Inclusionary Zoning and Housing Supply
Evidence from California's Palmer Fix *

Noah Kouchekinia†

University of California, Irvine
Department of Economics, Social Science Plaza A (SSPA) 3139, Irvine CA 92697, United States

June 11, 2026

Abstract

Inclusionary Zoning requires developers to set aside a share of new units at affordable rates to low/moderate income tenants. In California, more than a third of localities have adopted an ordinance. Mandatory inclusionary zoning acts as a tax on construction and, all else equal, reduces the overall supply of housing. This potentially reduces affordability in the aggregate. Quantifying this trade off has been difficult due to heterogeneous local policies, data quality issues, and the endogeneity of local inclusionary zoning adoption. I estimate the effect of inclusionary zoning using novel administrative data covering all cities and counties in California. I leverage a statewide change in preempting law for identification, performing a difference-in-difference around an exogenous reactivation of local inclusionary ordinances. I find the typical inclusionary zoning ordinance reduces annual new residential construction by nearly a third (31.8%). Additionally, I find the policies relocate multi-family developments within jurisdictions to lower rent areas, where the policy is effectively less severe. Accounting for the stringency of policies turns out to be critical, and heterogeneous effects by policy stringency may explain divergent results in the prior literature. I estimate the cost of generating an affordable unit with inclusionary zoning to be approximately $800,000 in "excess rents" paid by market rate renters as a result of the policy's constraint on supply. This exceeds the cost of directly incentivizing the creation of low-income housing in California through existing programs like LIHTC.

JEL Classifications: R52, R31, R23, R13, J68, I38

Keywords: Inclusionary Zoning, Affordable Housing, Spatial Sorting, Location Choice

*This research was supported by the California Policy Lab and the Horowitz Foundation.
†Email: nkouchek@uci.edu

1 Introduction

Inclusionary Zoning, a popular local housing regulation, requires developers to set aside a share of new units at affordable rates to low/moderate income tenants. Nearly 900 cities across the US have adopted an inclusionary ordinance, with the policy being most popular in high-cost metro areas in the Northeast and Pacific West. In California, more than a third of localities have adopted an ordinance, with the number of associated affordable units rivaling the largest federal housing policies (Wang and Balachandran, 2021).

These ordinances act as taxes on construction, and ought to reduce the overall amount of housing produced. However, existing evidence on Inclusionary Zoning's supply effects has been mixed. While some studies find no effect on the flow of new housing (Bento et al., 2009; Hamilton, 2021; Mukhija et al., 2010), others find a reduction (Means and Stringham, 2012), or mixed results (Phillips, 2024; Schuetz et al., 2011). Inclusionary Zoning's supply effects are difficult to study because local ordinances are highly heterogeneous and there was, until now, a lack of high quality data on their key attributes.

I estimate the effect of inclusionary zoning ordinances on housing supply using novel administrative data on all California cities, leveraging a statewide change in preempting law for identification. This advances the literature on several fronts. First, the new data measures inclusionary zoning ordinances' stringency. This allows me to estimate the intensive margin of inclusionary zoning's effect on housing supply. Second, by using an exogenous shock to localities' inclusionary zoning ordinances, I evade the confounding effects of simultaneously adopted local housing regulations. Additionally, the statewide shock allows me to use a standard difference-in-difference estimator without the onerous assumptions required for a staggered design (Goodman-Bacon, 2021). Finally, I extend my analysis to estimate inclusionary zoning's effect on the composition of housing constructed within adopting jurisdictions, showing it shifts construction to lower-rent neighborhoods.

I have assembled a novel panel of California localities' inclusionary housing policies using administrative reports. Through public records requests and some archival work, I have acquired over 2000 records from about 500 California localities, dating from the mid 90s through the present.

In using administrative data, I avoid high levels of measurement error in municipal land use regulation surveys (Lewis and Marantz 2019). I am also able to measure local policy with great detail. I recruited and trained readers who extracted 20 policy parameters of each jurisdictions inclusionary zoning ordinance. These capture the properties, breadth, and depth of each inclusionary ordinance. From these, a measure of policy stringency is constructed which captures what share of future rents landlords give up to comply with the policy. This is the implicit tax of inclusionary zoning.

I then perform a difference-in-difference estimation around a recent change in California policy. Inclusionary zoning was found to be an illegal form of rent control in Palmer v. LA (2009). The case abruptly halted the enforcement of local policies on the books. Policies remained dormant until a "Palmer Fix" was signed into law in 2017, reactivating all local inclusionary zoning ordinances. This allows causal identification of the effects of the ordinances using a difference-in-difference strategy around the reactivation date. This is preferable to previous approaches which compare housing production before and after policy adoption dates, which are endogenously selected and may correlate with other housing policy changes. Furthermore, they require a staggered difference-in-difference specification, which may be invalid in cases with dynamic and heterogeneous treatment effects.

Results show that for each 1 percentage point of implicit tax on rents, the flow of new housing is reduced 7.6%. This implies that the typical inclusionary zoning ordinance reduces annual new residential construction by 31.8%. Parallel trends leading up to policy reactivation reassuringly suggest the necessary identifying assumption holds. Results are robust to a binary specification splitting high stringency and low stringency policies, and a battery of alternative estimation strategies. Accounting for the stringency of policies turns out to be critical. Estimated effects are four times larger in the top quartile of IZ adopting jurisdictions than in the bottom quartile. Effects are statistically insignificant in the latter. This potentially explains the varied results of the prior literature.

I also examine how Inclusionary Zoning ordinances affect the composition and distribution of new housing supply. I find no effect on the composition of housing. Assessed values of newly built units, and the share of multifamily housing built are unaffected by the policy. I do find the policies shift development within jurisdictions, with larger reductions in new development in more expensive neighborhoods relative to less expensive neighborhoods. I also check for across jurisdiction spillovers from inclusionary ordinances, but find no evidence that IZ policies push development into neighboring localities.

I then turn my attention to implications for welfare. I am incorporating Inclusionary Zoning into a quantitative sorting model in the style of Diamond and Gaubert's (2022). The goal is to evaluate the trade off between subsidized rents, access to high income neighborhoods, and reduced housing supply for low income tenants. After calibrating the model to the pre-IZ equilibrium, I run policy counterfactuals to back out welfare effects. I find that inclusionary zoning ordinances as they currently stand are mildly beneficial to low income residents on the average, though the effect is uneven, with large gains to those lucky enough to get an IZ unit, and mild welfare losses for the rest. A hypothetical more stringent IZ policy generates large enough rent effects to more than counteract the benefits of the policy to low income residents through subsidized units.

Finally, I put my empirical estimates in policy context by comparing the cost of an IZ unit in terms of excess rents against other means of generating income restricted units.

2 Literature Review

While there is prior work which attempts to measure the effect on housing supply, the literature has failed to come to a persuasive consensus. Powell and Stringham (2004) kick off the debate with the descriptive finding that cities in the San Francisco Bay Area which adopt an inclusionary ordinance see a sharp reduction in building permits. This finding was followed by critique pointing to non-causal explanations for the correlation (Basolo and Calavita, 2004). A set of subsequent papers found no reduction in housing supply. Bento et al. (2009) conduct a comparison between inclusionary policy adopters and non-adopters in California, finding no effect on housing starts. The analysis is close in form to a difference-in-difference, however it lacks discussion of the assumptions required to interpret the estimates as causal. Mukhija et al. (2010) focuses specifically on the LA metro, again finding no correlation between IZ adoption and a reduction in housing production. Means and Stringham (2012) study California localities using a first-difference approach, binning cities' adoption decade. They find a statistically significant 7% reduction in housing construction, but standard tests supporting the parallel trends assumption are absent. More recently Hamilton (2021) examines inclusionary ordinances adopted in the Baltimore-Washington metro area, finding no reduction in housing supply.

Perhaps the highest quality work so far is that of Schuetz et al. (2011). They study the Boston and San Francisco metros, but their finding on housing supply is split, with a reduction found in Boston metro jurisdictions adopting the policy, but not in San Francisco. Phillips (2024) takes a different approach than the previously discussed analyses, instead using housing market models and previously estimated parameters to predict the effects of inclusionary ordinances. Critically, the paper predicts that effects on housing supply increase sharply with the stringency of the policy.

In general, the literature on this subject is plagued by dual challenges of endogeneity and heterogeneity. Cities select if and when they will adopt an inclusionary zoning ordinance. Conducting a difference-in-difference in a setting with staggered adoption is econometrically challenging and requires is generally not robust to dynamic or heterogeneous treatment effects (Goodman-Bacon, 2021). Furthermore, when a jurisdiction changes one housing policy, it is likely to change others. It is difficult to persuasively argue that adoption of inclusionary ordinances is not correlated with other changes in housing policy which could cause divergent trajectories in housing production, even if shown to be parallel before treatment. Cities adopt inclusionary ordinances that vary widely in their stringency. In the face of this complexity, nearly all prior analyses simply compare cities with a policy to those without. This makes it difficult to interpret effects or meaningfully compare effects outside of the specific contexts in which they are estimated.

Another strand of recent work has evaluated inclusionary zoning ordinances of a different form. Some jurisdictions directly compensate developers with tax credits who set aside below-market-rate units. These function more like local-LIHTCs. Ng and Wang (2025) evaluates how Seattle's program affects the spatial distribution of development. Soltas (2024) evaluates New York's ordinance, estimating the marginal cost of a below-market-rate unit. These estimate a fiscal rather than a housing market trade off to producing below-market-rate-units.

3 Data

Housing Elements

A key contribution of this paper is the use of Housing Element updates to measure local inclusionary zoning ordinances with great detail. The State of California requires that all cities and counties regularly report their local housing policies to state regulators. This process emerged from

California's mandate that local governments ensure sufficient housing production for community need. This mandate began in principle in 1969, but the level of state supervision, and the documentation local governments need to turn over to the state, has generally increased since then. Since the 1980s, local governments are required to submit updated housing elements to state regulators approximately once every seven years for review. More recently, state regulators have required that localities begin providing annual progress reports on housing construction.

