Second-Generation Rent Control and New Housing Construction: Evidence from Oregon's SB 608

William B. Thorne
The Graduate Center, City University of New York

September 2026

Job market paper

Abstract

This paper covers Oregon's recent passage of SB 608 in 2019, a softer modern rent control that caps the annual rent increases at seven percent on top of an allowance for inflation, makes an exception for new builds which remain free from rent control for their first fifteen years, and allows rents to adjust to the market rate between tenancies. I use parcel-year assessor records on both the Oregon and Washington sides of their shared border in a difference-in-differences design, and I find no aggregate decline in new construction. To find a proxy for rental construction, I use new builds on the parcels held by absentee owners; I estimate these fell by roughly 1.2 per 1,000 parcels relative to those on the Washington side of the border. However, these estimates are attenuated by the existence of only one clean post-treatment year in these data and a contaminating pre-treatment differential trend.

Chapter 1 of the author's dissertation, Second Generation Rent Control and New Construction, The Graduate Center, City University of New York, 2026. Page numbers below are the chapter's own.

1 Introduction

High housing costs have been an issue of public concern, and rent control has often been used to address housing costs for tenants. Rent control policies can be broadly classified into two main categories: first-generation rent control and second-generation rent control (Arnott 1995). There are a variety of differing features between the two generations as they are commonly understood; the new-construction exemption is not explicitly tied to either of these regime generation definitions (which are not thoroughly defined in the literature). By virtue of being a softer feature, it has become associated with the more lenient second-generation policies, but is present in numerous first-generation policies. The lack of an exemption for new construction may lead to a depression in the construction of new housing, and thus undercut the intended effect of the policy. First-generation rent control is quite well-represented in the literature, but analysis of newer second-generation rent control is more sparse, and the main effects are still not well understood (Kholodilin 2024). This paper presents evidence of the effects of second-generation rent control on new construction.

In February of 2019, Oregon passed SB 608, a rapid policy shift from a statewide preemption of rent control to implementing a second-generation rent control law. SB 608 crucially does exempt new construction from the rent control, and rather than setting strict rent caps, allows for some increases in rent provided they are within state limits, which is also common among second-generation rent control policies. I use the data from housing parcels on either side of the long Oregon-Washington state border to estimate the effect of SB 608 on new housing construction. I estimate the effect on new construction using a difference-in-differences design with housing in counties adjacent to the border to control for spatial variation in housing markets. By comparing units close to the border, I can be more confident that these housing units are likely to face otherwise similar market conditions, and be more likely to only differ by treatment of rent control.

New construction is exempt from the cap for 15 years, so none of the units observed here faces the regulation directly. What the design measures instead is whether the arrival of rent control makes builders and landlords less willing to enter the rental market. Oregon had banned rent control and then showed that the rules could change, and the rental margin is where this forward-looking response would appear.

