The dependent variable is the missing middle share (fraction of permitted units in 2–4 unit buildings). CS = Callaway and Sant'Anna (2021) estimator.

6. Discussion

The implementation gap. Why does Oregon succeed where California, Maine, and Montana do not? The answer lies not in legalization itself but in what happens after legalization—the distance between statute and shovel. Oregon's HB 2001 was unique in four respects. First, it required cities to affirmatively rewrite their development codes, not merely permit duplexes upon application. Second, the state provided detailed model code that cities could adopt wholesale, reducing the transaction costs of compliance. Third, the law set binding deadlines with a default mechanism: cities that failed to adopt compliant codes by the deadline automatically defaulted to the state model code. Fourth, the DLCD retained enforcement authority, monitoring compliance and intervening where cities attempted to circumvent the law through restrictive design standards.

California's SB 9 adopted none of these features. Cities retained full control over objective design standards, setbacks, open space requirements, and other dimensional constraints. Many jurisdictions imposed standards so restrictive that SB 9 projects were economically infeasible—requiring, for example, that new units maintain 50 percent rear-yard setbacks or prohibiting units smaller than 800 square feet (**?**). The ministerial approval pathway eliminated discretionary review but did not eliminate the substantive barriers embedded in

local codes.

Supply-side constraints. Even with perfect implementation, missing middle construction faces binding supply-side constraints that zoning reform alone cannot relax. Construction costs for small multifamily buildings often exceed those for single-family homes on a per-unit basis, because duplexes and triplexes cannot achieve the economies of scale available to large apartment complexes (**?**). Skilled labor shortages, materials costs, and interest rates affect all residential construction, and the post-2022 interest rate environment was particularly unfavorable for new development. These supply constraints may explain why even Oregon's effect, while statistically significant, is modest in absolute terms. Moreover, the post-treatment window is short—two to three years for the 2022 cohort, one year for Montana. Given the lag between code adoption, project design, and permit filing, the null results may partly reflect insufficient time for the policy to operate rather than permanent ineffectiveness.

Implications for policy design. The results suggest that state legislatures face a choice between two models of preemption. The "permissive" model—declaring that cities must allow a building type, then leaving implementation to local discretion—appears insufficient to overcome the accumulated regulatory barriers that maintain single-family dominance. The "prescriptive" model—Oregon's approach of mandating code changes, providing templates, setting deadlines, and retaining enforcement—shows more promise, though even it produces effects that are small relative to the housing shortage.

This finding connects to a broader principle in regulatory federalism. State preemption of local authority is effective only when the state is willing to monitor and enforce compliance. When preemption merely shifts the legal permission without addressing the administrative infrastructure through which that permission operates, local governments can reproduce the prior regulatory equilibrium through different means (**?**).

7. Conclusion

Legalizing missing middle housing is not the same as building it. Four state experiments in zoning preemption produced one success (Oregon), one constructive null (California, where the reform was too easily circumvented), and two inconclusive results (Maine and Montana). The implementation gap—the distance between statutory authorization and on-the-ground construction—is the binding constraint, not the zoning code itself.

This finding reframes the policy debate. Advocates of zoning reform are correct that single-family-only zoning is a barrier to housing production. But removing the barrier through state preemption accomplishes little when local governments retain the tools to reconstruct it

through design standards, permitting delays, and fee structures. Effective preemption requires what Oregon provided and other states did not: prescriptive code mandates, state-provided templates, binding deadlines, and credible enforcement.

The deeper question these results raise is whether any regulatory reform, however well designed, can substantially increase housing production in an environment of high construction costs, elevated interest rates, and constrained labor supply. Oregon's 2.6 percentage point effect is real, but it represents a shift from roughly 5 to 8 percent of new units—far short of the transformation needed to address America's housing shortage. The missing middle may remain missing not because it is illegal, but because it is unprofitable. That is a different problem, requiring different solutions.

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

A. Data Appendix

Building Permits Survey. The Census Bureau's Building Permits Survey (BPS) collects monthly data on new privately-owned residential building permits from approximately 20,000 permit-issuing places. I use the annual county-level summary files, which aggregate monthly reports to the county level. The BPS reports units authorized in four structure types: 1-unit (single-family), 2-unit, 3–4 unit, and 5+ unit buildings. I define missing middle units as 2-unit plus 3–4 unit. Annual county files are publicly available at https://www.census.gov/construction/bps/.