Housing elements are technical documents meant for communication between local planners and state housing regulators. They describe a jurisdiction's housing stock, community needs, current housing policies, and planned improvements to their housing policies. Local planners' objective is to persuade state regulators that their housing policies for the next seven years will allow enough production to meet the housing needs projected by the state. Some important context is that local governments near universally fall short of state expectations. While this could in principle lead to fines and loss of planning authority, the state has taken a more collaborative approach, encouraging changes to local planning.

Housing elements are submitted to the state by local governments and reviewed by the department of Housing and Community Development (HCD), which checks for accuracy and policy sufficiency. Local governments then respond with revisions until the housing element is approved. This process can take several years.

Most of the housing elements is forward looking. Local governments promise state regulators compliant levels of housing production and facilitating policy. Because housing production has routinely fallen below state expectations, it is understandable to be skeptical of the promises made in localities' housing elements. Note that I rely instead on the retrospective portion of each housing element, where local governments simply report the state of their current housing policy.

Using housing elements to retrieve local policy parameters provides key advantages over the most broadly used alternative method: surveying localities, sometimes supplemented by reviews of local ordinances. Lewis and Marantz (2019) show that such surveys often return inconsistent results. Because producing housing elements has an administrative purpose, planners have the time and motivation to ensure they are accurate. The fact that HCD reviews housing elements gives further reassurance. Furthermore, housing elements contain far more detail than can be included in any reasonable survey. This is particularly important in the context of inclusionary zoning regulations, which often give developers a menu of options on how to comply. Information on each option extended to developers is necessary to assess the policy's stringency. Housing elements provide the the data to achieve this. Finally, housing elements have been updated 6 times over the last several decades, allowing me to trace policy development over time.

Housing elements from the 5th and 6th update cycle (concluding in 2013 and 2023 respectively)are made publicly available by the California HCD. Housing elements for the 2nd through 4th update cycles were retrieved through a public records request. [1]

Housing elements take a fairly standardized structure, but report local housing policies in a narrative form. In order to extract inclusionary zoning policy details, a team of coders reviewed each housing element, recording a series of parameters describing the breadth, depth, exceptions of each jurisdiction's inclusionary policy. Note that many jurisdictions had multiple ways developers could comply with the IZ requirement, for such jurisdictions, all options were recorded. Coders were randomly assigned jurisdictions, and extracted IZ policy parameters from available housing elements in chronological order within jurisdiction. Coders extracted parameters from 1848 housing elements from 575 unique jurisdictions. All five coders extracted parameters from 10% of the sample, to allow for inter-coder reliability checks.

I restrict my analysis to inclusionary zoning ordinances which apply to new rental property development broadly. I exclude policies which have a narrow geographic focus like single streets or special development areas smaller than a zip code. I exclude a number of ordinances which apply only to condominium conversions. Most critically, I exclude ordinances which apply only to the development of owner occupied housing. The exogenous policy variation I use from superseding state law is specific to rental units. I also currently only make use of two waves of housing element data. I rely on 6th cycle housing elements, collected between roughly 2019 and 2023, to measure inclusionary zoning policy attributes. I narrowly rely on the 4th cycle housing elements to categorize jurisdictions based on whether they intended to adopt an inclusionary ordinances before the Palmer decision. Data coverage is near universal, however 11 jurisdictions need to be excluded due to exceptional policies. This includes Long Beach, which phased in inclusionary zoning neighborhood by neighborhood, and Napa (City and County) which aggressively reformulated their inclusionary zoning policy as a tax following the Palmer decision.

[1] State archives of housing elements from the 1st cycle have been destroyed.

Construction Data

My primary outcome of interest is the flow of new housing. The conventional way to measure this is the Census Bureau's Building Permit Survey (BPS), which records the number of units, unit types, and the assessed value of residential projects permitted by jurisdiction. Headline results make use of BPS unit counts. Permit data is is aggregated from the permit office levels to the jurisdiction level. Jurisdictions are set of self-governing census places in California, and California counties net-of self-governing census places.

Official permit counts are supplemented by data from CoStar. CoStar data contains the precise geographic location of new units. This is useful for analyses in which I calculate inclusionary ordinance stringency at the neighborhood level. It also provides more detailed building attributes, which allows me to check for compositional changes in new housing stock. CoStar is somewhat limited in its coverage, containing data only on larger scale rental developments run by institutional landlords.

Additional Data

I made use of several additional data sets for context and extensions. Jurisdiction economic and demographic attributes are summarized with data from the American Community Survey. Fair Market Rents and Area Median Incomes are sourced from directly from HUD. These local economic indicators, along with jurisdiction's inclusionary zoning policy parameters are used to construct a measure of policy stringency at the neighborhood level. Neighborhoods are defined by the intersection of ZCTAs (the level of geography at which Small Area Fair Market Rents are measures) and jurisdictions (the level of geography at which inclusionary policy parameters are determined).

4 Empirical Strategy

Measurement

I assemble the inclusionary zoning ordinance parameters extracted from each city's housing elements to produce an economically interpretable measure of stringency. In order to comply with an inclusionary policy, building owners must rent units at specified below market rate rents. The cost to land owners can then be expressed as the share of rent forgone or τ_{pn} . This share depends on both local economic conditions and the parameters of a localities inclusionary zoning ordinance. Localities follow HUD's criteria for affordability. A unit is affordable to an income tier if it costs no more than 30% of that tier's pre-tax income. Income tiers are defined as a share (δ) of a metro area's median income (AMI_m) as calculated by HUD. Moderate income tenants make no more than 125% of the AMI_m. Low income tenants make no more than 80% the AMI_m. Very and extremely low tenants earn 50% and 30% the AMI_m respectively. All else equal, policies which require units be affordable to lower income categories require land owners to forgo more rent, and are therefore more stringent. Policy makers refer to the tier of income an ordinance requires units be made affordable to as the "depth" of the policy.

Inclusionary zoning ordinances also specify the "breadth" of affordability, or the share of all built units which must be set aside ($\eta_{ji\delta}$). All else equal, the greater the breadth of an inclusionary ordinance, the more rent a landlord must surrender to comply with the policy. Typically, an ordinance will require that units be set aside for multiple affordability tiers. Usually many units are required to be set aside for low income tenants, and a small number of "deeply affordable" units must be set aside for very low income tenants.

In about half of jurisdictions, developers are given a menu of breadth and depth combinations from which they may choose. For example, a city might require a developer to set aside either 10% of units to low income tenants or 5% to extremely low income tenants. In order to assign a policy to a neighborhood, I assume that developers select the set of affordability requirements that minimizes their forgone rent. Table 1 summarizes inclusionary ordinances policy parameters in each neighborhood. The top panel shows parameters at the ordinance by neighborhood level. This summarizes the menu of compliance options available to developers. The bottom panel summarizes parameters of the least-stringent policy option in each neighborhood. In aggregate, wider but shallower compliance options (in which developers must set aside a large number of units for moderate income tenants) minimizes rent forgone. Developer decisions could differ from my assumed decision rule. There may be unobserved costs to maintaining moderate income units (perhaps higher vacancy), which would cause developers to opt for narrower and deeper compliance options.

The stringency of an affordable policy also depends on the market rent a landlord could charge for the unit, were it not subject to an affordability constraint. Note that this component of the inclusionary zoning ordinance is hyper-local, varying even within a jurisdiction. If a jurisdiction applies a uniform inclusionary ordinance across its entire area, the policy will be more stringent in high rent areas than in low rent areas. I use HUD's Small Area Fair Market Rent (FMR_n) to proxy median hyper-local market rents[2]. One could argue this measure may be too low; new units ought lease for more than the area median. However the measure is widely available, and sometimes used by localities to calculate the in-lieu fee it charges developers who are unable to actually construct the income-restricted units. Furthermore, as I will soon show, most of the variation in policy is due to jurisdiction's choice of policy parameters not local economic conditions.

$$\tau_{pn} = \min_i \sum_{\delta \in \{0.3,0.5,0.8,1.25\}} \frac{1.25 \times SAFMR_n - \eta_{pi\delta} \left(\delta \frac{3}{10} \frac{1}{12} AMI_m \right)}{1.25 \times SAFMR_n}$$

$$\begin{aligned}
\tau_{pn} &= \text{Share of rent forgone to comply with the Inclusionary zoning ordinance} \\
\delta &= \text{Share of median income defining an affordability tier} \\
FMR_n &= \text{2017 Small area fair market rent in a neighborhood (ZCTA)} \\
\eta_{pi\delta} &= \text{The share of units required to be set aside for income tier } \delta \\
AMI_m &= \text{The area median income for a metro area.} \\
m &= \text{index for metro areas} \\
p &= \text{index for jurisdiction or ``place''} \\
n &= \text{index for neighborhood} \\
i &= \text{index for inclusionary zoning ordinance options developer may select from}
\end{aligned}$$

The expression for inclusionary zoning ordinance stringency is above. It summarizes the cost to developers of complying with an ordinance in each neighborhood, taking into account the breadth and depth of the policy. It can be relatively easily extended to capture other elements of inclusionary ordinances. If an ordinance allows for off-site provision of affordable units, the FMR_n simply becomes $\min_n FMR$. To capture variation in the duration of affordability requirements, one can

[2] HUD's Small Area Fair Market Rents estimate the 40th percentile rent. To proxy the local median rent, I use 125% the Small Area Fair Market Rent.

add a discount rate and sum over the required number of years. However, the below simple form captures the most important variation in inclusionary zoning policy stringency. Note that while the above expression gives the inclusionary zoning ordinance stringency for a specific neighborhood, for much of the subsequent analysis I take a population weighted average over a jurisdictions neighborhoods to arrive at an average stringency for the jurisdiction as a whole (τ_p).

Figure 1 maps the stringency of inclusionary zoning ordinances across the state of California. It is already clear that there is a high degree of variation in policy stringency. There are some localities which have an ordinance which does not bind. This means developers have the option to set aside units to be rented at rates above the local market rent. On the other hand, several jurisdictions have quite stringent policies, requiring property owners to forgo more than 7% of total rent to comply with the policy. The map also shows the variety of jurisdictions which adopt an inclusionary ordinance. The range of treated jurisdictions includes principal cities like San Diego and San Francisco, but also rural areas covered by county ordinances. The most enthusiastic adopters of IZ ordinances appear to be major outer suburbs like Dublin in the East Bay Area or Irvine in Orange County.