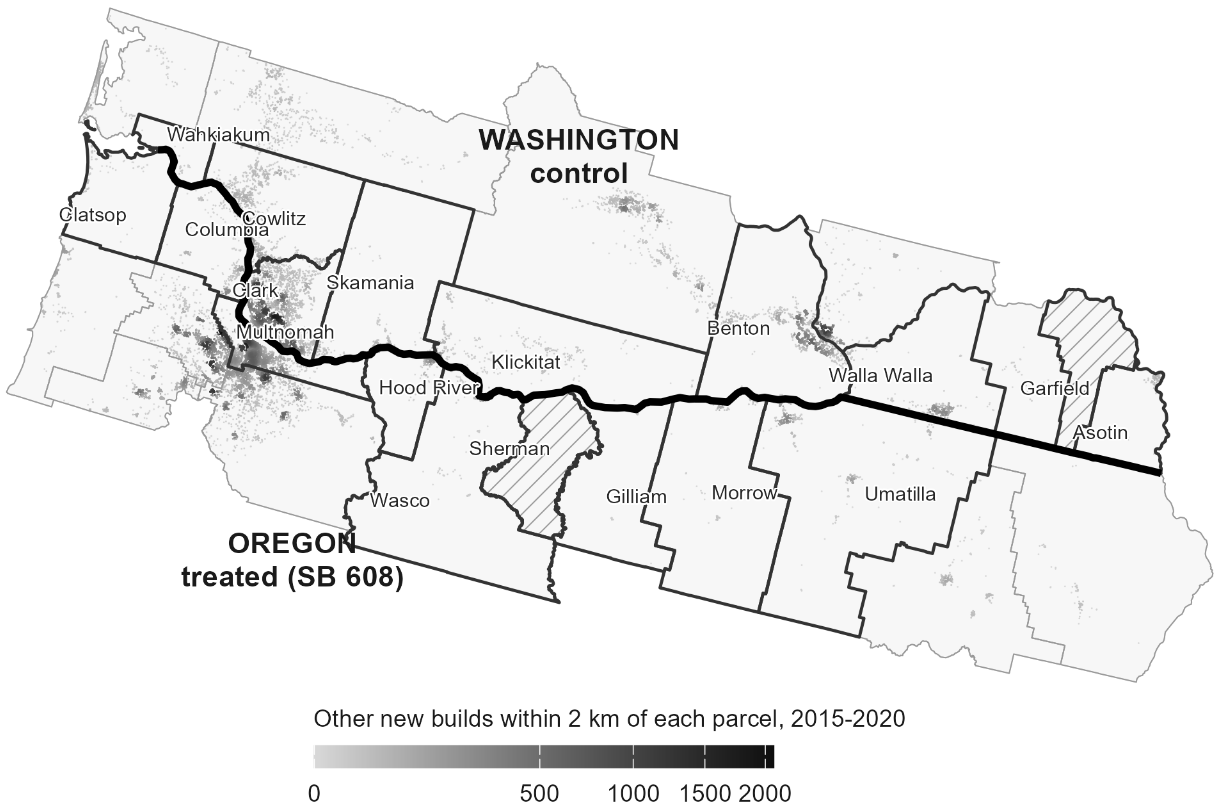

Figure 1: New construction and the study area, 2015–2020.
Notes: Each point marks a parcel with a recorded new build. Darker shading means that more builds lie within two kilometers of that parcel, and the count stops at the state border. The outlined counties form the estimation sample; the two hatched counties were excluded beforehand because they lack a live assessor feed. The map locates the observations rather than showing the policy comparison itself. Washington's border counties build at roughly twice Oregon's per-parcel rate throughout the pre-period, for reasons that predate SB 608. The design therefore asks whether that gap changes after the law, which is what the tables and event studies estimate.

I find no evidence of an aggregate decline in new housing construction. I find that there is a detectable decline on the rental proxy, new builds on parcels whose owner does not live at the address listed for the property, though with one clean post-treatment year this decline cannot be fully separated from a pre-existing differential trend. Oregon's rate of new builds on rental-side parcels fell relative to Washington: 1.24 fewer parcels each year had new rental-side builds per 1,000 parcels. The size of the effect of second-generation rent control is likely a function of the difference between uncontrolled and controlled sector prices, the elasticity of new housing supply, and the length of the new construction exemption, although future research is necessary to understand these effects. To this point, no study in the literature has measured the new construction effect of a rent control regime with the same set of design features. Among those that have measured the new construction effect of similar rent control regimes (new construction exempt and full vacancy decontrol), all examined municipal regimes in California (e.g., Murray et al. 1991; Goetz 1995).

2 Institutional background and policy timing

Oregon, like many states in the United States, maintained a rent control preemption law prior to the passing of SB 608. These preemption laws were passed to prevent cities within these states from passing rent control laws. These preemption laws coincided with a broader academic consensus that rent control had a broadly negative effect on housing supply (Arnott 1995), although the majority of the evidence on rent control emerged from studies of first-generation control, which was the predominant form in the first half of the twentieth century. Many cities across the United States have seen rising rents in recent decades, including the Portland metro area, which straddles the Oregon and Washington border, which led to the political environment where SB 608 was passed. SB 608 restricts rents from rising by more than seven percent plus the consumer price index, with the primary exception being any unit that is less than 15 years old. SB 608 also includes a variety of structures to protect tenants from eviction without a qualifying reason; qualifying reasons include owner-occupancy, demolition, conversion to a condo, or repairs and renovations. Tenants are also entitled to fees or payments to assist in relocation if a landlord does evict them.

Given the passage of the law in February of 2019, and the delay in construction due to the logistical and regulatory impediments of new construction, I treat 2019 as a transition year: the main estimates pool it with the pre-period, and the sensitivity model for this choice gives it its own term. This addresses a threat to the design: housing planned prior to SB 608 may have completed in 2019, and the short interval between the law's passage in February 2019 and year-end makes it unlikely that new construction in 2019 reflects a response to SB 608. Using the same logic, some 2020 completions are construction decisions that were made before the law. This pushes the aggregate estimate toward zero, but the rental proxy I estimate is less exposed, because the rent-or-sell decision is observed in the 2021 vintage, after passage. The main estimates therefore compare 2020 with 2015 through 2019.

3 Data

The underlying data for this analysis are parcel records by year; each parcel record has building attributes for that year. These records include the year built, the precise location of the parcel, and a swath of building characteristics. I have these data from 2010–2021, covering every county in Oregon and Washington. Pre-2015 assessor records show evidence of bulk data refreshes in both states, with measured new-build rates that swing by a factor of two or more between adjacent years, and are therefore excluded from the analysis. The data produced by county assessors are irregular, including recording lags and backfilling, but the design still attributes these units to their correct build year in data assembly (the outcome keys on each parcel's final recorded build year, so late-arriving records land in their true year). The private parcel-level data technically end in 2021; however, because of regular recording delays on the part of Oregon assessors, 2020 is the last confidently comparable year from these data. Including 2021 would look like a drastic Oregon building decline, when it is really just a recording lag. I therefore settle on 2015–2020.

The data used for the primary estimates are all parcels in Oregon or Washington state that are in a live-feed county adjacent to the border, with no further distance restriction (Figure 1). The panel formed from border counties has 553,563 parcels observed from 2015 to 2020; the rental-proxy rows use 310,024 of these from outside Multnomah County. Oregon parcels from 2020 onwards are considered "treated" in the main analysis.

The main outcome I track is the recording of a new build in the assessor records. This is found by the records of "Year Built" for each parcel in each assessor year. Due to lags in assessor record keeping, I record a parcel as a new build in a given year (and zero otherwise) whenever the parcel's final recorded build year matches that year. The square-footage row (measured in thousands of square feet) of Table 1 of new residential building per parcel, estimated on Columbia, Hood River and Multnomah against Clark, is the second outcome. It was estimated using this geography and window because county and regional records extend past the 2021 snapshot and continue as normal until they reach the still-provisional 2024 and 2025 records, leaving a 2015–2023 window. I code a given parcel as a rental proxy in cases where the parcel is non-owner-occupied, which is ultimately any time the owner mailing address does not coincide with the property address. Tenure is measured from the owner-occupier data in the 2021 vintage of the private parcel-level dataset.

