THE UNIVERSITY OF CHICAGO

# The Effect of Inclusionary Zoning on the Housing Supply in U.S. Cities: A Bunching Analysis

By

Zachary Goldstein

August 2023

A paper submitted in partial fulfillment of the requirements for the Master of Arts degree in the Master of Arts Program in the Social Sciences

Faculty Advisor: Pascal Noel

Preceptor: Marshall Jean

# Abstract

Inclusionary Zoning (IZ) is a policy used in hundreds of U.S. cities and counties in which developers are required to set aside a fraction of units in new residential buildings for low-income tenants and charge below-market rent. A concern about IZ is that it decreases the housing supply by making it less profitable to build. I investigate the regulatory burden IZ imposes on developers by exploiting a common feature of IZ policies, that they only take effect above a sharp threshold of unit count, such as 20+ unit buildings. I study the extent to which developers avoid having to comply with IZ by strategically constructing apartment buildings with unit counts right below these thresholds. I find moderate IZ-driven bunching below the thresholds in some jurisdictions, but in most cities and counties studied there is no bunching under the thresholds.

# Acknowledgements

Throughout my work on my MA thesis, I have been very fortunate to have the support of many faculty members and fellow students at the University of Chicago.

I would like to thank my advisor Pascal Noel for providing friendly encouragement and helping me find sensible answers to challenging questions.

I am grateful to my preceptor Marshall Jean for giving supportive and constructive feedback from beginning to end.

I thank Yanyan Shang for guiding me through the quantitative methods program over the past year.

I appreciate having had the opportunity to chat with professors like Damon Jones and Emily Talen about various aspects of my research, from the urban policy side to the econometrics of bunching estimation.

I am thankful to my classmates in the thesis workshop for helping me clarify my ideas and how best to communicate them.

I am especially grateful to Otis Reid, who very helpfully suggested an interesting research question of real-world importance that inspired the work in this paper.

Lastly, I would like to thank my family and friends for supporting and encouraging me during my time in MAPSS.

# Introduction

“Inclusionary zoning” refers to a policy requiring developers of new apartment buildings in a city to set aside a fraction of units for low-income renters, and to set rents for those units low enough that they are affordable to those renters. Often, inclusionary zoning policies apply to buildings above a certain size threshold, such as 25 units. The details of inclusionary zoning policies, such as that unit count threshold, the percent of affordable units required, and the degree of affordability, vary widely across the cities and counties that have them. Also, in some programs, developers can satisfy requirements by paying a fee or building affordable units at another site (Wang and Fu 2022).

Proponents of inclusionary zoning believe it is a way for the government to induce developers to build affordable housing without having to spend taxpayer dollars. But inclusionary zoning has also faced criticism that it reduces the supply of housing by making it unprofitable to build apartment buildings, since landlords are not able to charge as much in rent (Hamilton 2021).

Affordable housing in general is important as an economic issue due to the large component of household budgets spent on housing, as well as the high number of rent-burdened households in the United States. Inclusionary zoning in particular is important to study because of its potential as a solution, and because of the current widespread occurrence of inclusionary zoning policies across hundreds of American cities and counties (Wang and Balachandran 2021).

Studying the effects of inclusionary zoning on the supply of market-rate housing units is even important to the issue of affordable housing because of filtering, the process by which initially expensive housing gradually becomes more affordable over time. A recent study found that building a new market-rate apartment building for 100 residents causes 45-70 people to move out of lower income neighborhoods, which reduces demand pressure and increases affordability (Mast 2021).

My research question concerns evaluating a specific possible consequence of inclusionary zoning laws. Do they cause developers to respond by building apartment buildings with fewer housing units than they would otherwise? We might expect this to occur since inclusionary zoning regulations often only apply to buildings above a particular sharp threshold of number of units per building, for example apartment buildings with at least 10 units. Therefore, developers may respond to inclusionary zoning regulations by choosing to build smaller apartment buildings, especially right below the minimum threshold at which the regulations take effect. If this is the case, we should observe "bunching" in the distribution of apartment unit counts right below the threshold.

An important aspect of this research question is what policy regime is considered to be the imagined counterfactual. Sometimes, inclusionary zoning laws actually liberalize zoning laws compared to the pre-IZ zoning code by creating a "density bonus" - legalizing higher density than what was legal before as long as IZ requirements are met. Because of this, some may argue that inclusionary zoning actually increases the housing supply and leads developers to build apartment buildings with more housing units than they would otherwise. However, my framework for the research question is not going to focus on whether IZ reduces housing supply compared to the prior legal environment, but rather compared to the counterfactual in which developers could build buildings with the same number of units that are legal to build under IZ, but without any affordability requirements.

There are a few reasons for this choice. Firstly, a "density bonus" is really two policies packaged together. It means simultaneously raising the maximum number of units a developer can legally build in an apartment building while also imposing an affordability requirement or fee on the newly-legalized larger buildings. The first part, allowing developers to build at higher densities, is an important policy to study in its own right, but the affordability requirement is the distinctive feature of inclusionary zoning policies. It is useful to decouple these effects because they do not necessarily have to accompany each other in policy design. Sometimes cities legalize the construction of buildings with more units without an affordability requirement, and sometimes cities impose an affordability requirement on already-legal building types. Secondly, the effects of IZ compared to the pre-IZ policy regime are highly contingent on whatever the pre-IZ policy regime is, so lessons learned from answering such a research question using data from some cities are not as applicable to other contexts. For example, a change-over-time-based methodology might find an overall positive effect of inclusionary zoning on the housing supply due to raising the maximum unit count even if the affordability requirement itself had a negative effect. Thirdly, many studies using methods like difference-in-differences or fixed effects models that study change-over-time have already been conducted. It is a more unique contribution to the literature to focus on a counterfactual that resembles a jurisdiction's actual policy in every way except the affordability requirements.

This paper contributes to a literature on evaluating the effects of inclusionary zoning policies in United States municipalities and counties. One takeaway from previous studies is an inconsistency in findings about the effect of inclusionary zoning on the total supply of housing. An Urban Institute report attempting to synthesize the research concluded that "overall, evidence that IZ laws negatively affect private market prices and development is mixed. The type of impact these laws have appears dependent on the design of the policy, the neighborhood location, and the housing market in the area" (Ramakrishnan, Treskon, & Greene, 2019, p. 5).

For example, one study found a "marginally significant increase in multifamily housing starts" in California cities that adopted inclusionary zoning (Bento & Lowe, 2009). A study focused on Los Angeles and Orange County California found no statistically-significant adverse effect on housing supply (Mukhija et al., 2010). Similarly, a study of Greater Washington D.C. inclusionary zoning programs found no evidence that inclusionary zoning reduces the housing supply (Hamilton, 2021). A study of inclusionary zoning in various jurisdictions in suburban Boston found decreased production of housing there, but the same study also examined policies in the San Francisco metropolitan area and found no statistically-significant effect on the quantity of new housing development there (Schuetz, Meltzer, & Been, 2011). Other researchers also studied inclusionary zoning in the California jurisdictions and found that inclusionary zoning decreased new housing production there (Powell & Stringham 2004, Means & Stringham 2012).

A recent paper studied mandatory inclusionary housing in London, specifically evaluating a policy change in which the minimum threshold was lowered from fifteen to ten unit buildings. They found evidence of "strategic behavior", specifically a reduction in 10-14 unit buildings and an increase in projects with nine or fewer units. When these effects were aggregated, they found no net loss of new housing units (Li & Guo, 2022).

Another recent paper is notable for moving beyond a limited geographic area of case studies to evaluate inclusionary zoning in cities across the United States. They focus on the outcome of direct affordable unit production in compliant buildings, and use a correlational analysis to study which attributes of policies are associated with high affordable unit production. They find that "jurisdictions with policies that were mandatory, older, covered the entire jurisdiction, or had more complex income requirements designed to reach lower income levels had significantly higher production of affordable units" (Wang & Fu, 2022, p. 550).

These studies generally rely on some kind of pre-post analysis using panel data to identify the effects of the policies. Often this is enhanced with difference-in-differences, controlling for observed covariates in multiple regression analyses, and spatial/temporal fixed effects models (Bento & Lowe, 2009; Schuetz, Meltzer, & Been, 2011). These are sensible techniques for inferring causality given the infeasibility of conducting a randomized experiment, and offer some useful information. However, there are some limitations to these methodologies that point to the need for the alternative econometric approach that this paper takes.

As with many studies using these methods across applied microeconomics, there are concerns about unmeasured confounding and the unreliability of the parallel trends assumption. In the case of inclusionary zoning, we may be concerned that the implementation of inclusionary zoning is caused by some feature of a city's housing market that also affects the housing supply but is difficult to measure. We may also be concerned that inclusionary zoning may be associated with other affordable housing policy changes occurring at similar times, which could obfuscate the effects of any single policy change.