Figure 2 shows the distribution of IZ policy stringency conditional on a location having an IZ policy. Panel (a) shows the distribution at the neighborhood level. There are many neighborhoods in which the local ordinance does not bind. Panel (b) shows policy stringency aggregated up to the jurisdiction level. At the jurisdiction level, the distribution of stringency is roughly uniform between 0% and 8% of rent forgone, with a few high stringency outliers.

Table 2 shows demographic and economic characteristics of jurisdictions binned by their inclusionary zoning ordinance stringency. Jurisdictions with more stringent policies have higher household incomes, higher rents, and higher home values. They also have higher educational attainment, a higher share of white residents, and higher democratic vote shares. There do not appear to be differences in the age of their housing stock or their shares of renter or owner occupied housing units. Perhaps most interestingly, as one moves across the columns from non-IZ adopting places to the most stringent group, the Wharton Residential Land Use Regulation Index increases monotonically. Housing unit supply elasticity as measured by Baum-Snow and Han (2024) generally decreases with policy stringency. Jurisdictions with more stringent policies also have a harder development environment for both policy and non-policy reasons.

There are two potential sources for the variation in policy stringency. As I showed in the expression for τ_{pn}, IZ stringency depends on ordinance parameters and local economic conditions. I compare these two sources of variation in Figure 3. On the x-axis, I recalculate the stringency of each jurisdiction's inclusionary zoning ordinance, not based on their actual local economic conditions, but using state averages. This isolates the stringency due to the jurisdictions policy parameters. On the y-axis, I show the realized stringency based on local rents and incomes minus the recalculated stringency using state averages. This isolates the effect of local economic conditions on policy stringency. Center dots show a population-weighted average of a jurisdictions' neighborhoods. Vertical lines show the range between the 25th and 75th percentile of a jurisdiction's neighborhoods. The plot shows a positive relationship between two sources of policy variation. The jurisdictions which adopt stringent policy parameters (deep or wide set-aside requirement) are also the jurisdictions with economic conditions that make a given IZ policy more costly to developers (high market rent to area median income ratios).

Estimation

I estimate the causal effect of inclusionary zoning ordinance stringency on housing production using a difference-in-difference strategy. While prior literature has differenced around the endogenous adoption of an inclusionary ordinance by a city, I instead difference around a statewide legal event. In 2017, the California legislature passed the "Palmer Fix" which reactivated dormant inclusionary zoning ordinances. It overturned a 2009 court ruling which found inclusionary zoning requirements illegal under the state's rent control laws. This meant that on-the-books inclusionary ordinances which had gone unenforced from 2009-2017 were enforceable starting in 2018. The simultaneous reactivation of local inclusionary zoning policies allows me to perform a conventional two-period difference-in-difference, sidestepping the econometric complexities of the staggered case (Goodman-Bacon, 2021). My baseline specification includes the continuous measure of stringency aggregated to the jurisdiction level (τ_p) described above. $Post_t$ has a value of one for observations in 2018 and onward. Year (ω_t) and jurisdiction (γ_p) fixed effects are included. The baseline specification is presented below.

$$ln(UnitsPemitted_{pt}) = \beta\,(Post_t \times \tau_p) + \gamma_j + \theta_t + \epsilon_{pt}$$

I aim to capture the effect of the policy on latent residential construction, so I weigh observations by pre-period permitting activity. Standard errors are clustered at the jurisdiction level.

I estimate several additional models for context and robustness. I substitute the $Post_t$ indicator for a set of year dummies, producing the event study to check for pretrends. My estimation strategy relies on parallel trends between treated and untreated jurisdictions, and for which parallel trends ahead of treatment provide suggestive evidence.

Cities without inclusionary zoning ordinances serve as controls in the baseline specification. They stand in for the counter-factual where cities with inclusionary zoning ordinances never have them reactivated. However, as one can see in Table 2, there are systematic differences between cities which adopt inclusionary ordinances and those that don't. This is an issue for identification if these differences lead to a break in parallel trends around the time of treatment. I leverage the richness of the Housing Element data to improve the analysis. As a robustness exercise I restrict the sample to jurisdictions which deciding on, planning to implement, or had already implemented an inclusion zoning ordinance in their final housing element before the Palmer decision. This includes cities which had intended to implement an inclusionary ordinance, but abandoned the effort when they were found unenforceable. These are natural controls for cities which did in actuality adopt the policy.

I also run several models in which I discretize treatment. Recent work has argued against conventional continuous treatment difference-in-difference estimation (Callaway et al., 2024). I show robustness of my headline results by comparing discrete treatment bins to untreated jurisdictions, and high stringency jurisdictions to low stringency jurisdictions. I also plot event studies of these discrete specifications, allowing me to provide suggestive evidence for the strong parallel trends assumption: counterfactual trends are parallel between each level of treatment intensity. This is needed to interpret continuous difference-in-difference specifications causally.

For additional robustness, I re-estimate my headline regression using alternative weighting: no weights, population weights. I also re-estimate using alternative estimation strategies: Poisson Pseudo Maximum Likelihood, Negative Binomial, and the procedure recently proposed by Callaway et al. (2024))

I extend my analysis by looking at how inclusionary zoning ordinances impact the composition and location of residential construction. I rerun my headline analysis with alternative outcomes like the assessed value of constructed units and the share of units built which are in multifamily structures. I then estimate my models at the neighborhood level using CoStar data on the construction of large multifamily developments. I add jurisdiction fixed effects, estimating effects across neighborhoods of the same jurisdiction. This makes use strictly of of intra-jurisdictional variation in stringency driven by local rents. Figure 9 maps variation in stringency across San Francisco neighborhoods as an example.

Finally, I extend my analysis by looking at spillovers. It is possible that developers, faced with stringent policies in one jurisdiction, simply relocate construction activity to neighboring or close-substitute jurisdictions. If this is the case, my headline estimates would be biased downward due to a SUTVA violation. This would have the effect of making the impact of inclusionary ordinances appear larger than they are in reality. First, I measure the impact of inclusionary ordinances only on permitting in jurisdictions with no ordinance. Log permit rates are regressed on the average stringency of the neighboring jurisdiction's IZ policies. In addition to checking for spillovers across geographic neighbors, I look for spillovers across jurisdictions which are close substitutes in terms of labor market access. Jurisdictions are close substitutes if the places their current residents commute to are correlated. As an illustrative example, I map the neighboring and close-substitute jurisdictions of Riverside in Figure 8. After checking for spillovers in non-IZ jurisdictions, I estimate spillovers to all jurisdictions by regressing permit counts on weighted averages of neighbor's policy stringency, a jurisdictions own stringency, and the interaction of the two.

5 Results

The headline result and key robustness checks are reported in Table 3. In column (1), I estimate the effect of an inclusionary zoning ordinance's stringency on housing production. I find that for every percentage point increase in the stringency of an inclusionary zoning ordinance, the flow of new housing supply decreases by 7.6%. For context, a typical binding inclusionary ordinance requires a developer to forgo about 4.2% of rent, implying a 31.8% reduction in housing production.

The event study version of this baseline specification is displayed in Figure 4. There is no significant pre-trend leading up to the 2017 reactivation of inclusionary policies, with a clear decrease in housing production post-treatment. Note that I restrict the pre-period to 2014-2017. This is to avoid the effects of the original Palmer decision and the Great Recession. An extended event study with coefficients plotted back to 2003 is reported in Figure 5. Housing production in cities with inclusionary zoning ordinances is elevated in 2010 through 2013. This is likely because these localities had inclusionary zoning ordinances until 2009, and developers responded to the elimination of the policy with more production. Note that these years are excluded from the baseline analysis. Additionally, before the original Palmer Decision, housing production in IZ jurisdictions is lower than jurisdictions without IZ policies, mirroring the effects I find in the post-period. While trends do not look parallel around and before the 2009 Palmer decision, parallel trends were not to be expected in this period as localities were not experiencing parallel policy environments.

In column (2) of Table 3, I augment my continuous measure of inclusionary zoning ordinance stringency with a simple indicator for the presence of an inclusionary policy. I find no extensive margin effect. All of the effect on housing supply loads onto the measure of IZ policy intensity. This highlights the importance of capturing the stringency of inclusionary policies. In column (3) I restrict my sample to only localities with an inclusionary zoning ordinance. I am able to estimate significant effects of a similar magnitude to my headline result relying only on the variation in intensity across adopting localities. Finally, in column (4) I restrict the sample to localities which, prior to the Palmer decision in 2009, had either already adopted an inclusionary zoning ordinance or indicated in their housing element that they were planning to implement one. On that sample, I estimate a result quite similar to the headline result in column (1). The final two models provide reassurance that the detected effect on residential housing supply is not driven by selection into inclusionary ordinance adoption.

Robustness

In Table 4, I discretize the stringency of jurisdictions' inclusionary zoning ordinances in various ways. In column (1) I regress log permits counts on a indicator that a jurisdiction has an inclusionary ordinance, finding a estimated 33% reduction in permitted units, though with less precision than my headline results. In column (2) I group inclusionary ordinances into 4 bins by stringency. I find an economically large, but statistically insignificant, 15% decline in housing production in localities with lower two bins ($\tau_p < 5\%$), but very large and statistically significant effects for policies in the higher two bins ($\tau_p \geq 5\%$). In column (3) I compare policies above and below the mean policy against the non-adopting localities. In column (4) I compare jurisdictions with above median policies to jurisdictions with below median policies, restricting the sample to only jurisdictions with some inclusionary ordinance. In both cases I find that adopting a very stringent inclusionary ordinance leads to an over 50% decline in housing production. Overall, this table shows that a) my headline results are robust to various sensible ways of discretizing the analysis and b) identification of effects are driven by localities with more stringent inclusionary ordinances.

Another reason to run discrete models is to plot their corresponding event studies. Looking at pre-trends across groups of different treatment intensities gives suggestive evidence for the strong-parallel trends assumption required for continuous difference-in-difference. In Figure 6, I plot the event study version of the model listed in Table 4 column (2). It shows the trajectory of residential permitting in four bins of inclusionary zoning relative to non-adopting jurisdictions. Reassuringly, pre-period trends look parallel to the untreated jurisdictions for all groups. Each group therefore appears to be parallel with one another. I test this more directly in Figure 7, where I plot the event study version of the model presented in Table 4 column (4). Here, I show the trajectory of residential permitting for jurisdictions with above median stringency inclusionary ordinances relative to jurisdictions with below median stringency inclusionary ordinances. The sample is restricted to jurisdictions with an inclusionary ordinance. There appears to be parallel trends in high and low treatment intensity jurisdictions ahead of the 2018, with a sharp divergence afterward. Because trends appear parallel across different levels of treatment, and between treated and non-treated jurisdictions, I am confident in assessing my continuous difference-in-differences estimates as causal.