4 Research design

The advantage of the parcel-level data, rather than broader aggregates at the state or county level, is that it restricts the comparison to units that share in many unobserved characteristics and experience otherwise similar market conditions. I estimate a difference-in-differences of parcels in counties adjacent to the Oregon-Washington border. In my main specification, I specifically estimate the differential change in the probability of new construction on Oregon parcels after SB 608, compared to those in Washington. This estimate does not necessarily imply that parcels beyond this sample would have the same magnitude of response to SB 608. For this design to be meaningful, I must assume that housing markets in these border counties experience similar economic pressures and shocks, and thus move in similar ways aside from the SB 608 policy change. Including all the available counties is a trade-off between sample size and comparability across a border. Dube et al. (2010) suggest it is preferable to stay close to the border to ensure comparisons are more similar.

The main specification places no distance restriction on parcels within the border counties. For robustness, each parcel has a calculated distance in kilometers from the border, d_i, and I re-estimate restricting the sample to those where $d_i \leq B$ for several values of B. The border-county sample keeps the basic cross-border comparison while leaving enough counties for the cluster-based inference used in the paper, but this trade-off between comparability and sample size is always present. The pre-treatment period includes all years 2019 and earlier, and the post-treatment year is 2020.

The main unit of observation is a parcel-year, meaning each parcel is represented each year, and identification of the estimated effects stems from changes in these counties over time. The included fixed effects are county α_c, and year γ_t. This accounts for time-invariant heterogeneity for each county, and the broader macroeconomic shocks in each year. Therefore the main difference-in-differences design can be written as: $NewBuild_{it} = \beta(Oregon_i \times Post2020_t) + \alpha_c + \gamma_t + \varepsilon_{it}$, with clustering at the county level to account for county assessor behavior and other local market conditions.

As discussed in Conley and Taber (2011), I cannot afford to cluster at the level of treatment, the state level, and without significant numbers of other controls I am left with the next best alternative of clustering at the county level. To widen the control group significantly, even if practically possible, would also sacrifice crucial comparability when it comes to housing markets. The headline estimates use 12–14 clusters, which places them near the lower edge of the range discussed by MacKinnon and Webb (2017) and Roodman et al. (2019), and so I also give cluster-robust standard errors together with wild-cluster-bootstrap *p*-values using Webb weights (Cameron et al. 2008; MacKinnon and Webb 2017; Roodman et al. 2019; Webb 2023). Randomization inference is also reported alongside as a different check on the same finding: for the headline samples the treated label is reassigned to every same-sized subset of the sample's counties (7 of 14, 6 of 13, or 6 of 12). After the randomized treatment is assigned, the difference-in-differences coefficient is computed under each newly randomized assignment. The reported p-value is then the two-sided share of placebo coefficients which are at least as large in magnitude as the one that was actually observed, with the null being "no effect in any county." Young (2019) finds that randomization tests overturn a meaningful proportion of conventionally significant joint results, but the headline estimate here remains detectable under county-level exchangeability, not at the level of policy assignment.

The model including 2019 explicitly as a transition year can be written: $NewBuild_{it} = \beta_{19}(Oregon_i \times 2019_t) + \beta_{20}(Oregon_i \times Post2020_t) + \alpha_c + \gamma_t + \varepsilon_{it}$. In either case, β represents a change in the probability of new construction on a parcel, rather than a number of housing units, and Table 1 reports 1,000β, the change in new builds per 1,000 parcels. The event studies reported in Figure 3 replace the two interaction terms with one Oregon-by-year term for every year except for an omitted reference year. Each coefficient then is that year's Oregon-minus-Washington gap relative to the omitted reference year (which is 2018 in the main series).

5 Results

Table 1 shows no clear aggregate rise or fall in housing construction, and its measurement is not very precise. Before 2020 the border counties recorded 6.2 new builds per 1,000 parcels a year in Oregon and 13.7 in Washington, and then 1.4 and 2.8 on the Multnomah-excluded rental proxy (Figure 2). The multifamily margin is relatively underpowered, as we can see in Table 1's CRVE p-values. On the rental proxy there is a higher level of precision, and I find a rental-proxy decline in Oregon relative to Washington state across multiple specifications and robustness checks: -1.24 builds per 1,000 parcels on the border counties (cluster-robust p .005, bootstrap .013, randomization .043) and -1.46 on the clean panel (randomization p .002).

Table 1: Main difference-in-differences estimates per 1,000 parcels, Oregon to Washington, with county and year fixed effects; samples span 2015–2020 except as stated.

			size .136 / .079 / .043			
Margin	Sample	Estimate	CRVE p	WCB p	RI p	Clusters
All new construction	Border counties	-0.99	.278	.644	.694	14
All new construction	Clean panel (12 live counties)	-0.76	.794	.837	.868	12
Rental proxy	Border counties	-1.24	.005	.013	.043	13
Rental proxy (robustness)	Border counties + Multnomah	-1.47	<.001	.014	.021	14
Rental proxy	Clean panel (12 live counties)	-1.46	.011	.008	.002	12
Multifamily buildings	Border counties	0.02	.574	–	.802	14
Multifamily buildings	Clean panel (12 live counties)	-0.14	.098	.065	.050	12
New residential sq ft (000s) per parcel	4-county segment, 2015-2023	-1.27	.778	.666	.750	4
Permits, 5+ unit buildings (asinh)	Border counties, 2015-2025	-0.53	.320	.330	.268	15

Notes: A rental proxy is a new build on a parcel coded absentee in the 2021 records. The Border counties samples place no distance restriction; the Clean panel (12 live counties) samples restrict to parcels within 5 kilometers of the border in twelve live-feed border counties. The robustness row adds Multnomah County with recovered build dates; it survives dropping all Portland-jurisdiction parcels, and the estimate survives dropping any single county. The three p-value columns are cluster-robust (CRVE), wild-cluster bootstrap (WCB), and exact randomization inference (RI); the cluster count is the number of border counties used for inference, and is one lower on the border-county rental-proxy row because Multnomah County is held out of the tenure split. The four-county segment is Columbia, Hood River and Multnomah against Clark; the permit series covers fifteen of the nineteen total counties adjacent to the border (Gilliam and Wasco lack pre-period permits), which is a different set of counties than the parcel sample. For clarity: the size line above each p-value column reports that test's measured clean-panel rejection rate at a nominal 5 percent under 5,000 simulated null draws. The permits row is from the Census Building Permits Survey; asinh is the inverse hyperbolic sine. The four-county segment allows only four randomization assignments, so its smallest attainable randomization p is .25.