In terms of external validity, a weakness of much of the literature is that it narrowly focuses on a few case study regions. The outcomes in those regions may or may not generalize to inclusionary zoning elsewhere due to differences in policy design and the underlying housing market conditions (Basolo 2004).

# Data

I use the Inclusionary Housing Database to obtain information about inclusionary zoning programs across the United States, including when they began and what kinds of thresholds they include (Grounded Solutions Network, 2020). The biggest source of data on housing I use is the CoreLogic property database, which includes millions of properties across the U.S. The main information I use from this source is the city/county, year built, and number of units.

I supplement this with state and local open data for populous jurisdictions with inclusionary zoning programs but limited CoreLogic coverage. I especially focus on jurisdictions where the inclusionary zoning programs have been in place for at least a few years and they use a unit-based threshold, and am limited to places with available open data that includes the information I need for this analysis, like number of units per building. I use ten residential building permit databases and property characteristics databases covering the following jurisdictions:

1. Denver, CO
2. Fairfax County, VA
3. Honolulu, HI

4. Minneapolis, MN
5. Montgomery County, MD
6. New Jersey
7. Portland, OR
8. San Diego, CA
9. San Francisco, CA
10. San Jose, CA

After the data is limited to buildings from inclusionary zoning jurisdictions built or permitted after their inclusionary zoning policies began, 70% of the observations are from CoreLogic, and 30% are from the supplemental datasets.

Table 1: Data Summary

| | |
|---|---|
| Number of programs | 141 |
| Number of buildings: | 33,795 |
| Most common states, by buildings: | CA (41%), MA (22%), WA (12%) |
| Most common cities, by buildings | San Diego, CA (17%), Boston, MA (9%), Cambridge, MA (7%) |
| Most common states, by programs: | MA (50%), CA (23%), WA (4%) |

Figure 1

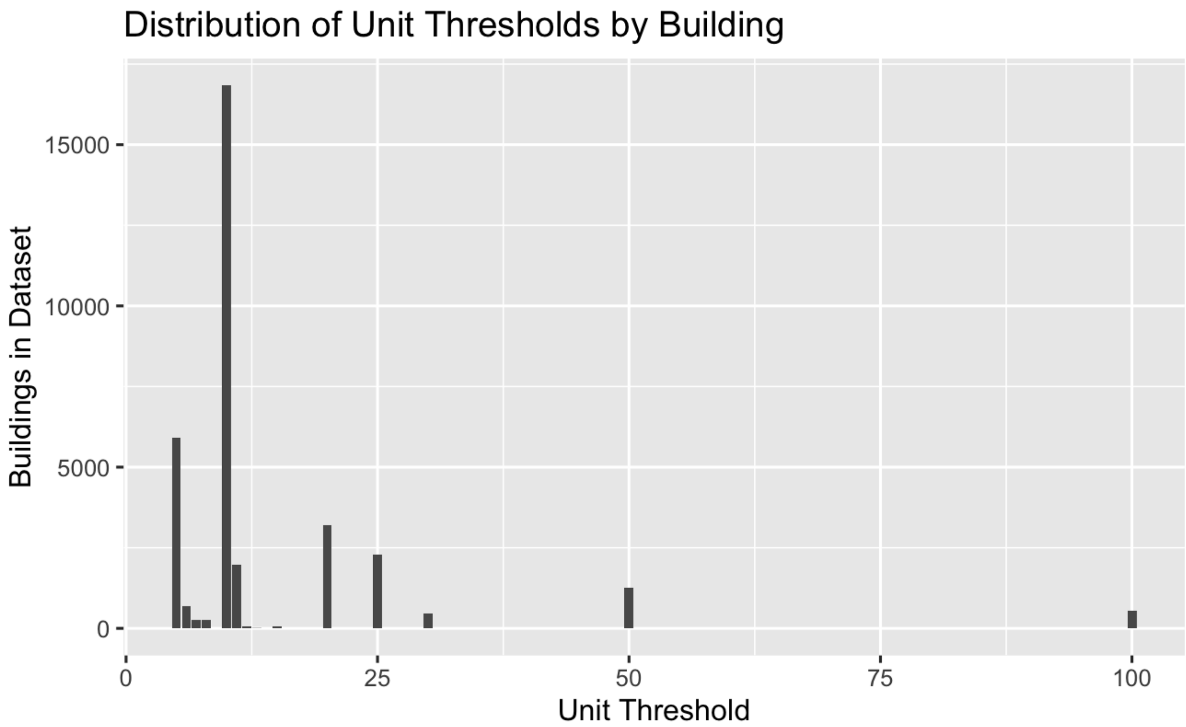


Figure 1 shows how many buildings in the data are in jurisdictions with inclusionary zoning programs that take effect above a certain threshold. The largest sample size is for policies that take effect above 10 units, but there are some as low as 5 units and some as high as 100 units.

<u>Figure 2</u>

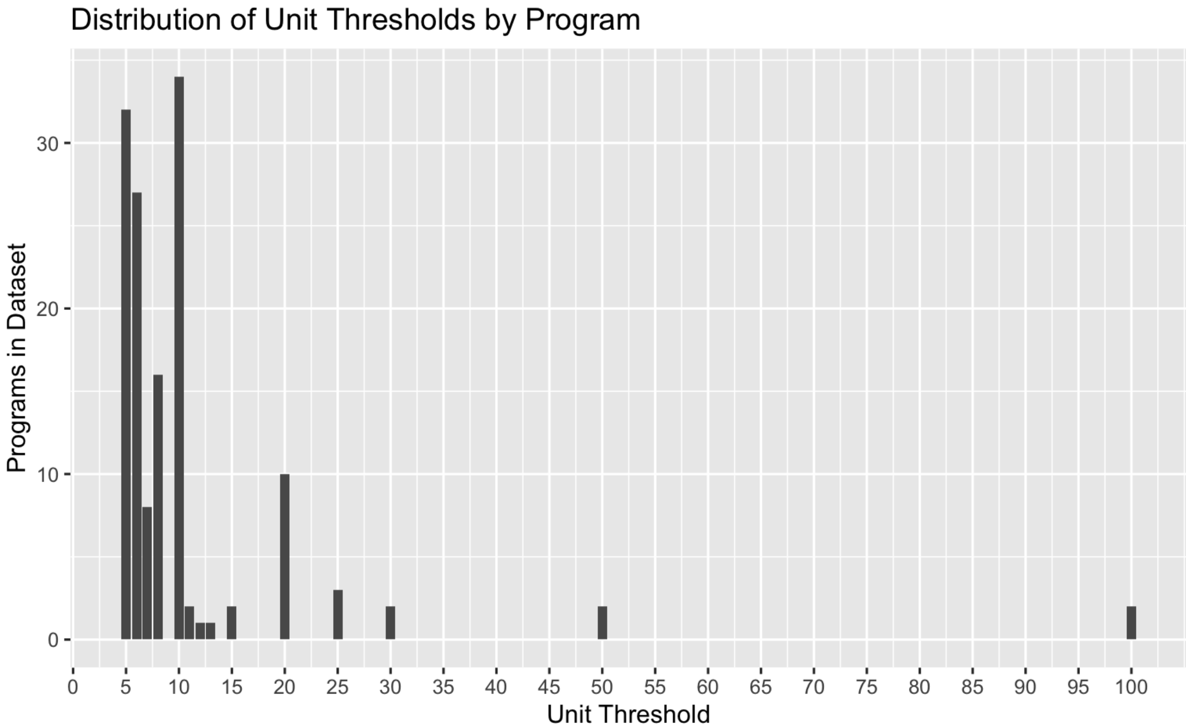


Figure 2 shows the distribution by program, which looks different from Figure 1 because some cities have more buildings in the dataset than others. Most programs have thresholds between 5 and 10 units.