Although parallel trends hold between IZ localities of various stringency and the non-IZ localities, they clearly have different characteristics as described on Table 2. This raises the concerning possibility that estimated treatment effects are driven by some coincidental divergence in trends based on the observable differences between high stringency IZ, low stringency IZ, and non-IZ localities. To mitigate this concern, I test the robustness of my discretized treatment models to an inverse propensity weighting strategy. In Table 5, I estimate the propensity of localities to adopt inclusionary zoning using 2013-2017 ACS demographic and housing stock characteristics, labor and housing market conditions, and 2016 Democratic vote share. For exposition, I first estimate a simple linear model in column (1). In column (2) I estimate a logistic model with an indicator for adopting any IZ policy on the left hand side. In column (3), I estimate an ordered logistic model with quartiles of IZ stringency on the left hand side. With this model, I am able to estimate the probability of a locality adopting a IZ ordinance at each level of stringency. I then use these probabilities to reweight non-IZ adopting localities to serve as a more similar set of controls for each quartile of treated IZ localities. In practice this means upweighting control jurisdictions with the higher population, higher income, high democratic vote share, and older housing stock. A seperate model is estimated for each treatment quartile in Table 6. Estimates utilizing inverse propensity weighting are slightly smaller, but indistinguishable from the discrete treatment estimates in Table 4.

I test the robustness of my results to various alternative estimation strategies in Table 7. In column (1), I weight jurisdictions by population, rather than pre-treatment build rates. The estimated effect on housing supply is slightly smaller. It also has a subtly different interpretation. It represents the effect for the average person residing in an IZ jurisdiction, rather than the average effect on latent construction. In column (2) I rerun my headline regression unweighted. Effects are small and insignificant. This is likely driven by the large number of small jurisdictions with sparse building permit data. In jurisdictions where building is rare, changes due to policy are hard to discern from noise.

In my headline regressions I use log permit counts as my outcome. This means jurisdiction-years with 0 building permits are dropped from the sample. This represents about 9% of the sample unweighted, but only 0.7% of the sample weighted by pre-period permit counts. To confirm that these dropped zeros do not drive my results, I instead implement Poisson and Negative Binomial count models. These are displayed in columns (5) and (6) of Table 7. My research setting does not naturally lend itself to conventional Poisson estimation. Poisson regression is suited for outcomes which are counts of independent events. Since permitted units arrive in clusters, they will not follow the distributional assumptions required for conventional Poisson estimation. For this reason I estimate the Poisson model by pseudo-maximum likelihood, which is robust to the overdisperson likely present in my main outcome. The negative binomial model handles overdispersion parametrically, estimating a statistically significant overdispersion parameter. Both count models produce estimates similar to my headline specification.

Finally, I estimate a continuous difference-in-difference using the estimation strategy suggested by Callaway et al. (2024). Their preferred average treatment effect on the treated (ATT) and average causal response on the treated (ARCT) are presented in Table 5 column (6). Their preferred ARCT measure is quite similar to my headline coefficient. Note that their estimator requires a balanced panel, so as a comparison I repeat my headline regression in Table 7 column (6) with the same balanced sample restriction. Overall, my head line result is robust to a range of sensible estimation strategies.

Extensions

I now turn my attention to estimating how inclusionary zoning effects the composition and distribution of new housing. In Table 8 I estimate how the assessed value of permitted units, the share of units which are in multifamily structures, and the share of units in large multifamily structures responds, finding no effects. Declines in housing production appear to be uniform across housing type.

I use CoStar data with the precise location of new large multifamily developments re-estimate my results at the neighborhood level. I find a statistically significant 9% reduction in CoStar units for each percentage point increase in inclusionary zoning stringency. This model is presented in Table 8 column (4). I then add jurisdiction by year fixed effects, so that the estimation is conducted entirely within place. These place by year fixed effects absorb most of the variation in inclusionary zoning stringency. A cross-sectional neighborhood level regression of stringency on jurisdiction fixed effects yields an R^2 of 0.85. Nonetheless, with some of the remaining variation in policy stringency seems to induce a response from developers. Within jurisdiction, neighborhoods with higher stringency see a statistically significant decrease large multifamily residential development relative to lower stringency neighborhoods. Recall that within locality, neighborhood to neighborhood variation in stringency is driven entirely by differences in market rents. This means that inclusionary zoning is disproportionately restricting development in the higher rent neighborhoods of cities.

I estimate potential spillovers from jurisdiction to jurisdiction Table 9. As shown in column (1), there is no detectable effect on residential construction on untreated jurisdictions whose neighbors have inclusionary zoning ordinances. In column (2) rather than looking for effects from adjacent jurisdictions, I look at effects from jurisdictions who are close substitutes in terms of labor market access. Still, no spillover effects are found. If anything, point estimates suggest a slight decrease in housing production. Spillover estimates remain insignificant when spillovers are allowed between treated jurisdictions, as shown in columns (3) and (4). This shows that the headline results are not being inflated by positive spillovers onto untreated units.

6 Quantitative Sorting Model

I calibrate a quantitative urban model in the style of Ahlfeldt et al. (2015) in which residents choose their housing consumption (l_{ijo}^{θ}), location of residence (indexed by i), and location of work-site (indexed by (j) from a set of discrete geographies. Geographies have heterogeneous amenities (B_i), productivities (A_j), and inclusionary zoning policy stringency (τ_i). In order to evaluate the effect of inclusionary zoning on different parts of the income distribution, I extend the basic model in the style of Diamond and Gaubert (2022), with two types agents, skilled and unskilled $\theta \in \{U, S\}$. The types are heterogeneous in a) their productivity (cob douglas production parameters a and b) b) their consumption share of housing (α^{θ}) and c) their probability of acquiring non-market housing (λ).I advance the literature on quantitative spatial models by adding a realistic inclusionary zoning to the model. In locations with an IZ ordinance, developers must build non-market rate units at the ordinance specified rate. The below market rate housing is then rented at a discount to workers of the type with lower earnings.

Production and trade is simplified relative to richer Quantitative Spatial Models. Firms produce a non-differentiated frictionlessly tradable good using low skill and high skill labor, using a Cobb-Douglas CRS technology. There is a skill neutral workplace specific exogenous productivity, which is a fundamental of the neighborhood.

$$Y_j = A_j \left(\left(L_j^U \right)^a + \left(L_i^S \right)^b \right)$$

A firm located in j then solves the problem

$$\max_{L_j^U L_j^S} A_j \left(\left(L_j^U \right)^{\frac{\rho-1}{\rho}} + \zeta \left(L_i^S \right)^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}} - w_j^U L_j^U - w_j^S L_j^S$$

Worker's consume the perfectly tradable good (c_{ijo}^{θ}) and an untradable good (housing) h_{ijo} with Cobb-Douglas utility. In addition, they derive utility from local amenities B_i and disutility from commuting d_{ij}. As a proxy for non-homothetic preferences in housing, I follow Diamond and Gaubert (2022) and allow the expenditure share on housing to differ by skill group $\alpha^U > \alpha^S$. This will drive workers to sort differently across the two skill groups. The idiosyncratic utility term ζ_{ijo} is drawn from a Frechet distribution with scale parameters $Ti, Ej > 0$ and shape parameter $\epsilon > 0$:

$$U^{\theta}_{ijo} = \frac{B_i \zeta_{ij}}{d_{ij}} \left(\frac{c^{\theta}_{ijo}}{1-\alpha^{\theta}} \right)^{1-\alpha^{\theta}} \left(\frac{h^{\theta}_{ijo}}{\alpha^{\theta}} \right)^{\alpha^{\theta}} \qquad \text{where } F(\zeta_{ijo}) = \exp(-T_i E_j \zeta^{-\epsilon}_{ijw})$$

Workers select their residence i, their worksite j and their consumption bundle $\{c^{\theta}_{ijo}, h^{\theta}_{ijo}\}$ to maximize their utility subject to their budget constraint $w^{\theta}_j = c^{\theta}_{ijo} + r_i h^{\theta}_{ijo}$. Based on the properties of the Frechet distribution, the probability of selecting to live in location i and commute to location j (absent inclusionary zoning is)

$$\pi^{\theta}_{ij} = \frac{T_i E_i \left(d_{ij} r_i^{\alpha^{\theta}} \right)^{-\epsilon} \left(B_i w^{\theta}_j \right)^{\epsilon}}{\sum_i \sum_j T_i E_i \left(d_{ij} r_i^{\alpha^{\theta}} \right)^{-\epsilon} \left(B_i w^{\theta}_j \right)^{\epsilon}}$$

Conditional on living in i, the probability of commuting to j and visa versa

$$\pi^{\theta}_{ij|i} = \frac{E_j \left(d_{ij}\right)^{-\epsilon} \left(w_j\right)^{\epsilon}}{\sum_j E_j \left(d_{ij}\right)^{-\epsilon} \left(w_j\right)^{\epsilon}} \qquad \pi^{\theta}_{ij|j} = \frac{T_i \left(d_{ij} r_i^{\alpha^{\theta}} \right)^{-\epsilon} \left(B_i\right)^{\epsilon}}{\sum_i T_i \left(d_{ij} r_i^{\alpha^{\theta}} \right)^{-\epsilon} \left(B_i\right)^{\epsilon}}$$

This conditional commute decision is relevant in our context because the low-earning workers who receive inclusionary units (who are of mass λN^U) do not select residential locations, only their workplace. The total number of workers in a location then is the sum of high skill workers selecting that location ($\pi^S_i N^S$), the low skill workers who are not provided an inclusionary unit which select that location ($\pi^U_i (1-\phi) N^U$), and the amount of inclusionary housing produced at a location, distributed at the same density as market rate units in that location $\left(\nu_i H_i \frac{\pi^S_i N^S + \pi^U_i (1-\phi) N^U}{1-\nu_i H_i} \right)$

$$\begin{aligned} N_i &= \pi^S_i N^S + \pi^U_i (1-\phi) N^U + \nu_i H_i \frac{\pi^S_i N^S + \pi^U_i (1-\phi) N^U}{1-\nu_i H_i} \\ &= \frac{1}{1-\nu_i} \left(\pi^S_i N^S + \pi^U_i (1-\phi) N^U \right) \end{aligned}$$