I do not measure with enough precision in the matched-sample decomposition, the identifiable split of clean-panel builds into rental-proxy builds and its complement, the owner-occupied or unknown tenure builds. Which is to say: I cannot detect whether there is a reallocation from renting to owning or a net reduction in building. An additional vulnerability is that a parcel that builds units intending them for rent in an earlier year and converts those units to be owner-occupied (selling them as condos or single-family homes) by the time the 2021 snapshot is observed will not be accurately reflected by this measure, a channel consistent with Diamond et al. (2019).

Tenure measurement tends to be more difficult in the case of rental housing for a practical reason: while ownership is clearly reported for each parcel, its actual use is not. Additionally, there are ambiguous cases: the owner of a parcel with a duplex who rents out one half might not show up in these data; roughly one in five duplex parcels is coded owner-occupied. I do not know the use of the other half of the duplex without assumption. Neither this measure nor any other I estimate is sufficiently powered to say whether this means there was an actual drop in the share of parcels that built for rent, or not.

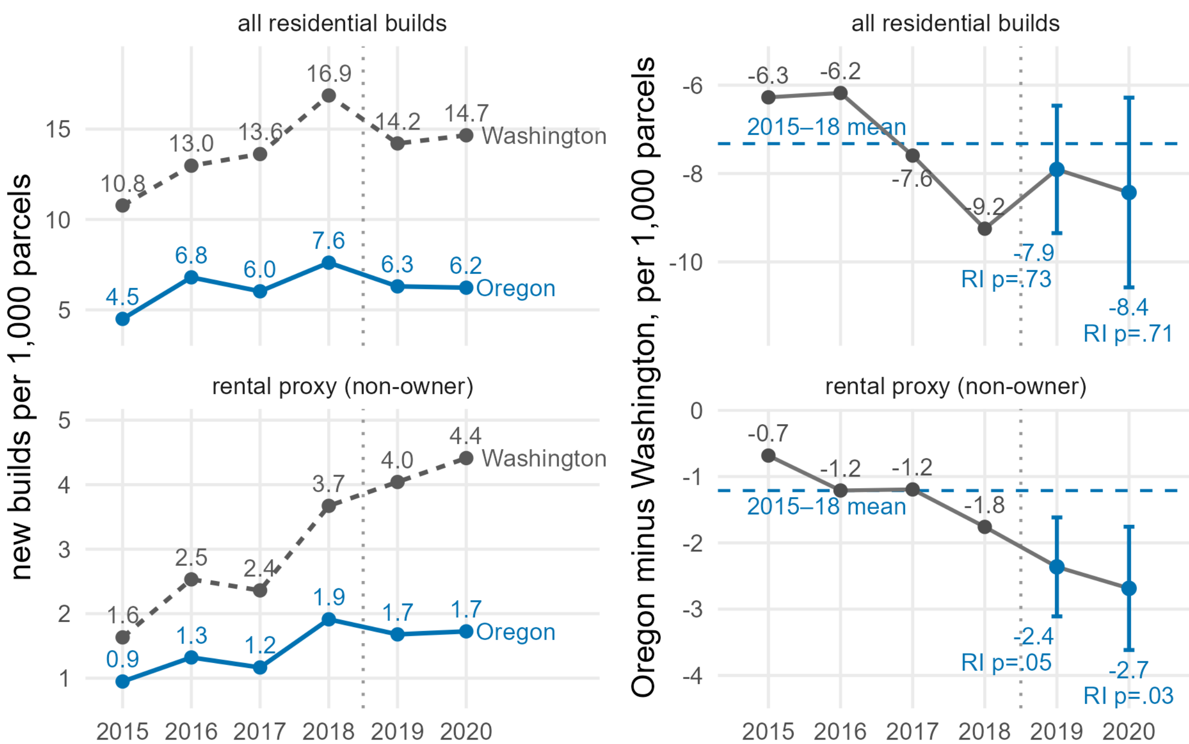

Figure 2: Main-specification event study, 2015–2020.

Notes: Each point is an Oregon-by-year coefficient in builds per 1,000 parcels, measured relative to 2018, which is shown by the open point. The specifications include county and year fixed effects; the bars are cluster-robust confidence intervals, and the labels report exact randomization-inference p-values. The dotted line separates the 2019 transition year and the 2020 post-treatment year from the pre-period. Pre-2018 coefficients are positive in both panels: the joint pre-trend test rejects, while no single pre-year is distinguishable from zero under randomization inference; see the reference-year discussion in the text.

I find that there is evidence of a pre-trend in these data (Figures 2 and 3), and the test that detects it has sufficient power: on the rental proxy the three-period model gives -1.15 builds per 1,000 parcels for 2019 (cluster-robust p .006, bootstrap .033, randomization .046) beside -1.47 for 2020 (p .005, .013, .034), and the 2020 gap over 2019, -0.32, is not distinguishable from zero (p .138, .066, .319). A conservative reading of this evidence could be that there is not enough evidence to concretely determine that the rental decline observed is truly due to rent control. Using an alternative set of reference years in Figure 3, the rental-proxy decline persists while the aggregate series is null throughout the post period. The 2018-referenced aggregate pattern is being produced by a 2018 dip instead of a post-2019 change. The results are not particularly sensitive to treating 2019 as a transition year.