<u>Figure 3</u>

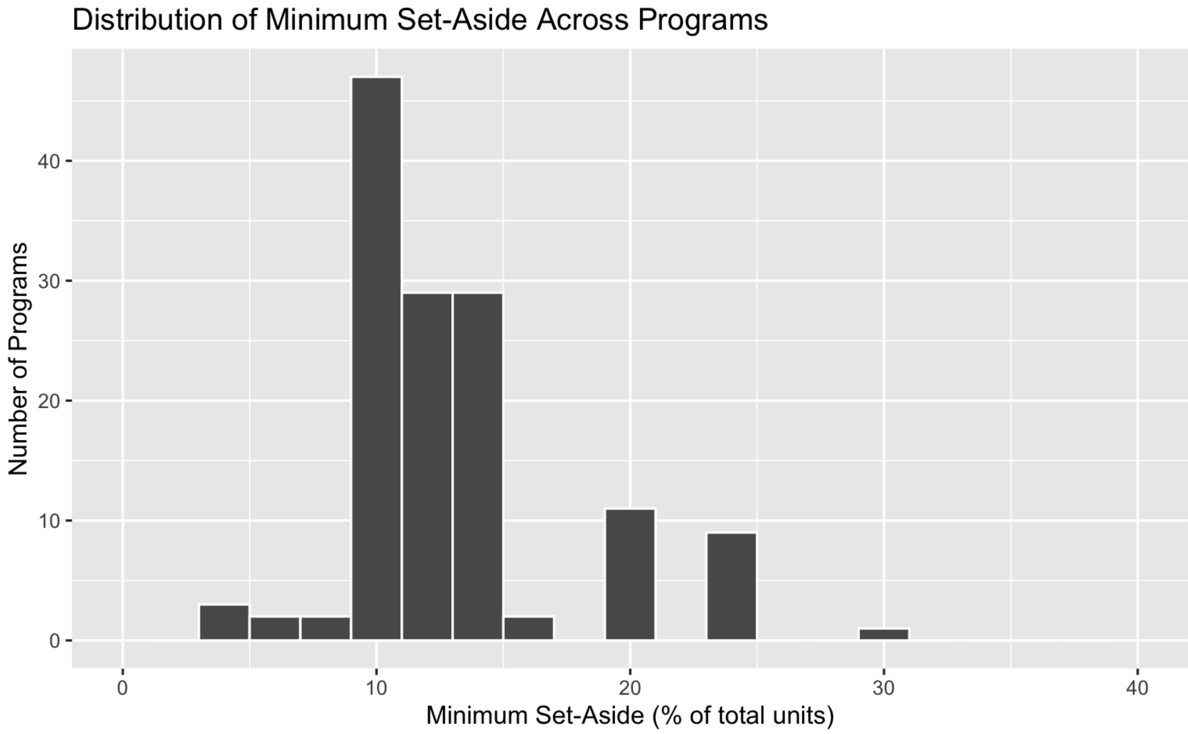


Figures 3 and 4 show information about the details of inclusionary zoning programs in the data. Figure 3 shows the minimum percentage of units in an apartment building that must be affordable and set aside for low-income tenants. Most programs set the percentage between 10 and 20. Figure 4 shows the distribution of a statistic measuring the degree of affordability that the Inclusionary Housing Database calls “lowest income served.” Usually, inclusionary zoning programs specify a range of incomes based on percentages of the median income in the metropolitan statistical area. Sometimes, a program might have multiple ranges, for example 5% of units must be affordable to households making 30-50% of Area Median Income (AMI), and another 5% must be affordable to households making 60-80% of AMI. The lowest income served is the lowest upper-bound among all the ranges. So in the hypothetical example, the upper-bounds are 50% and 80% so the lowest income served is 50%. The reasoning for this is that developers can always choose to price apartments at the upper end of the ranges, so that’s what really matters. There is also a "highest income served" statistic available in the database, for the highest upper-bound among all the ranges. The most common value of lowest income served in the data is 80% AMI.

Figure 4

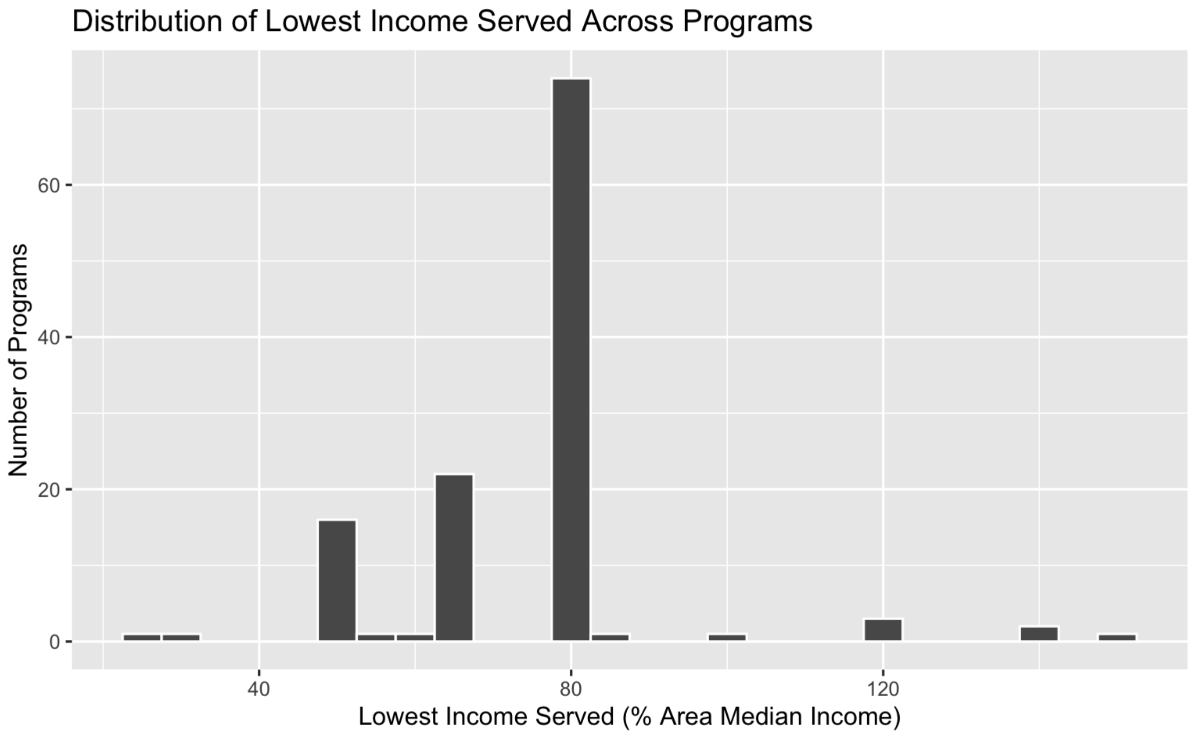


# Multi-City Analysis

Figure 5

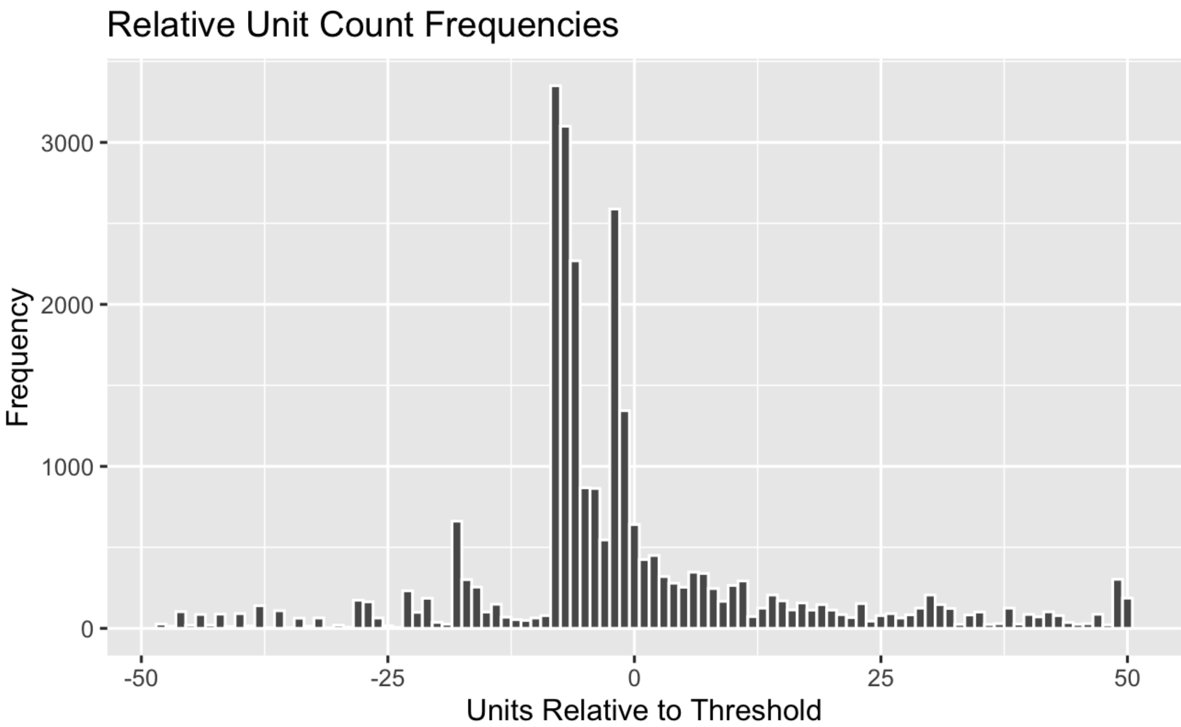


Figure 5 plots "relative unit counts" which are calculated by subtracting the minimum threshold at which inclusionary zoning takes effect in a city from the number of units an apartment building has. For example, if inclusionary zoning applies to buildings with 10+ units, and a building has 15 units, the relative unit count is 5. This can allow us to aggregate data from programs with different unit counts and check for bunching below the thresholds with higher precision.