Land is supplied by atomistic absentee landlords and the aggregate land supply function for a city is $\Lambda_i = \bar{\Lambda}_i \left(p_i^{\Lambda}\right)^{\eta_i}$. Perfectly competitive property developers convert land into housing using a linear production function. $H_i = K\Lambda_i$. The developer statutorily bear the cost of the Inclusionary Zoning requirement. Here, I use a simplified version of the policy in which a share of units ν_i must be set aside as affordable for Low Income Tenants, defined as 80% the Area (county) Median Income[3]. For these units, a developer may not charge more than 30% of 80% median income. Therefore, the total share of rent foregone by the landlord is the share of units which must be set aside times the difference between the market and inclusionary rents.

$$\tau_i = \nu_i \left(\frac{r_i - (0.3)(0.8)\text{med}(w_{ijo})}{r_i}\right)$$

In the presence of an inclusionary zoning requirement, the developer's problem becomes

$$\max_{\Lambda} \left(1 - \tau_i\right) r_i K \Lambda - p_i^{\Lambda} \Lambda$$

Which yields a first order condition stating the price of housing must be proportionate to the price of land $p_i^{\Lambda} = (1 - \tau_i) K r_i$. Plugging this, along with the production function back into the aggregate land supply function, we get the housing supply function

$$K^{-1} H_i = \bar{\Lambda}_i \left((1 - \tau_i) K r_i\right)_i^{\eta}$$

$$H_i = \bar{\Lambda}_i K^{\eta_i - 1} (1 - \tau_i)^{\eta_i} r_i^{\eta_i} = \bar{\Lambda}_i K^{\eta_i - 1} \left(1 - \nu_i \left(\frac{r_i - (0.3)(0.8)\text{med}(w_{ijo})}{r_i}\right)\right)^{\eta_i} r_i^{\eta_i}$$

In other words, housing supply is increasing in the available land in a location Λ_i, the building technology K, the price r_i, and the elasticity of land supply η_i, which in this simplified model is also the elasticity of housing supply. Housing supply is decreasing in the stringency of the inclusionary zoning requirement τ. The stringency of inclusionary zoning depends on wages, market rents, and the local set aside requirement.

Note that in the presence of inclusionary zoning, total housing supply must be split into market

[3] In practice, inclusionary zoning may include set asides for extemely low, very low, and moderate income tenants as well. However, in a model with only two worker types, multiple tiers does not enrich the model much. To capture some of the variation in stringency from cities choice of IZ policy "depth" I could adjust the "breadth" of a city's IZ policy until the stringency matches what we observe in practice. However, for now the simple total of all unit set aside requirements is used, acting as if they were all set aside at the low income tier

rate housing $H_i^m = (1 - \nu_i)H_i$ and income restricted housing $H_i^r = \nu_i H_i$. The restricted housing is distributed to residents in the lower earning group.

Equilibrium in this model can be described as a set of wages $\{w_j^U, w_j^S\}$, rents $\{r_j\}$, location/commute choice probabilities $\{\pi_{ij}^U, \pi_{ij}^S\}$, and lottery probability $\{\lambda\}$ such that worker and firms have optimized and markets clear. The market clearing conditions are

The equilibrium is pinned down by the above conditions, along with profit maximization and zero profit in the housing development and tradable goods sectors, the free mobility condition across California, and exogenously determined populations $\{N^U, N^S\}$ within California. For the "open-city" version of the model, free mobility between California is assumed, with endogenous population of low and high earning workers pinned down by a pair of exogenous outside utility levels $\{\bar{U}^U, \bar{U}^S\}$. The reality is somewhere in the middle, with workers being able to leave California, through with great cost. Therefore, the welfare effects of IZ will be bound between zero (the open city case, where the only margin of adjustment is population) and the closed city welfare estimates.

The model is calibrated to a no Inclusionary Zoning baseline using 2017 data. Commute patterns by skill are measured in LODES[4]. Rents are measured using 2017 small area fair market rents from HUD where available. Otherwise county level fair market rents are used in rural areas.

The parameters for the model are listed in Table 10 alongside their estimation strategy. The production share of high and low skill labor are calibrated to the wage bill across skill groups as measured in the ACS public use microdata for 2017[5]. Low and high skill housing shares of consumption (α^L and α^H) are also estimated from the ACS by taking the person weighted gross rent to household income ratio over respondents by skill.

A partial commuting elasticity ($\kappa^\theta \epsilon^\theta$) is estimated from gravity regressions of commute frequencies on drive times between residence-workplace pairs as well as residence and workplace fixed

[4]LODES does not directly provide commuting flows by skill level. LODES has: a) Education×Wage(bin)×Residence b) Education×Wage×Workplace c) Wage×Residence×Workplace counts. What I need but do not directly observe is d) Education×Workplace×Residence counts. I use iterative proportional fitting to find a set of Education×Wage×Residence×Workplace counts which are consistent with margins (a-c), and then collapse the wage bins to get a proxy for (d). The margins do not provide enough constraints to uniquely identify d). The algorithm converges to the margin-consistent set of flows that is closest to the prior I give it. For now, I am using the prior that 80% of the high wage and 20% of the low wage (¡$40k) flows are college educated 30+, and visa versa for low skill workers. Zeros are persistent through the fitting algorithm, so I pick a prior that puts at least some flow in all plausible cells. This approach incorporates the maximum amount of information from LODES to back out flows by skill level.

[5]To be precise, the shares are estimated by first finding the wage for each skill group by each place of work public use microdata area, then taking a weighted average over the place of work public use microdata areas using commuter counts from LODES as weights.

effects.

$$\pi_{ij}^{\theta} = \left(\kappa^{\theta}\epsilon^{\theta}\right)(DriveTime)_{ij} + \gamma_i + \zeta_j + u_{ij}$$

The ϵ^{θ} parameters can then be separately estimated using the variance matching strategy from Ahlfeldt et al. (2015). This allows me to back out the κ^{θ} parameters as well.

Finally, there are two neighborhood specific provided parameters. Housing supply elasticities by neighborhood are borrowed from Baum-Snow and Han (2024). The IZ set-aside requirement is supplied form my housing element data.

The model's geography is California "neighborhoods" with neighborhoods being unique intersections of ZCTAs and localities, with localities being municipalities or counties net of self governing municipalities.

Once the model is calibrated to the 2017 pre-IZ equilibrium, we can retrieve the neighborhood specific productivity and amenity values. These are plotted in Figure 10. Additionally, we can see spatial sorting across workplaces and residence by skill level in the baseline equilibrium. This is displayed in Figure 11. Complimentary in production pull workers together, but non-homothetic housing consumption causes them to sort away from one another. As the high income type is less rent sensitive, it will gravitate towards expensive high amenity neighborhoods, while the low skill income type prefers neighborhoods with lower rent for floorspace.

With the baseline calibration completed, I can run counterfactuals by setting the number of units required to be set aside by the IZ policy to be above zero in at least some neighborhoods. First, I try the realistic case, setting IZ requirement in line with local ordinances as they stand today. Some of the effects are plotted in Figure 12. Rents increase in localities with an IZ requirement. The average across all IZ adopting jurisdictions is a 4% increase. Additionally, the share of low skill workers is altered by the IZ ordinance. While the IZ policy creates some units reserved for low income workers, it also increases rents, which they are more sensitive, making the overall effect ambiguous. It appears that in the neighborhoods with the most stringency IZ policies, the rent effect dominates.

For the "closed state" version of the model, where there is no migration in or out of California, I can report welfare effects of the policy counterfactuals. In the first column of Table 11, I report the welfare effects for each skill group of municipal IZ ordinances as they currently stand. Low skill residents who win the option to take an IZ unit are almost 3% better off, while low skill residents who do not win a unit face mild welfare losses due to the rent effect of the policy. Higher income residents face an even smaller welfare loss. They are less sensitive to rents, may benefit from the general equilibrium labor market effects. As a thought experiment, I also try a more extreme counterfactual in which every city in the state adopts a 15% IZ policy. The larger rent effect of this more extreme policy means that even though IZ unit winners are much better off, low skill workers on the whole are worse off. High skill workers actually benefit slightly,again through the general equilibrium wage effects.

This model is a work in progress. Additional mechanisms and counter-factuals should be tested. The model currently lacks agglomeration economies or endogenous amenities, skill bias or otherwise. These are important drivers of sorting, and should be added. Additionally, I aspire to microfound the non-homothetic housing consumption with a neighborhood specific, minimum unit size, reflecting zoning's role in shaping housing consumption patterns.

It is important to keep in mind the different time scales studies by the quantitative model and the reduced form estimates. The model is static. Welfare differences reflect the difference between a world build completely without IZ requirements and a world in which all building occurred under the requirements. The static model is useful for evaluating the long term effects of an inclusionary zoning policy, after most real estate has been redeveloped or renovated under the requirements. The reduced form estimates, by contrast, can only tell us the immediate effects of the policy.

7 Discussion and Conclusion

I evaluate inclusionary zoning ordinances' effect on housing production, finding they cause a large drop in the flow of new housing supply. I advance the literature by using high quality novel data, which allows me to measure the stringency of inclusionary zoning ordinances. I utilize a recent change in superseding state law which provides exogenously timed changes to local inclusionary ordinances, improving upon the staggered approach followed by most prior literature. Finally, I show how inclusionary ordinances alter the geography of housing development within jurisdictions.

There are direct implications of these findings for current California policy. Under the "Palmer

Fix" legislation which reactivated local inclusionary zoning ordinances, state regulators had the power to review and suspend ordinances which prevent cities from meeting market rate housing production benchmarks. Although cities routinely fail to meet their obligations under state law, no such action has been taken to date. By showing that inclusionary zoning ordinances constrain housing supply, these findings should empower state regulators to review ordinances at cities which fail to meet housing goals.