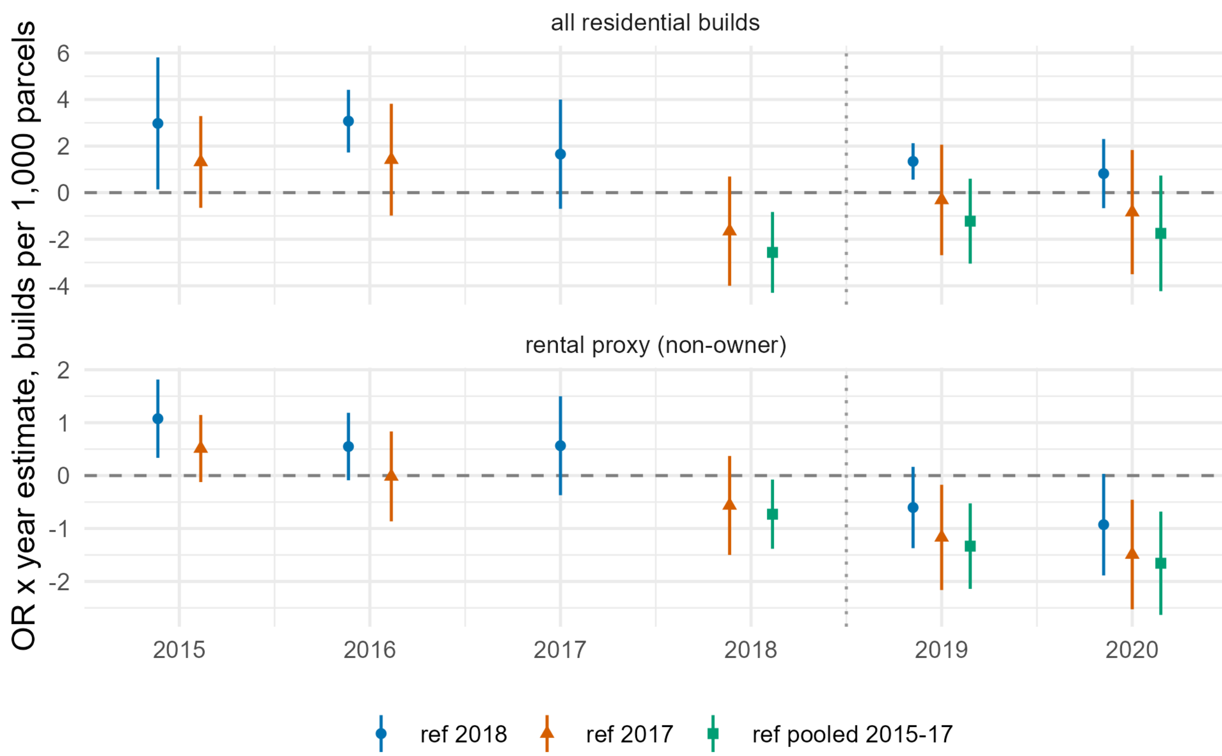

Figure 3: Reference-year sensitivity of the event study. Under the 2018 reference, the joint test that “the 2015 to 2017 coefficients are zero” is rejected with cluster-robust inference.
Notes: Each series re-estimates the event study with a different omitted reference: 2018, 2017, or the pooled 2015–2017 mean. Under the alternative references the aggregate series is null throughout the post period, identifying the 2018-referenced pattern as an artifact of a 2018 dip; the rental proxy declines in the post period under all three. Its apparently flat 2015–2017 pre-period is an artifact of pooling those years into the baseline: resolved individually the pre-period is not flat, and under the 2018 reference it moves at a rate comparable to the post-period. No single pre-year is distinguishable from zero using randomization inference. Bars are 95 percent cluster-robust intervals.

The headline rental-proxy estimate is relatively small and should be read in light of the fact that aggregate construction could not be measured to fall. Therefore this is in line with an interpretation that there was no clear disincentive to build housing as a consequence of rent control.

There are two events on the Washington side of the comparison which fall inside the relevant time window and bias the estimates in opposite directions. Washington state had SB 5334, which intended to encourage condominium production, in 2019, and in 2020 halted non-essential construction for four to five weeks with no corresponding stoppage for construction in Oregon: the first of these potentially raises Washington state's build rate after 2019 and as a consequence pushes the aggregate and owner-side estimates toward a decline. On the other hand, the second would lower Washington's 2020 rate and would push every 2020 difference toward a possible increase.

6 Robustness and validation

6.1 Bandwidth and spatial architecture

It might be intuitive that the effects of rent control depend on distance from the border, and thus I constructed several distance-based bandwidths. Figure 4 reports the estimates for bands running from one to fifty kilometers, followed by the border counties without a distance cutoff and the widest sample these parcel data allow. That last sample is not statewide: the parcel store covers thirty of the region's seventy-five counties, twenty-four of which have live assessor feeds. The aggregate estimate is not statistically distinguishable from zero in any of these samples. The rental-proxy estimate remains negative throughout, and expanding beyond the border counties does little to change it. Therefore the main specification disregards this bandwidth cutoff in favor of a larger sample size. At a one-kilometer band the sample is so small that there is a meaningful lack of precision in the estimate, but as bandwidths get larger, precision improves markedly and the estimates remain quite stable. For robustness, dropping each county one by one leaves the rental-proxy result negative and statistically detectable throughout, though its magnitude varies. The Census permit series is the check of units (rather than builds on parcels) that the parcel data cannot adequately supply. The effect on units is null: the permits for buildings of five or more units in the border counties only move by -0.53 in inverse hyperbolic sine units, with p between .27 and .33 on all three tests (Table 1).