<u>Figure 6</u>

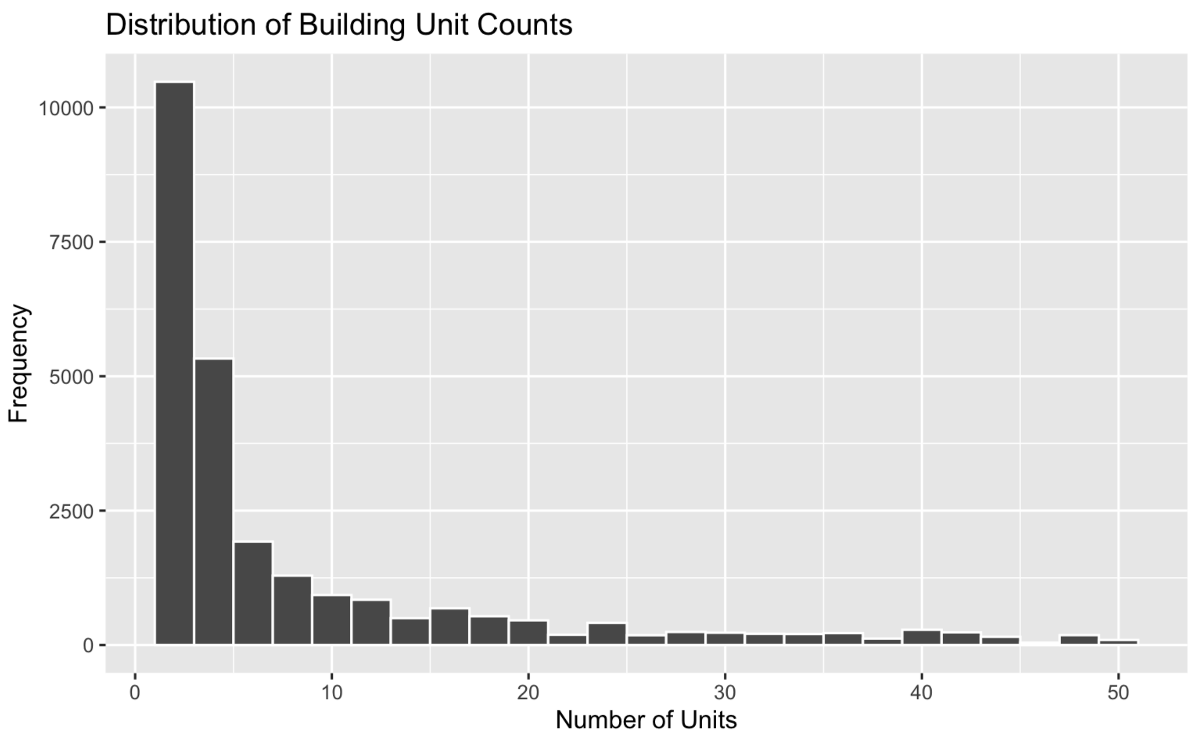


Figure 6 displays the raw distribution of unit counts across buildings in the whole dataset, not adjusted for threshold. It appears that in the data as a whole, there is a high number of small buildings with 2-4 units and the frequency of a building having any specific unit count declines rapidly as the unit count increases. Therefore, the distribution might be better analyzed separately by unit count threshold. On the other hand, the sample size is lower when analyzing each threshold group separately. We can compromise by grouping the data into four buckets with similar unit count thresholds.

Figure 7: Unit Count Distributions Separated by Threshold Level

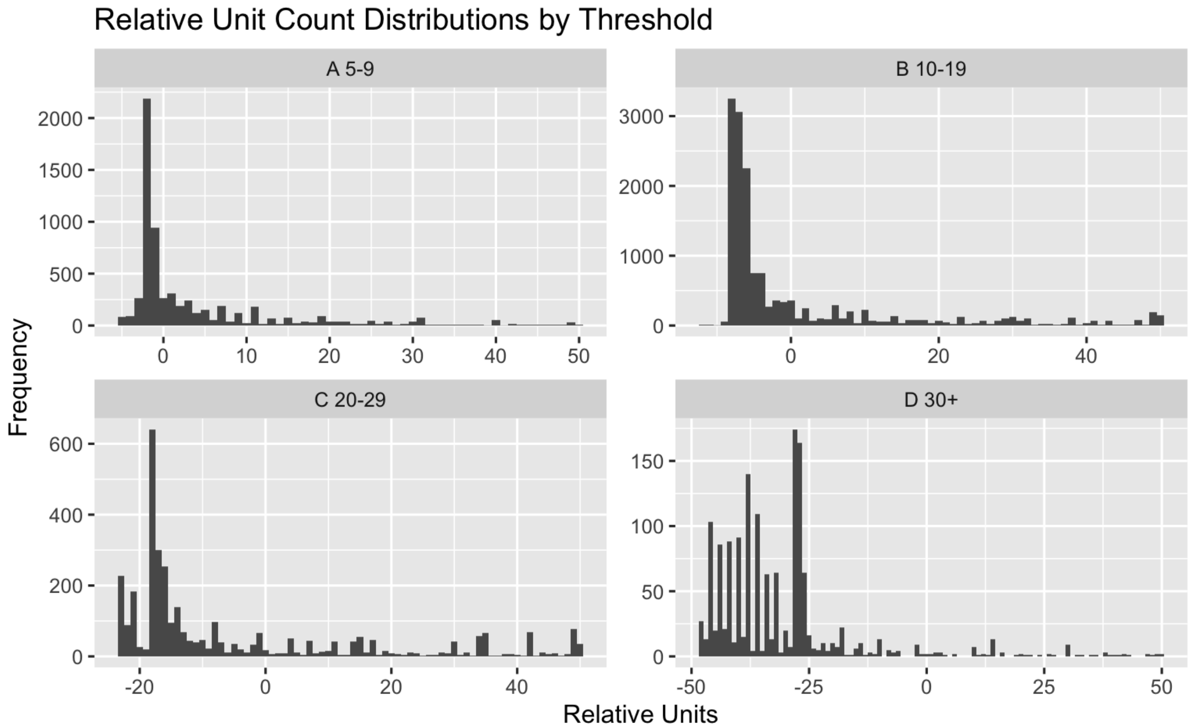


Visually, we do not see much bunching below the threshold of zero on any of the plots in Figure 7. However, it is useful to answer the question more precisely and objectively using hypothesis testing.

The McCrary Sorting Test is commonly used to check for excess mass in a distribution to one side of a threshold (McCrary 2008). In this study, we can use the McCrary Test to measure the statistical significance of any bunching we observe below program thresholds. If we find that the density is continuous at the threshold value, that suggests developers are not responding to the policy by strategically building apartment buildings right below the threshold. If we find a significant discontinuity, that indicates developers may be behaving in the hypothesized way.

Table 2: McCrary Test Results for Programs Grouped by Unit Count Threshold Ranges

| group | n_buildings | n_programs | binsize | theta | se | p | z |
|---|---|---|---|---|---|---|---|
| A 5-9 | 7106 | 83 | 0.900 | -0.548 | 0.026 | 0 | -21.369 |
| B 10-19 | 18946 | 40 | 1.052 | -1.524 | 0.023 | 0 | -67.425 |
| C 20-29 | 5486 | 13 | 2.987 | -1.207 | 0.043 | 0 | -28.321 |
| D 30+ | 2257 | 6 | 4.006 | -2.531 | 0.131 | 0 | -19.311 |

Table 2 shows that the distributions of building unit counts from all four groups of inclusionary zoning programs, organized by ranges of unit count thresholds (5-9, 10-19, 20-29, and 30+ units) have extremely highly significant discontinuities at the thresholds according to the results of the McCrary tests. However, this alone is not enough to conclude that there is bunching in the hypothesized way, because the distributions could simply be discontinuous at many points, not just at the threshold.

To see whether the discontinuities are especially large at the thresholds compared to other points, we can run a placebo check analysis in which we re-run the McCrary Tests at every possible point along the x-axis, record the "theta" value for the size of the discontinuity, and compare the actual value to the distribution of placebos.