More importantly, there are implications for policy makers hoping to generate affordable housing units and considering inclusionary zoning. Inclusionary zoning creates a trade off between generation of affordable housing units and a substantial decrease in market rate housing production. A back of the envelope calculation can illustrate the magnitude of this supply constraint. Taking the product of the typical IZ stringency across adopting California Cities (4.3pp.), the headline estimated effect on the flow of new housing (7.6%), an off the shelf demand elasticity from Mense (2025) of -0.2, the typical rent of households in IZ jurisdictions before treatment ($1,700), the number of households in IZ adopting cities (just over 1 million) and the number of months in our analysis window (60), I can find that renting households in IZ adopting jurisdictions have paid about $6.97 billion dollars in additional rent due to the supply effects of inclusionary zoning. On the other hand, mandatory inclusionary zoning ordinances have produced approximately 8,990 low income units, during the same window of analysis. This implies a per-affordable-unit cost of $775,000 in terms of excess rents paid by market rate renters. This is substantially more than the marginal cost of directly subsidizing a low income housing unit through the Low Income Housing Tax Credit, which is estimated to be $441,000 in California (Reid et al., 2026). Additionally, policy makers concerned with equity should be particularly concerned with excess rents charged to renters, who are on average lower income. Forgone tax revenue from LIHTC, on the other hand, can be made up for with a progressive tax system. Overall, the hidden costs of inclusionary zoning stack up poorly against alternative means to produce low income housing.

References

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6):2127–2189.

Basolo, V. and Calavita, N. (2004). A critique of the reason foundation study on inclusionary housing policy in the san francisco bay area.

Baum-Snow, N. and Han, L. (2024). The microgeography of housing supply. *Journal of Political Economy*, 132(6):1897–1946.

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

Callaway, B., Goodman-Bacon, A., and Sant'Anna, P. H. (2024). Difference-in-differences with a continuous treatment. Technical report, National Bureau of Economic Research.

Diamond, R. and Gaubert, C. (2022). Spatial sorting and inequality. *Annual Review of Economics*, 14(1):795–819.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of econometrics*, 225(2):254–277.

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

Lewis, P. G. and Marantz, N. J. (2019). What planners know: Using surveys about local land use regulation to understand housing development. *Journal of the American Planning Association*, 85(4):445–462.

Means, T. and 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., and 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.

Ng, Y. R. and Wang, E. S. (2025). Housing affordability, supply, and spatial misallocation.

Phillips, S. (2024). Modeling inclusionary zoning's impact on housing production in los angeles: Tradeoffs and policy implications. *Turner Center for Housing Innovation (ternercenter. berkeley. edu)*.

Powell, B. and Stringham, E. P. (2004). Housing supply and affordability: Do affordable housing mandates work? *Reason Foundation policy study*, 318:1–45.

Schuetz, J., Meltzer, R., and 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.

Soltas, E. J. (2024). The price of inclusion: Evidence from housing developer behavior. *Review of Economics and Statistics*, 106(6):1588–1606.

Wang, R. and Balachandran, S. (2021). Inclusionary housing in the united states: dynamics of local policy and outcomes in diverse markets. *Housing Studies*, pages 1–20.

8 Exhibits

Figure 1: Inclusionary Zoning Stringency by Jurisdiction

(a) Statewide

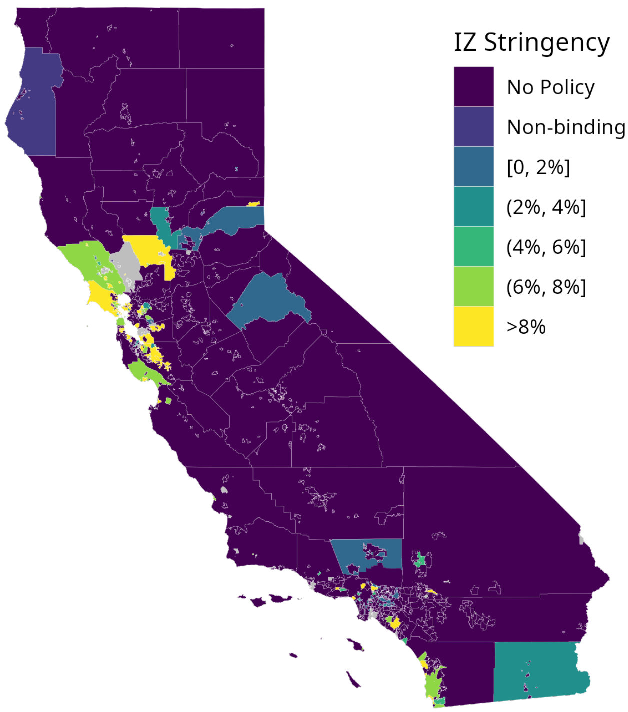

(b) San Francisco Bay Area

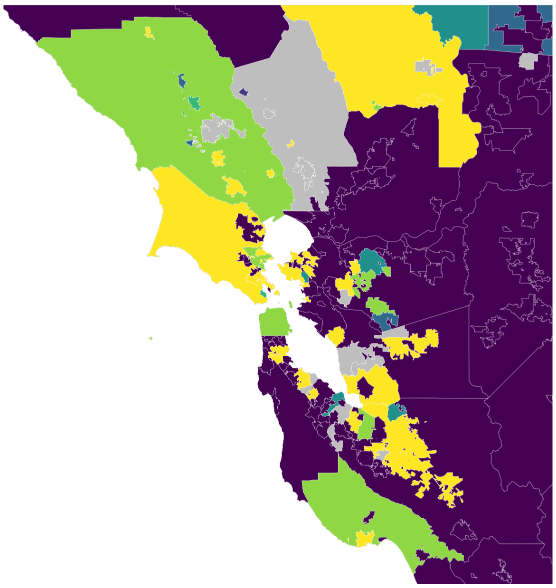

(c) Coastal Southern California

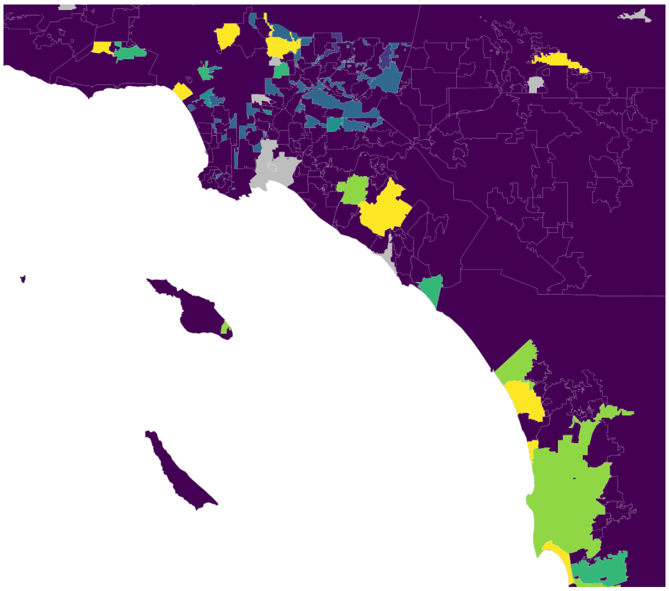

Figure 2: Inclusionary Zoning Stringency Descriptive Figures

(a) IZ Stringency (τ)

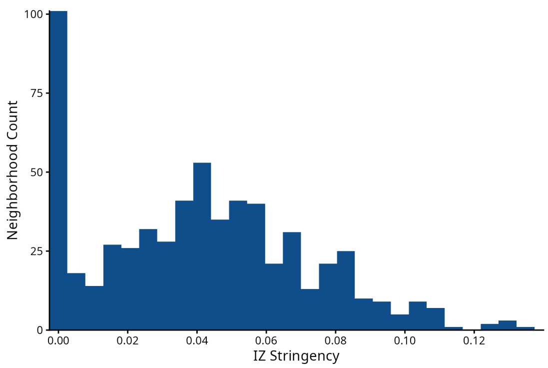

(b) Average IZ Stringency (τ) by Place

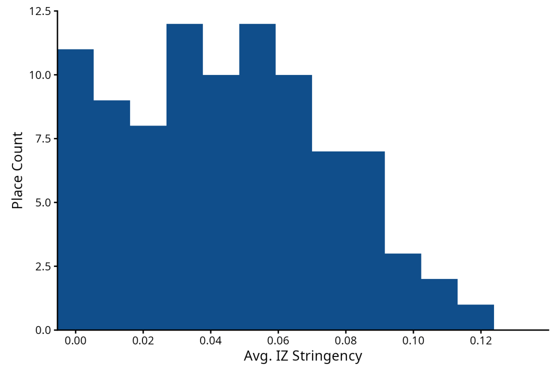

Figure 3: Stringency Decomposition

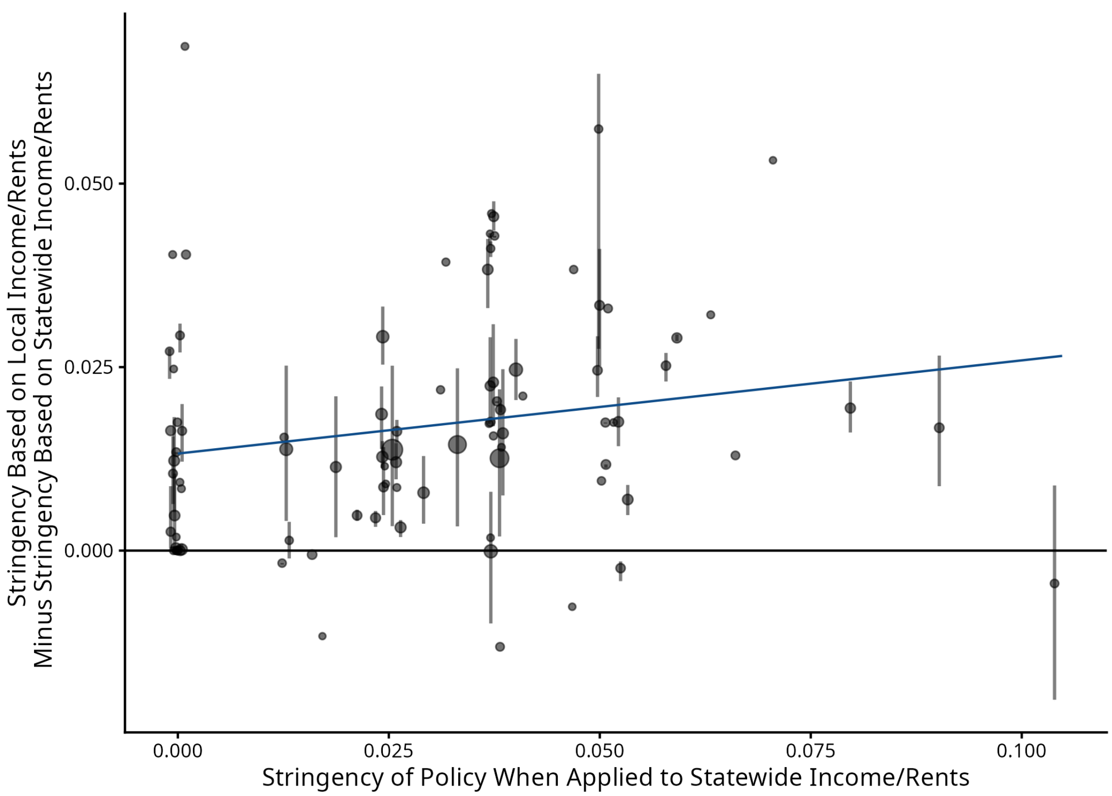

Dots scale proportionate to place population. Vertical error bars show range from 25th to 75th percentile of place zip codes. Line of best fit shown in blue. A small amount of uniform noise has been applied to the x-axis variable for exposition.