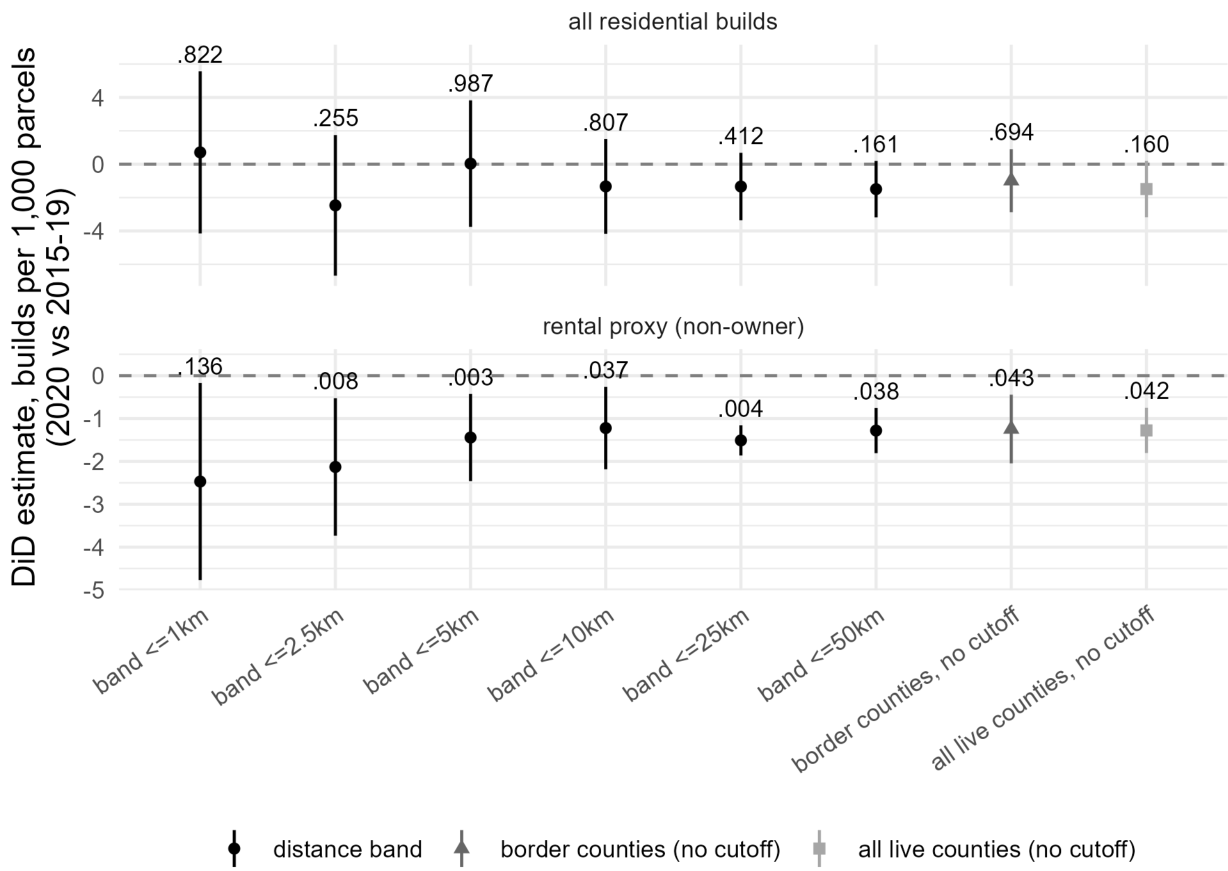

Figure 4: Bandwidth sensitivity of the border comparison, 2015–2020.
Notes: Points are difference-in-differences estimates in builds per 1,000 parcels for all residential builds (top panel) and the rental proxy (bottom panel), across distance bands from 1 to 50 kilometers, the border-county sample with no cutoff, and the widest sample these data support. That widest series is not statewide: the parcel store covers 30 of the region's 75 counties, 24 with live assessor feeds. The aggregate estimate is not distinguishable from zero in any sample. The rental-proxy estimate is negative in every sample and, under cluster-robust inference, distinguishable from zero at every bandwidth; randomization inference, reported alongside, distinguishes it from 2.5 kilometers outward. Widening beyond the border counties leaves it essentially unchanged.

This design is vulnerable to an issue raised in Dube et al. (2010): cross-border comparisons which do not account for local economic conditions can produce spurious negative effects due to spatial heterogeneity. When adopting their proposed solution, border-segment-by-year fixed effects, I then only compare parcels within the same 25-kilometer stretch of the border; I find that every segment-year specification remains negative and distinguishable from zero under cluster-robust inference, but is not adjudicable, because randomization inference does not retain enough valid assignments there to give power.