Figure 8: McCrary Test Placebo Checks

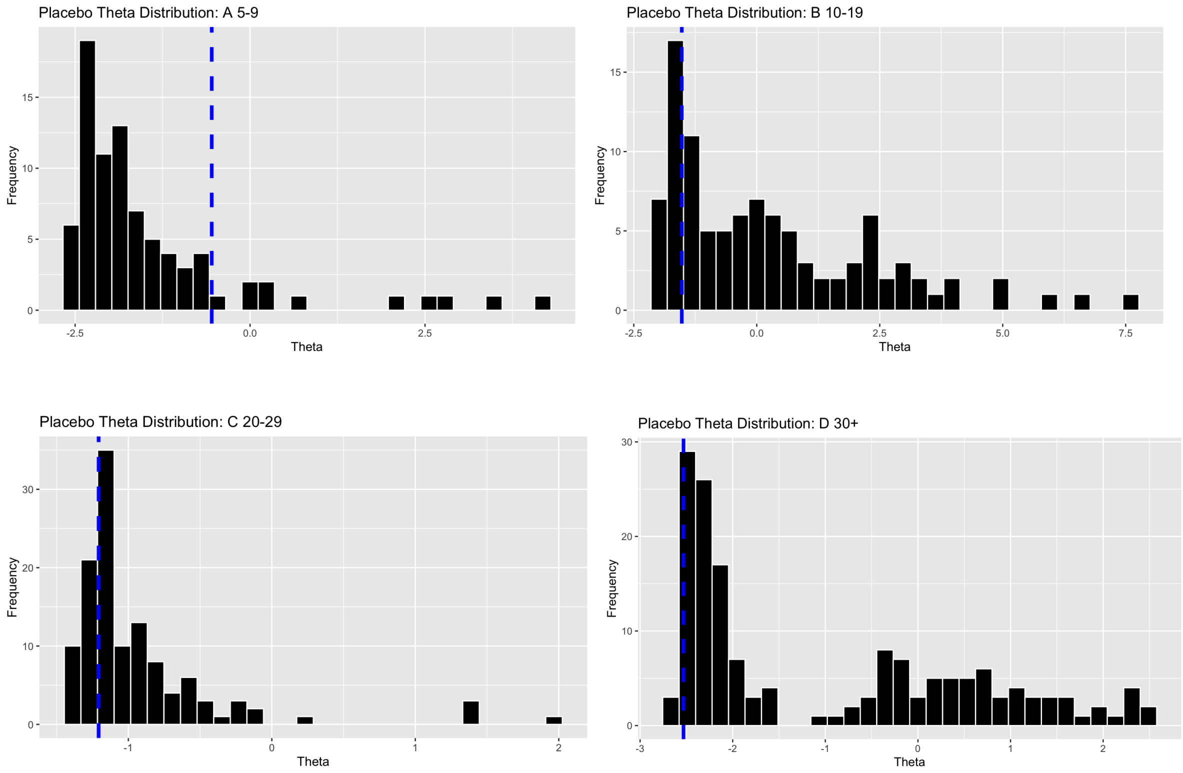


Figure 8 shows the distribution of the placebo theta values for each group of data consisting of buildings in cities with similar unit-count thresholds for their inclusionary zoning policies. The dotted blue lines show the value of the actual theta statistic. When an effect is real, you would expect to see the actual theta value to one side of almost all of the placebos. In Figure 8, we see that the observed theta value is unremarkable for groups A, B, and C, but the value in D does seem to be smaller than almost all of the placebo values.

<u>Table 3: Placebo Check P-Values</u>

| Threshold Range | p (two-tailed) | p (one-tailed) |
|---|---|---|
| A 5-9 | 0.95181 | 0.87952 |
| B 10-19 | 0.48000 | 0.22000 |
| C 20-29 | 0.31405 | 0.28099 |
| D 30+ | 0.05590 | 0.05590 |

The placebo checks show that the discontinuities are not significant at all in the first three groups, but is borderline significant with p=0.056 in the group with the highest thresholds of 30 units or more. I ran two-tailed tests, so the p-value is calculated as the percent of placebo thetas which have greater absolute value than the absolute value of the observed theta. I also ran one-tailed tests for robustness, and the conclusions are the same.

We can also scan the dataset for individual cities and counties with the clearest signs of bunching.

Table 4: Cities and counties with the most bunching

| City/County | State | Government Type | Unit Threshold | N Bunch | N Pre-Bunch | N Post-Bunch | Total Buildings | Bunch Minus Pre-Bunch |
|---|---|---|---|---|---|---|---|---|
| PORTLAND | OR | Town or city | 20 | 54 | 16 | 3 | 945 | 38 |
| REDMOND | WA | Town or city | 10 | 31 | 4 | 2 | 2299 | 27 |
| KENMORE | WA | Town or city | 20 | 19 | 0 | 0 | 416 | 19 |
| CHARLTON | MA | Town or city | 6 | 27 | 17 | 7 | 69 | 10 |
| FAIRFAX | VA | County | 50 | 9 | 0 | 1 | 1205 | 9 |
| EVERETT | WA | Town or city | 8 | 8 | 0 | 0 | 52 | 8 |
| SAN FRANCISCO | CA | Consolidated city-county | 25 | 14 | 7 | 5 | 851 | 7 |
| SUNAPEE | NH | Town or city | 8 | 9 | 5 | 3 | 35 | 4 |
| WEST HOLLYWOOD | CA | Town or city | 11 | 7 | 3 | 0 | 94 | 4 |
| MEDWAY | MA | Town or city | 6 | 3 | 0 | 0 | 4 | 3 |

Table 4 shows cities and counties that appear to have the most bunching according to the following procedure. “N Bunch” is the number of buildings with unit counts at the point where we would expect to find bunching, namely 1 unit below the threshold, or in the case of even-numbered thresholds, 2 units below the threshold as well. That complication is due to the disproportionate tendency for buildings to have even-numbered unit counts. Bunching may also occur at the highest even-numbered unit count below the threshold as a result. “N Pre-Bunch” is the number of buildings right below that, 2-3 units below an odd-numbered threshold or 3-4 units below an even-numbered threshold. Normally, we would expect “N Pre-Bunch” to be higher than “N Bunch” because of the inverse relationship between a unit count and the frequency of buildings with that number of units (see Figure 6). So if “N Bunch” is higher than “N Pre-Bunch”, that suggests anomalous bunching below the threshold. The difference statistic in the last column measures that. “N Post-Bunch,” the number of buildings exactly at the minimum threshold, is also useful information because if it is relatively high, that contradicts the hypothesized behavior - it suggests developers would rather build a single extra unit than be free from the affordability requirement. The table is sorted in descending order by the raw difference metric, without any normalization by total number of units. Portland, Oregon has the most bunching by this metric, but as is evident from the fact that the difference metric quickly drops down to just 3 or 4 units as you go down the table, the vast majority of cities and counties in the dataset have none or almost none of the bunching pattern consistent with the hypothesis.