Figure 4: Event Study: Effect of IZ Stringency on Permitted Units

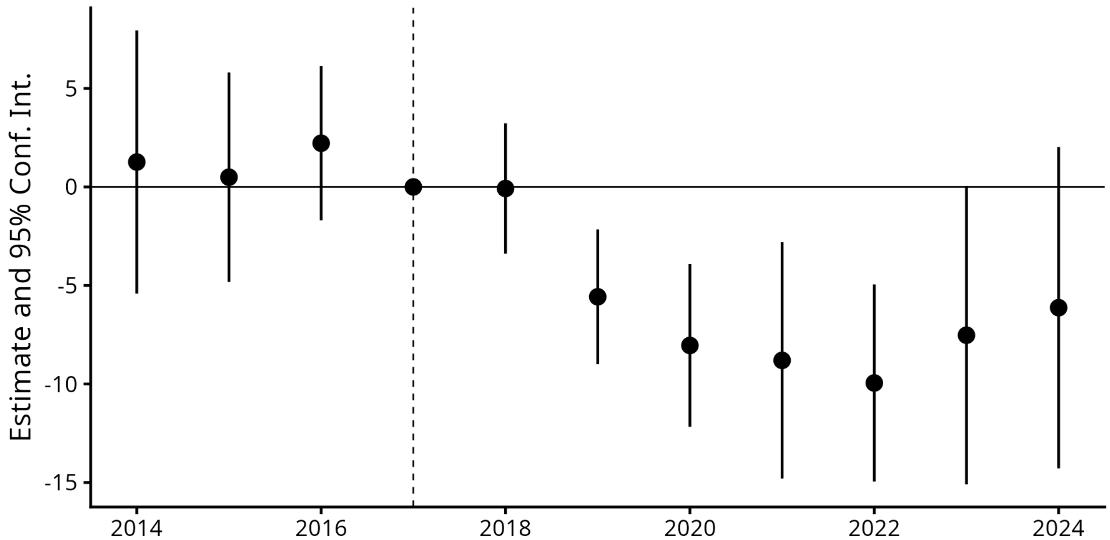

Dashed line marks enactment of "Palmer Fix" legislation, reactivating local inclusionary ordinances. Brackets show 95% confidence intervals. Corresponds to Table 3 column (1), but with extra pre-period estimates shown in grey.

Figure 5: Extended Event Study: Effect of IZ Stringency on Permitted Units

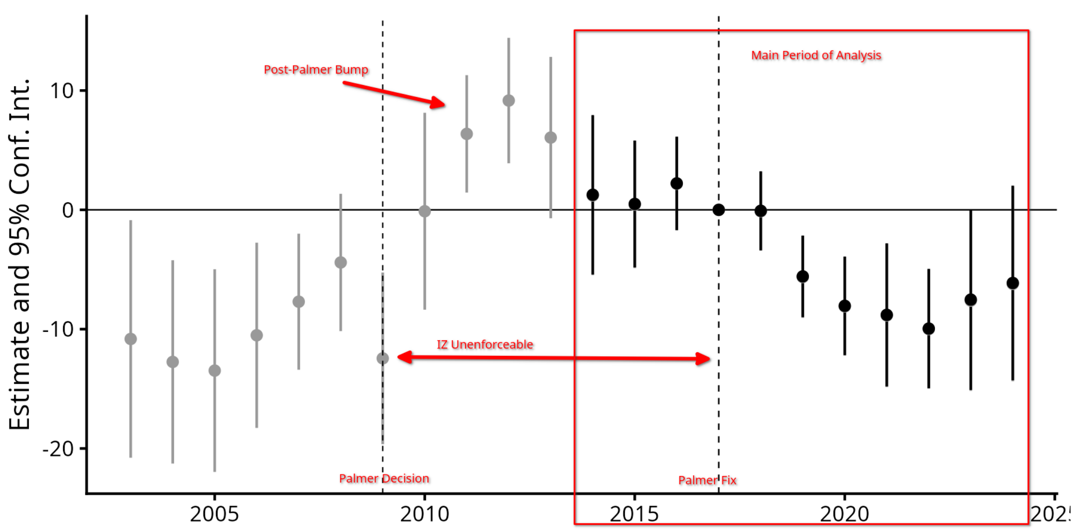

First dashed line marks the Palmer ruling, which made local inclusionary ordinances unenforceable. Second dashed line marks enactment of "Palmer Fix" legislation, reactivating local inclusionary ordinances. Brackets show 95% confidence intervals. Corresponds to Table 3 column (5)

Figure 6: Event Study with Discrete Treatment Levels

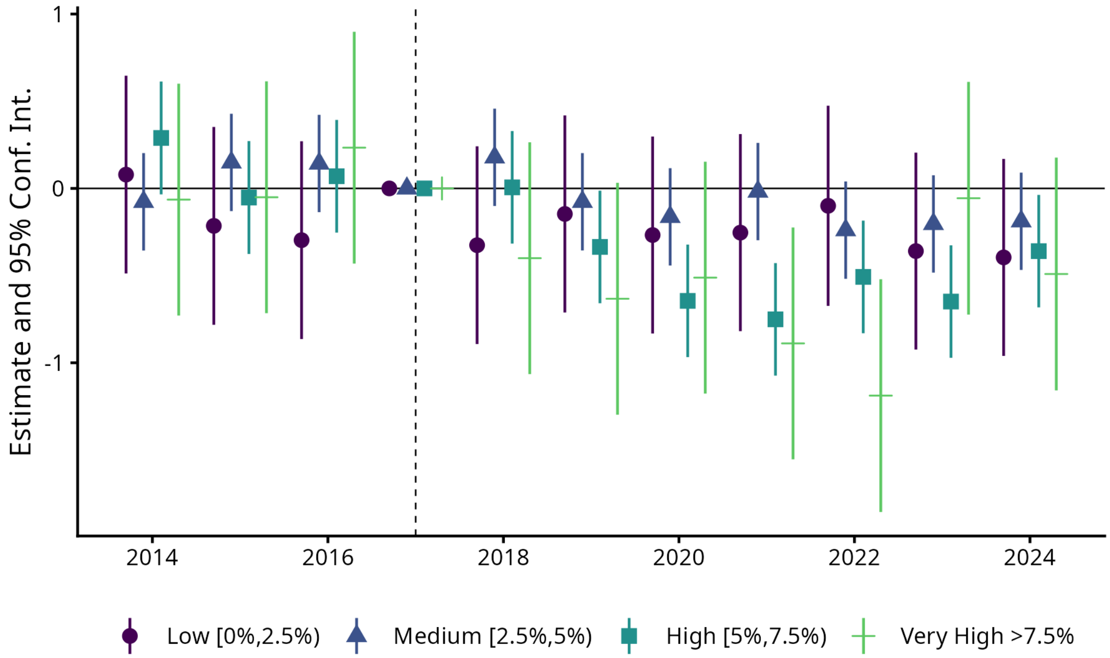

Dashed line marks enactment of "Palmer Fix" legislation, reactivating local inclusionary ordinances. Brackets show 95% confidence intervals. Corresponds to Table 4 column (2). Excluded comparison group are localities without an inclusionary ordinance.

Figure 7: High Stringency vs. Low Stringency Event Study

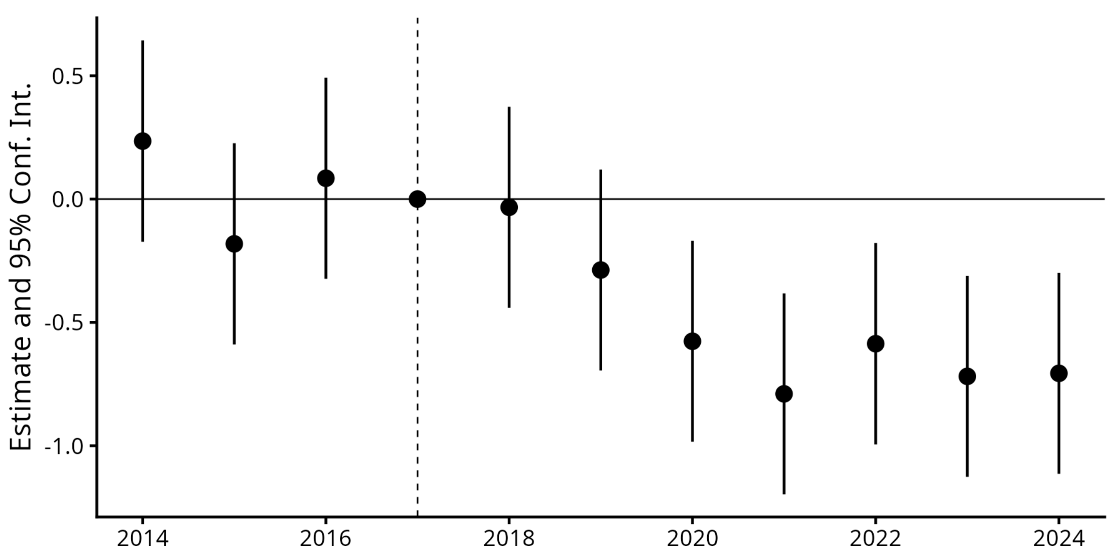

Dashed line marks enactment of "Palmer Fix" legislation, reactivating local inclusionary ordinances. Brackets show 95% confidence intervals. Corresponds to Table 4 column (4). Excluded comparison group are localities without an inclusionary ordinance.