6.2 The 2021 data hole and the Multnomah-in specifications

The 2021 private parcel-level data have a hole, a result of the delay in Oregon assessors publishing their records: when the 2021 file was compiled, very few Oregon records for that year were yet present. This motivated the assembly of additional data directly from county assessors to patch this hole and recover post-period effects in new construction (unfortunately, no owner-occupier data were included in this patch). Upon repair, 2021 no longer shows the sharp Oregon decline the unrepaired data imply, and 2022–2024 extend the new construction post-period, but do not alter the aggregate null. This repair is limited to six counties, four in Oregon and two in Washington, of which only Multnomah and Clark appear in the border sample. I estimated construction of new single-family buildings and multifamily buildings, with data exclusively from assessors (2008–2024) and, under randomization inference, could not detect any divergence post rent control, though the comparison is not powered to detect one. This is consistent with neither multifamily nor single-family noticeably declining, and paired with the rental-proxy results leans my interpretation towards shifts within occupancy, and not housing construction decline. All Multnomah-in rental-proxy specifications I estimate reject under the bootstrap, but one (the clean panel with each parcel's tenure code read from the parcel-data vintage that follows its build year, and where that code is blank, from the next vintage that has one) does not clear the randomization inference threshold (randomization p .086 against bootstrap p .016), and therefore I cannot claim a detection on that specification.

7 Mechanisms and interpretation

The 15-year exemption in the rent control policy was designed to allow developers of new rental housing to still be able to profit during that window. Oregon asking rents grew less quickly than

Washington rents, but with some separation starting in 2018, and only a small gap emerging in the years after rent control (Figure 5); this should not be taken as clear evidence of the effect of rent controls, although these data are not well suited to answer this question. The one rent measure that does reject is statewide rents with the Seattle counties excluded (-6.5%, bootstrap p .032). This still sits on a gap that was already opening from before 2019, and so the same "anticipation or trend" problem as the rental proxy appears again.

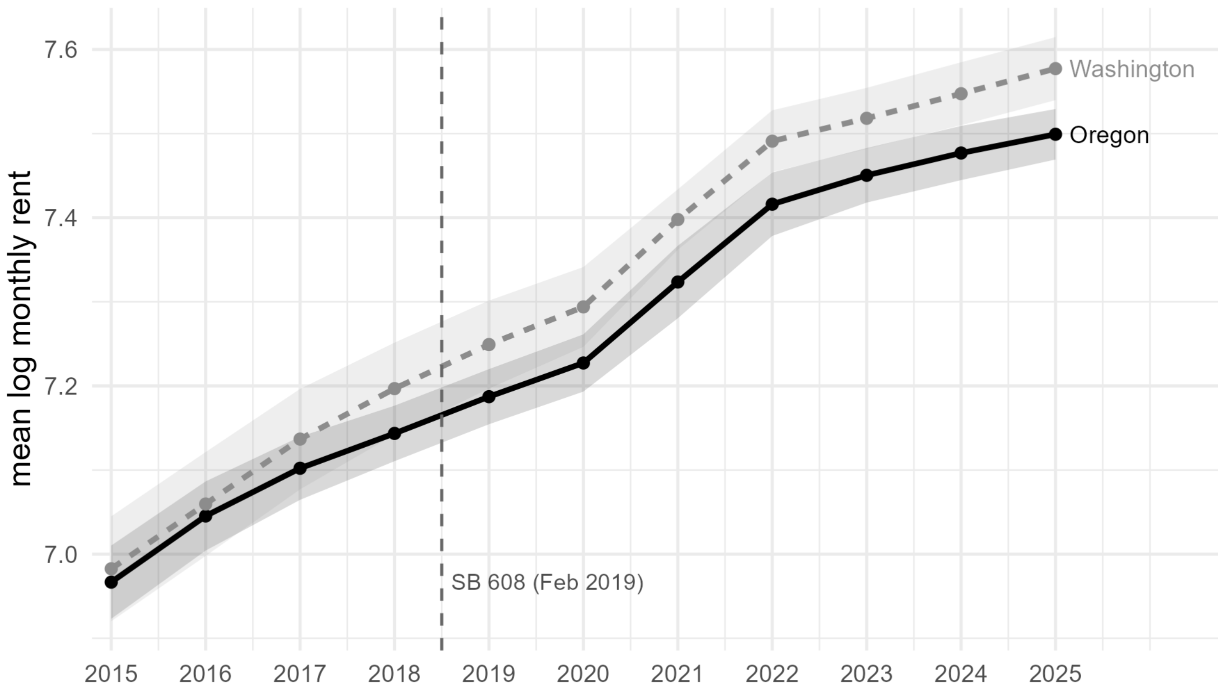

Figure 5: Market asking rents, Oregon versus Washington.
Notes: County means of log rent from Zillow's index of market asking rents, for the balanced panel of 8 Oregon and 9 Washington counties, 2015–2025; shading is one across-county standard error and the dashed line marks SB 608. Oregon rents grew slightly slower than Washington rents after 2019, and the statewide difference is not distinguishable from zero under randomization inference.

Conversions per one thousand at-risk rentals were higher in Washington state than Oregon in the pre-period, and there was a marked decline from 2018 to 2019 in both states; the positive estimate in Table 2, 28.3 per 1,000 rentals, is representing Washington's faster fall than Oregon. It is quite likely that this is an artifact of the data or a change in recording procedure. The rent-bifurcation rows of Table 2 test the exemption's potential consequence directly: the gap between rents for new and old units does not open from 2019 onwards, at -0.1 percent without controls and +1.3 percent with hedonic controls, but neither is distinguishable from zero. Given the loose cap form of this rent

Table 2: Mechanism-margin estimates.