Figure 9

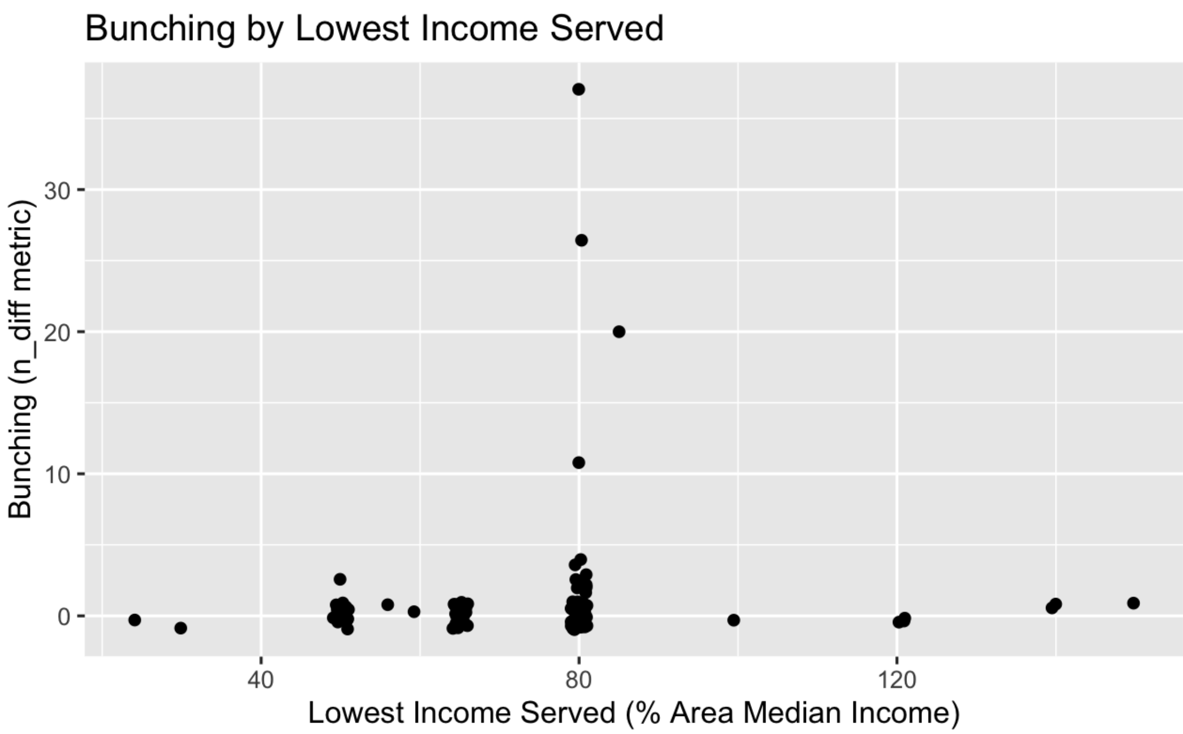


Figures 9 and 10 show the relationship between elements of inclusionary zoning policies and the degree of bunching observed across cities. The y-axis shows the "bunch minus pre-bunch" difference metric for bunching, abbreviated as "n_diff" in the labels, which I lower-bounded at zero. The x-axis displays the program attributes discussed in the "data" section, the percent of units set-aside for low income tenants, and the degree of affordability required. A problem with both visualizations is overlap from points with the same values on both axes, which I addressed by adding a small amount of random "jitter" so the dots do not overlap as much. My conclusion from both plots is that the handful of cities with notable bunching do not have especially heavy affordability requirements based on these metrics and compared with other programs in the database. It does not seem like we can explain the unusual degree of bunching in cities like Portalnd, Oregon based on either of these fundamental program attributes.

Figure 10

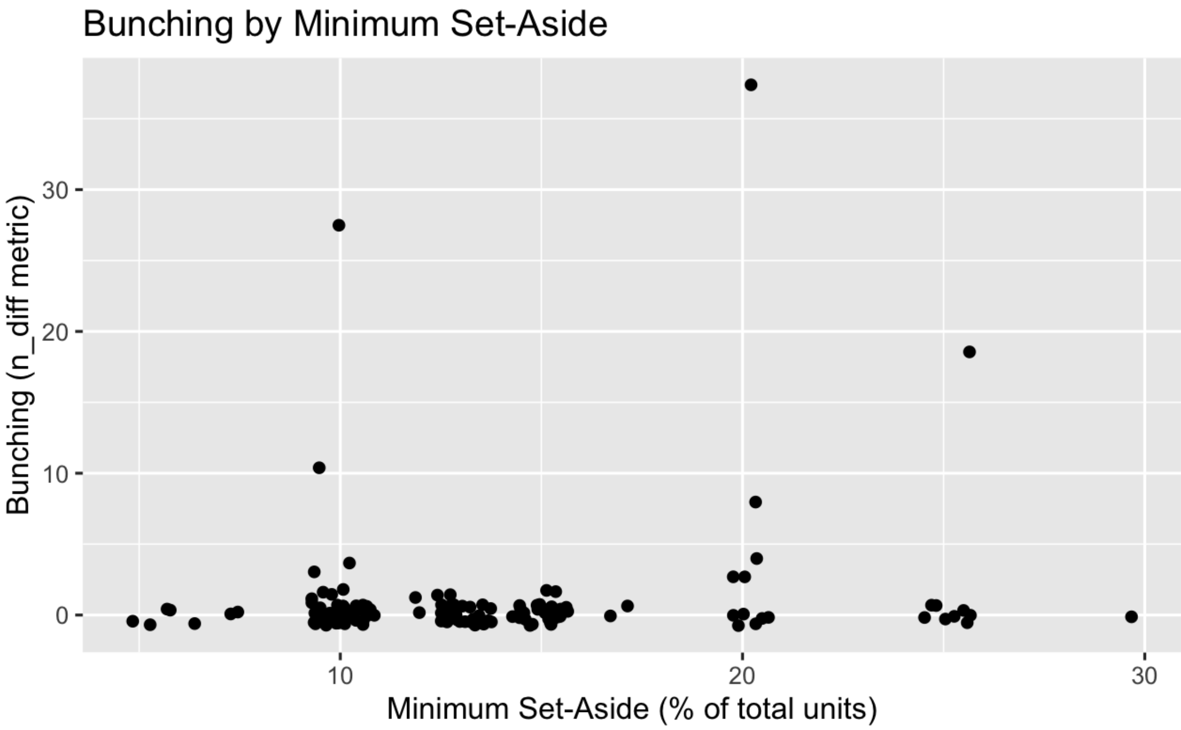


# Case Study: San Francisco, CA

It is also useful to analyze single case studies in more depth to explore the research question in more detail and handle individual peculiarities in program design, such as having complex programs with multiple thresholds. San Francisco, California adopted an inclusionary zoning policy in July 2016 with a moderate affordability requirement for buildings between 10 and 24 units, and a higher affordability requirement for buildings with 25 units. To check whether any bunching in the distribution of apartment building unit counts from after the policy took effect might be explained by some other factor besides inclusionary zoning, we can examine the distribution from before the policy was implemented to see if there is bunching below the thresholds. Based on Figure 11, there was only a small amount of bunching below either threshold.