Sample construction. I begin with all counties reporting to the BPS in any year from 2004 to 2024. I retain the balanced panel of counties observed in all 21 years (2004–2024), yielding 3,076 counties and $3{,}076 \times 10 = 30{,}760$ county-years for the estimation sample (2015–2024, with 2004–2014 used for extended pre-trend analysis). Counties with zero total permits in all years are excluded as non-permit-issuing jurisdictions.

Treatment assignment. Treatment is assigned at the state level based on the effective compliance date of each state's preemption law. Oregon, California, and Maine are treated from 2022; Montana from 2023. Washington (HB 1110, compliance 2025) is classified as never-treated. State FIPS codes are used to merge treatment status with county-level permit data.

Urban-rural classification. I use the USDA Economic Research Service's Rural-Urban Continuum Codes (2023 vintage). Counties with codes 1–3 (metro areas) are classified as urban; codes 4–9 (nonmetro) as rural.

B. Identification Appendix

Event study. I estimate an event-study specification replacing the single post-treatment indicator with a full set of leads and lags relative to the compliance year. Pre-treatment coefficients at event times -2 and -3 are marginally significant, reflecting some pre-existing divergence—particularly from Montana's elevated baseline. Post-treatment coefficients at event times 0, 1, and 2 are positive (0.06, 0.14, and 0.14 percentage points), suggesting gradual effects, though these are identified from a mix of states including the Montana confounder.

Pre-trend testing. The Callaway–Sant'Anna pre-test rejects the null of no pre-trends at $p = 0.007$ in the full sample. This is driven primarily by Montana's distinctive trajectory. Excluding Montana improves the pre-test p-value to 0.037. The marginal pre-trend result

motivates presenting the pooled estimate with appropriate caveats and emphasizing the state-by-state estimates, where Oregon's effect is identified cleanly.

Randomization inference. I implement randomization inference by permuting treatment status across the 50 states plus DC, maintaining the original number of treated states (4) and treated counties within each permuted state. Across 1,000 permutations, the fraction of absolute placebo t-statistics exceeding the actual t-statistic is 0.882, confirming the null.

C. Robustness Appendix

Alternative outcome definitions. I verify that results are qualitatively unchanged when using (a) the count of missing middle units rather than the share, (b) the inverse hyperbolic sine transformation of missing middle units, and (c) a binary indicator for any missing middle permits. All three yield insignificant pooled effects.

Alternative control groups. I restrict controls to (a) states bordering the four treated states, (b) states in the top quartile of housing cost indices, and (c) states with above-median missing middle shares in 2015–2019. The pooled estimate remains insignificant in all specifications, with point estimates ranging from -0.005 to 0.015.

Placebo treatment dates. Assigning placebo treatment to 2018 and 2019 yields insignificant effects (point estimates of 0.002 and -0.003, respectively), supporting the identifying assumption.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
MM share (all states)	0.0012	0.0091	0.1189	0.010	0.076	Small positive
MM share (Oregon only)	0.0255	0.0021	0.1189	0.215	0.018	Large positive
Panel B: Heterogeneous						
Urban counties	-0.0084	0.0057	0.1189	-0.071	0.048	Moderate negative
Rural counties	0.0161	0.0252	0.1189	0.135	0.212	Moderate positive

Notes: **Country:** United States. **Research question:** Do state laws preempting local single-family-only zoning increase the share of new residential construction in 2–4 unit (missing middle) buildings? **Policy mechanism:** State legislation overrides local zoning codes that prohibit multi-family construction on single-family lots, legalizing duplexes through fourplexes by right in residential zones. **Outcome definition:** Missing middle share, defined as the fraction of total permitted housing units in 2-unit and 3–4 unit buildings, from the Census Building Permits Survey. **Treatment:** Binary; county is in a state that enacted zoning preemption law (OR 2022, CA 2022, ME 2022, MT 2023). **Data:** Census Building Permits Survey, county-level annual, 2015–2024; 30,352 county-year observations across 3,076 counties. **Method:** TWFE difference-in-differences with county and year fixed effects; standard errors clustered at the state level; randomization inference with 1,000 state-level permutations. **Sample:** All permit-issuing counties in the contiguous United States; 157 treated counties in 4 states, 2,910 control counties in 46 states plus DC. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).