Channel	Measure	Estimate	p-value	Primary test	RI p
Market rents (net)	ZORI, OR vs WA statewide, 2015–2025	-4.0%	.210	WCB (Webb)	.211
Market rents (net)	ZORI, drop Seattle	-6.5%	.032	WCB (Webb)	.040
Rent bifurcation	PUMS DDD, built 2010+ vs pre-2000	-0.1%	.976	SDR replicates	–
Rent bifurcation	PUMS DDD, with hedonic controls	1.3%	.561	SDR replicates	–
Stock conversions	Absentee-to-owner rate per 1,000 rentals	28.3	.020	WCB (Webb)	.067
Stock conversions	Stock absentee share	-0.012	.274	WCB (Webb)	.333

Notes: Market-rent rows are Oregon-versus-Washington differences in log asking rent, in percent; rent-bifurcation rows are triple-difference estimates of the new-versus-old rent gap from Census rental microdata, with variance from the Census replicate-weight method; conversion rows are difference-in-differences on tenure transitions in the existing stock, from assessor occupancy codes. The reported p-value is the primary one for each row, from the wild-cluster bootstrap where a cluster bootstrap applies and from replicate weights for the microdata rows; the final column reports exact randomization inference alongside, which is not computed for the microdata design. The positive conversion-rate estimate reflects recorded conversions collapsing in both states in 2019 and rebounding in 2020, twice as deeply in Washington; a move that synchronized reads as a records disturbance rather than Oregon behavior. All three families switch treatment on in 2019, the year SB 608 passed, unlike Table 1's 2020. The rent rows are estimated from county-year panels of 17 counties, the bifurcation rows on household-level Census microdata, and the conversion rows on county-year panels of six counties (restricted to just border counties). SDR replicates are the Census successive difference replication weights.

control, this is to be expected in the short run: only after years of uncontrolled rents rising faster than seven percent plus inflation would the difference between controlled and uncontrolled rents start to build to such an extent that it would be measurable with coarse data and statistical tools.

8 Conclusion

This paper addresses a gap in the current literature on the effect of rent control with a CPI-plus cap form, full vacancy decontrol and a new construction exemption on new housing construction and on rental construction. With regard to new housing construction, there is no evidence that this rent control policy reduces the total amount of new housing construction, and in the case of rental construction in particular there is evidence of a decline concentrated on the rental proxy, though with one clean post-treatment year it cannot be fully separated from a pre-existing differential trend.

What this could mean is that while the rental proxy declined, there was a rise in owner-occupied construction, which would be consistent with common null or positive findings for new construction overall in the second and third generation literature (as covered in Chapter 2), though these data cannot distinguish this from the alternatives, as I discussed earlier under Results. Another effect worth considering is that of the elasticity of supply for housing construction, which may be affected by any number of building restrictions commonly present in American cities.

What is yet to be understood and remains unmeasured in this paper is whether the implementation of even a very soft rent control policy such as SB 608 might still be interpreted by landlords and developers as a shift towards a more risky rental market environment, which dissuades new rental construction.

Future work on soft rent control policies, with detailed microdata that measures the bite of the rent control policy, the supply elasticity of new housing construction, and the effect of the exemption for new construction in rent control policies, would be able to answer some of the unresolved questions that relate to this topic. Additionally, it is not clear whether the relative rental-proxy decline would just be a temporary short-run phenomenon or would permanently shift construction away from the rental sector.

References

Arnott, Richard. 1995. "Time for Revisionism on Rent Control?" *Journal of Economic Perspectives* 9 (1): 99–120. https://doi.org/10.1257/jep.9.1.99.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller. 2008. "Bootstrap-Based Improvements For Inference With Clustered Errors." *Review of Economics and Statistics* 90 (3): 414–27. https://doi.org/10.1162/rest.90.3.414.

Conley, Timothy G., and Christopher R. Taber. 2011. "Inference with 'Difference in Differences' with a Small Number of Policy Changes." *Review of Economics and Statistics* 93 (1): 113–25. https://doi.org/10.1162/REST_a_00049.

Diamond, Rebecca, Tim McQuade, and Franklin Qian. 2019. "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco." *American Economic Review* 109 (9): 3365–94. https://doi.org/10.1257/aer.20181289.

Dube, Arindrajit, T William Lester, and Michael Reich. 2010. "Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties." *Review of Economics and Statistics* 92 (4): 945–64. https://doi.org/10.1162/REST_a_00039.

Goetz, Edward G. 1995. "A Little Pregnant: The Impact of Rent Control in San Francisco." *Urban Affairs Quarterly* 30 (4): 604–12. https://doi.org/10.1177/107808749503000408.

Kholodilin, Konstantin A. 2024. "Rent Control Effects Through the Lens of Empirical Research: An Almost Complete Review of the Literature." *Journal of Housing Economics* 63 (March): 101983. https://doi.org/10.1016/j.jhe.2024.101983.

MacKinnon, James G., and Matthew D. Webb. 2017. "Wild Bootstrap Inference for Wildly Different Cluster Sizes." *Journal of Applied Econometrics* 32 (2): 233–54. https://doi.org/10.1002/jae.2508.

Murray, Michael P., C. Peter Rydell, C. Lance Barnett, Carol E. Hillestad, and Kevin Neels. 1991. "Analyzing Rent Control: The Case of Los Angeles." *Economic Inquiry* 29 (4): 601–25. https://doi.org/10.1111/j.1465-7295.1991.tb00850.x.

Roodman, David, Morten Ørregaard Nielsen, James G. MacKinnon, and Matthew D. Webb. 2019. "Fast and Wild: Bootstrap Inference in Stata Using Boottest." *The Stata Journal: Promoting Communications on Statistics and Stata* 19 (1): 4–60. https://doi.org/10.1177/1536867X19830877.

Webb, Matthew D. 2023. "Reworking Wild Bootstrap-Based Inference for Clustered Errors." *Canadian Journal of Economics/Revue Canadienne d'économique* 56 (3): 839–58. https://doi.org/10.1111/caje.12661.

Young, Alwyn. 2019. "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results*." *The Quarterly Journal of Economics* 134 (2): 557–98. https://doi.org/10.1093/qje/qjy029.