Figure 11

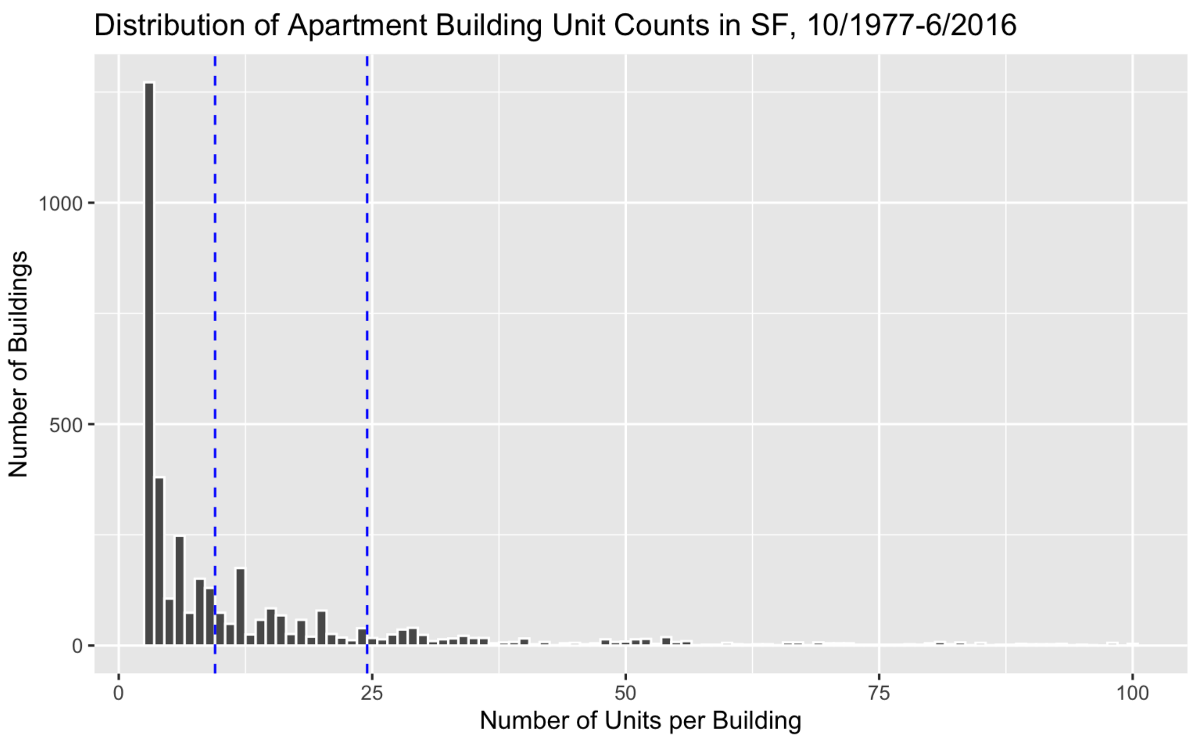


Figure 12

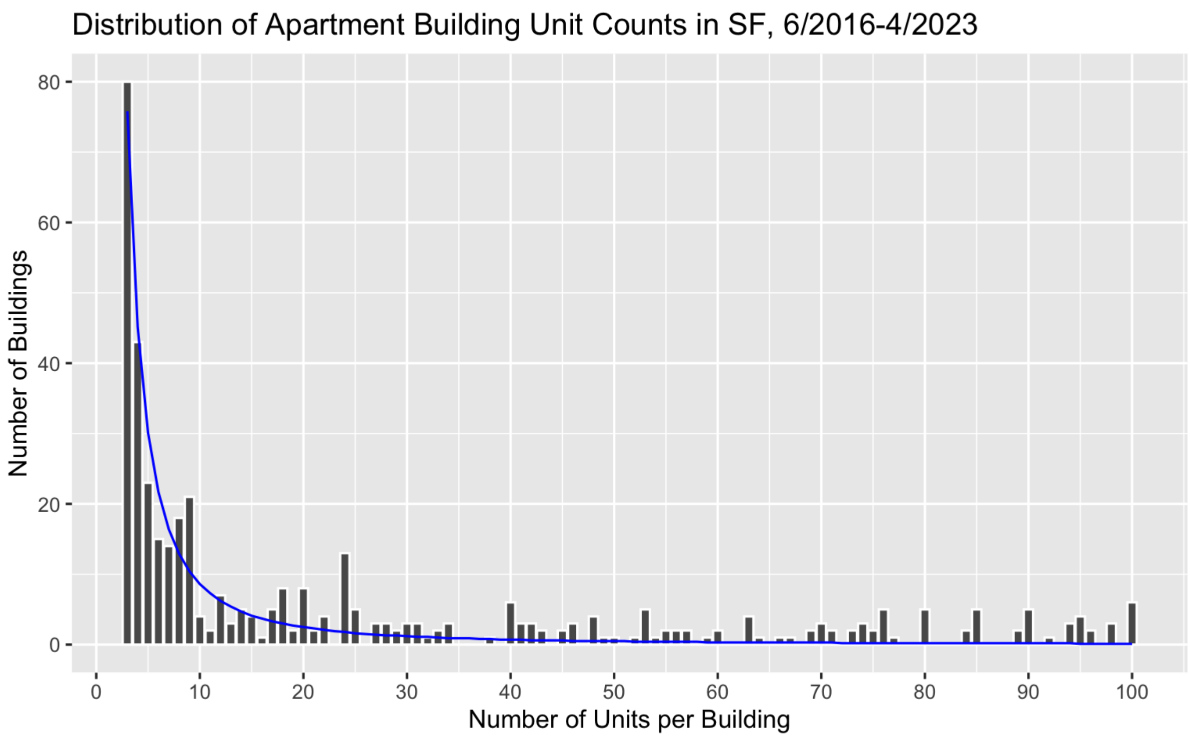


Figure 12 shows the distribution of building unit counts in the period after inclusionary zoning was implemented. You can see noticeable bunching right below both the 10-unit and 25-unit thresholds. The blue line in Figure 12 is an inverse power law regression model fit to the entire distribution. The inverse power law model was selected after comparing its fit with alternatives, including an inverse squared model and an exponential model.

<u>Figure 13</u>

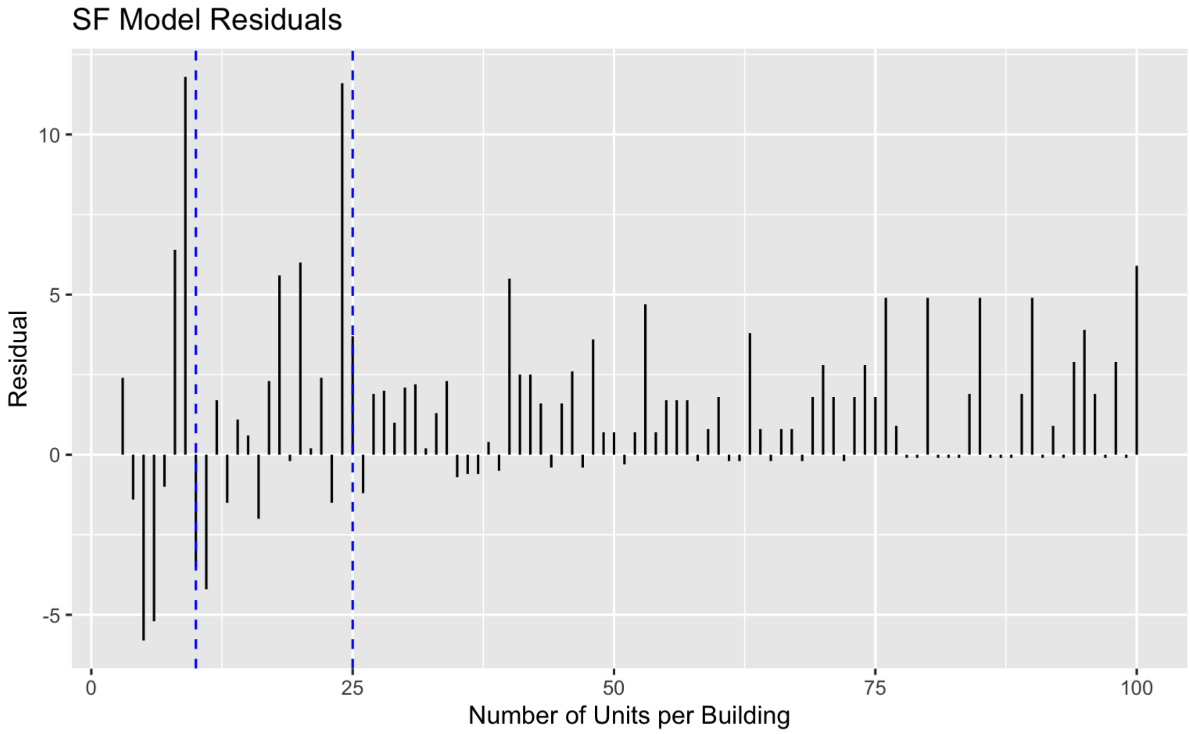


Figure 13 displays residuals from that model, which can be used to identify deviations from the overall trend defined by the inverse power distribution that are suggestive of an anomaly, such as the hypothesized bunching below the thresholds. The blue dotted lines mark the thresholds, and you can see that the two greatest residuals (including as measured by absolute value) are directly to the left of the dotted lines.

<u>Figure 14: Bootstrap Results for 25-Unit Threshold Test</u>

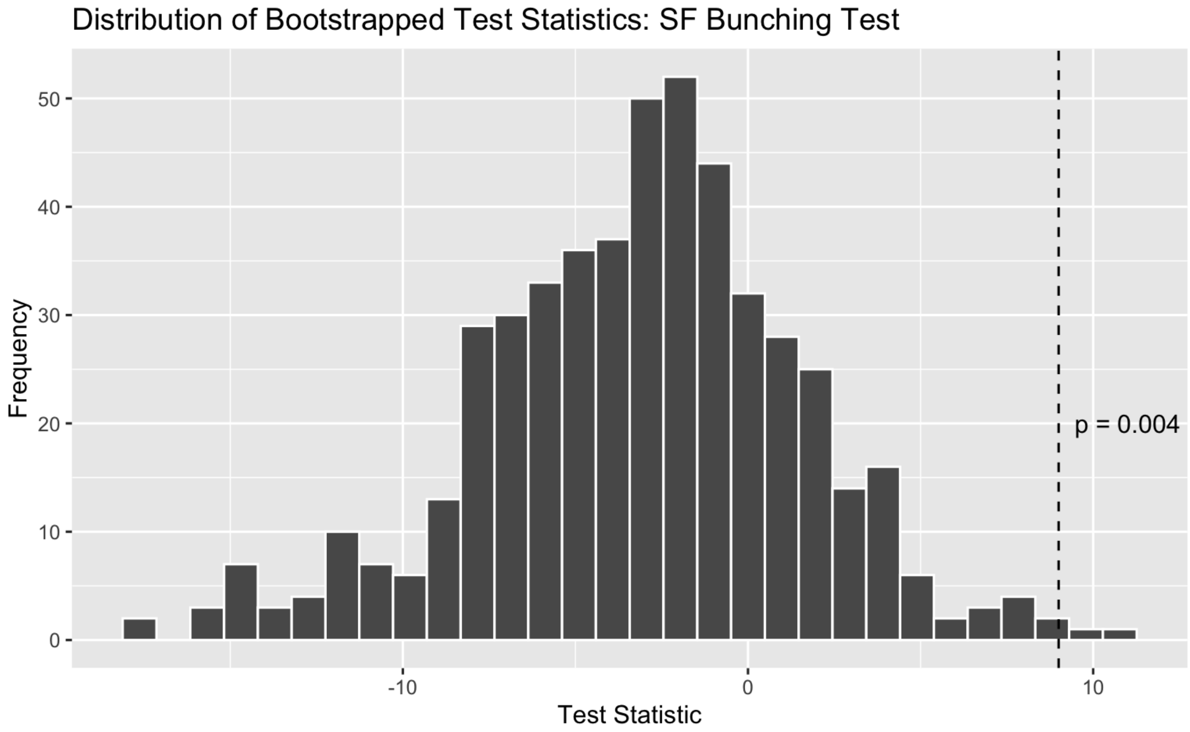


In addition to the visual evidence, we can run a hypothesis test for the significance of the bunching. For the test, I will focus on the 25-unit threshold since that one determines whether the heaviest inclusionary zoning requirement applies. Using bootstrapping with residuals resampling of the power law model, we obtain a one-tailed p-value of 0.004 showing that a residual as large as what we observe would be very unlikely to occur by chance if there were no anomalies. This indicates that the bunching is a significant feature of the distribution and not a chance occurrence.

The additional bunching below the 10-unit threshold provides even stronger evidence that the bunching is real. If I use the sum of the residuals below the 10-unit and 25-unit thresholds as the test statistic, I get a p-value of zero from a bootstrapped hypothesis test with 500 resamples. This means that bunching below the two thresholds, when considered together, is so significant that none of the bootstrap simulations produced bunching as high as what we observe in the residuals from the model fit to the actual data.

Figure 15: Bootstrap Results for Double Threshold Test

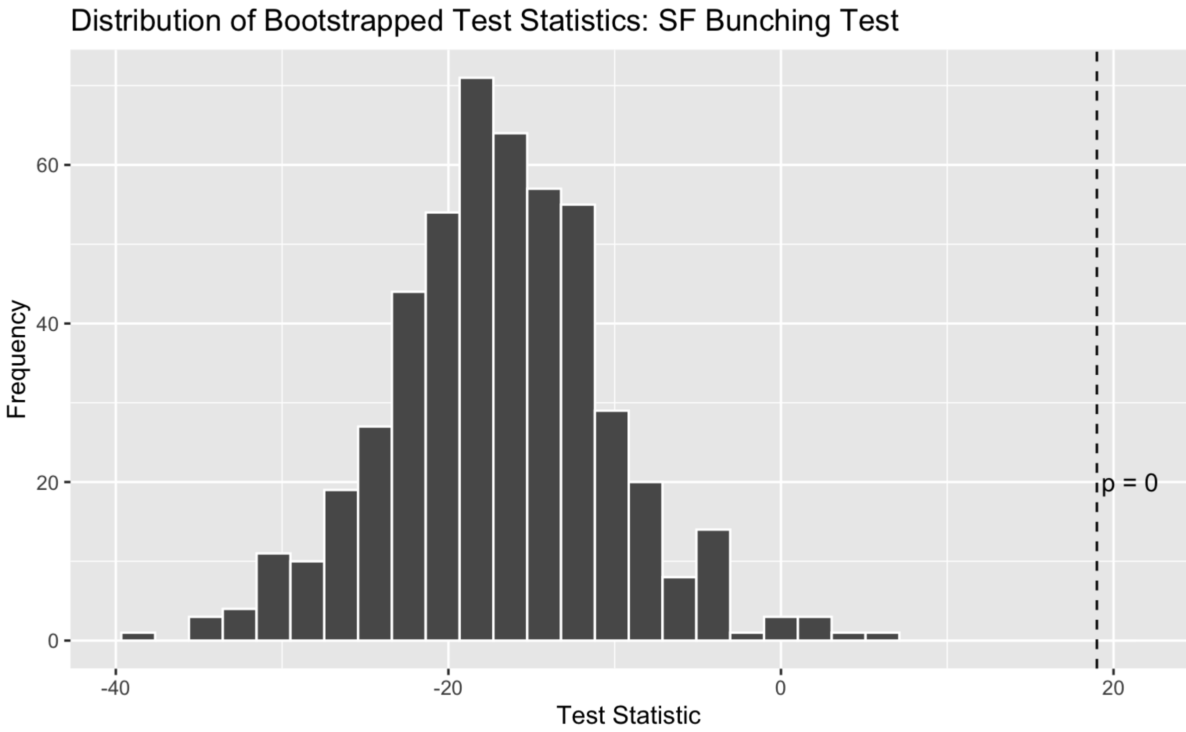


# Discussion

Imagine a hypothetical city where there is a very high demand for housing, a highly-burdensome inclusionary zoning requirement that applies to buildings above a unit count threshold, and minimal restrictions on residential building construction otherwise. In such a city, it would be highly likely to see bunching below the threshold. Under the right conditions, it is clear developers will try to evade the regulation. So the question of this paper is not whether such behavior is possible under the right conditions, but rather how developers behave given the details of the housing markets and policies in the U.S. cities where inclusionary zoning has been in place in recent decades.

Specifically, I hypothesize that a housing developer's response to inclusionary zoning and whether there is bunching under the threshold to avoid compliance depends on a few factors.

1. Is this a high-demand housing market with lots of demand for apartments? (if yes, we should see more of a bunching effect)
2. Are the other regulations besides inclusionary zoning friendly enough to developers that it is possible to build many below-threshold buildings? (if yes, we should see more of a bunching effect)
3. Is the inclusionary zoning especially burdensome on developers in terms of the percent of units required to be affordable and the degree of affordability? (if yes, we should see more of a bunching effect)
4. Are there other, easier ways for developers to avoid inclusionary zoning such as paying a small fee or getting a waiver? (if yes, we should see less of a bunching effect)

Based on the empirical findings of this paper, it seems like in the vast majority of U.S. cities with inclusionary zoning policies, the conditions are such that developers do not strategically build right below unit count thresholds. The relative importance of each factor I laid out, or others, is outside the scope of this paper. But this paper's findings suggest that in most cases, developers do not build right below the threshold to an unusually frequent extent.

On the other hand, the few case studies identified in which developers do appear to exhibit the hypothesized strategic behavior, such as San Francisco, CA and

Portland, OR, suggest that the conditions necessary for such behavior are not far outside of the realm of real-world market and policy conditions in the U.S. They offer support for the idea that under certain conditions, developers will avoid having to comply with an inclusionary zoning policy by lowering the unit count of a proposed apartment building right below a minimum threshold. But even in the case studies where we observe bunching, the effects are modest in magnitude. And a preliminary investigation into the relationship between the degree of bunching and two key program attributes measuring the degree of affordability required did not reveal any predictive relationships.

# Conclusion

One area for future research would be investigating what factors moderate whether inclusionary zoning policies lead to bunching below thresholds, as well as the extent to which inclusionary zoning policies reduce the supply of housing. For example, structural modeling could consider the relationship between various factors hypothesized to be moderators and the outcomes of interest.

Another area for future research would be to apply the methodology of analyzing bunching below thresholds to other housing policies that take effect above thresholds besides inclusionary zoning.

Lastly, this paper studies behavior along the intensive margin, the number of units per apartment building. As a complement of this research, it would be useful to study how inclusionary zoning affects new housing development along the extensive margin, whether a developer decides to build a new building at all or not. Researchers could use property data from many cities to study the effects of inclusionary zoning on the housing supply using a difference-in-differences or panel data modeling approach, as has been done in previous research with a more limited set of case study cities.

# Bibliography

Basolo, V., & Calavita, N. (2004). A Critique of the Reason Foundation Study on Inclusionary Housing Policy in the San Francisco Bay Area.

Bento, A., Lowe, S., Knaap, G. J., & Chakraborty, A. (2009). Housing market effects of inclusionary zoning. Cityscape, 7-26.

Grounded Solutions Network. (2020). Inclusionary Housing Database. Retrieved from http://inclusionaryhousing.org/map/

Hamilton, E. (2021). Inclusionary zoning and housing market outcomes. Cityscape, 23(1), 161-194.

Kleven, H. J. (2016). Bunching. Annual Review of Economics, 8, 435-464.

Li, F., & Guo, Z. (2022). How Does an Expansion of Mandatory Inclusionary Housing Affect Housing Supply? Evidence From London (UK). Journal of the American Planning Association, 88(1), 83-96.

Mast, E. (2021). JUE Insight: The effect of new market-rate housing construction on the low-income housing market. Journal of Urban Economics, 103383.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. Journal of econometrics, 142(2), 698-714.

Means, T., & Stringham, E. P. (2012). Unintended or Intended Consequences? The Effect of Below–market Housing Mandates on Housing Markets in California. Journal of Public Finance and Public Choice, 30(1-3), 39-64.

Mukhija, V., Regus, L., Slovin, S., & Das, A. (2010). Can inclusionary zoning be an effective and efficient housing policy? Evidence from Los Angeles and Orange Counties. Journal of urban Affairs, 32(2), 229-252.

Powell, B., & Stringham, E. P. (2004). Housing Supply and Affordability: Do Affordable Housing Mandates Work?. Reason Foundation policy study, 318, 1-45.

Ramakrishnan, K., Treskon, M., & Greene, S. (2019). Inclusionary zoning: What does the research tell us about the effectiveness of local action?.

Schuetz, J., Meltzer, R., & Been, V. (2011). Silver bullet or Trojan horse? The effects of inclusionary zoning on local housing markets in the United States. Urban Studies, 48(2), 297-329.

Wang, R., & Balachandran, S. (2021). Inclusionary housing in the United States: dynamics of local policy and outcomes in diverse markets. Housing Studies, 1-20.

Wang, R., & Fu, X. (2022). Examining the Effects of Policy Design on Affordable Unit Production Under Inclusionary Zoning Policies. Journal of the American Planning Association, 88(4), 550-564.