Figure 8: Spillover Weights for Riverside

(a) Simple Adjacency

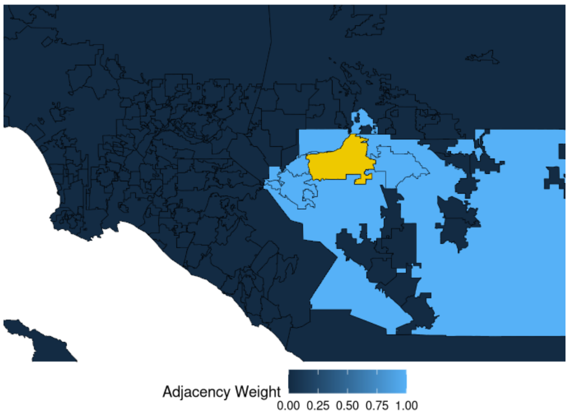

(b) Labor Market Access Substitutes

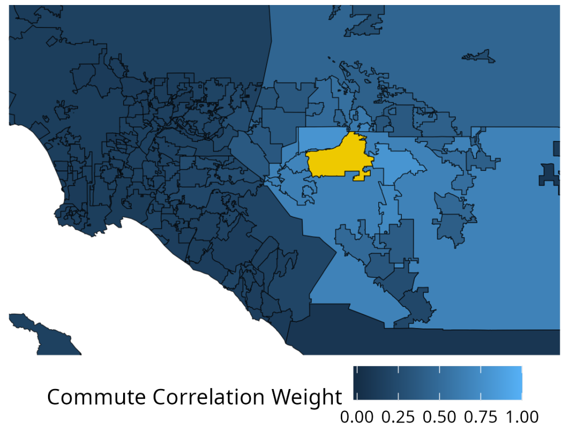

Figure 9: Neighborhood Level Stringency: San Francisco

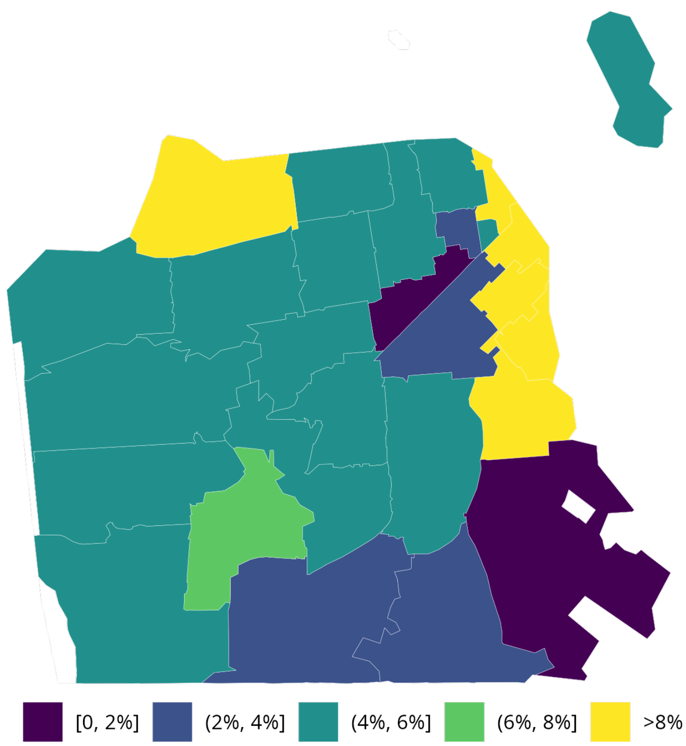

Figure 10: Exogenous Fundamentals of Calibrated Quantitative Spatial Model

(a) Productivity

(b) Amenity

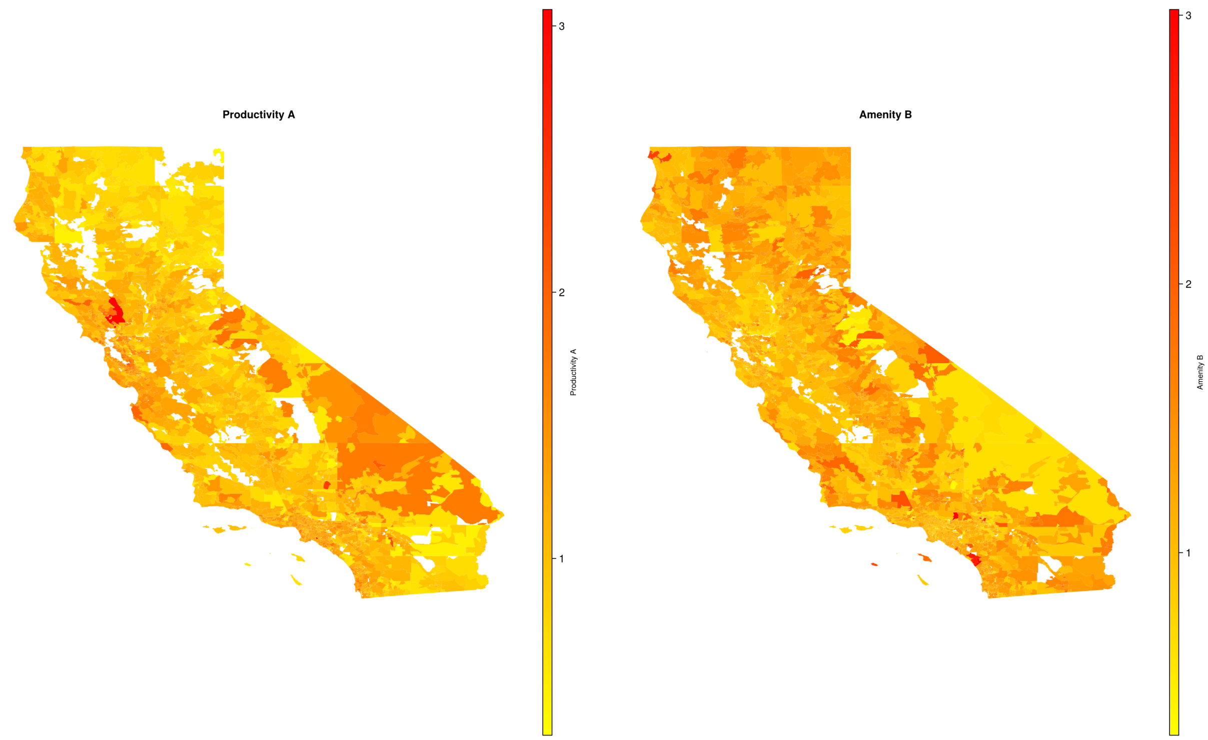

Figure 11: Spatial Sorting in Model

(a) High Skill Share of Workers at Baseline

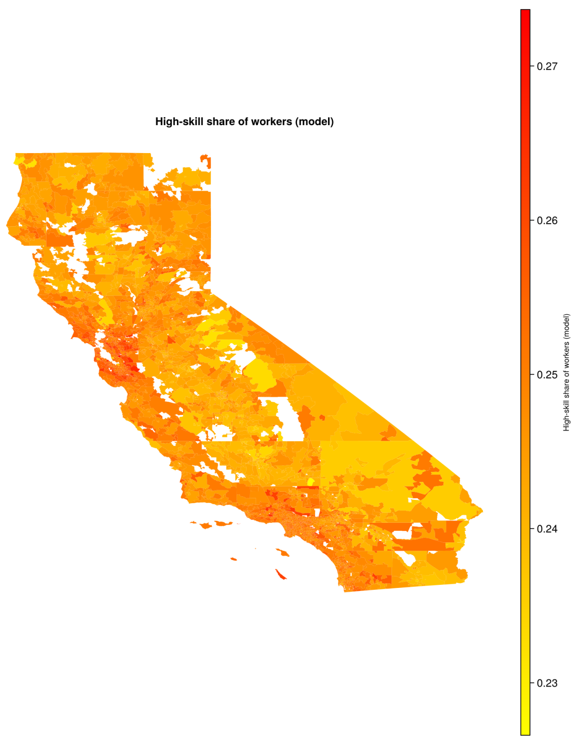

(b) High Skill share of Residents at Baseline

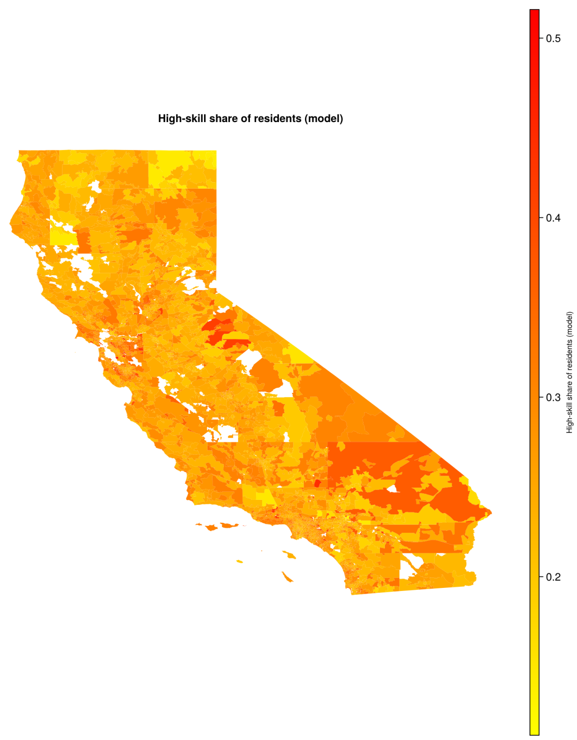

Figure 12: Current IZ Ordinances Counter-factual Effects

(a) Effect on IZ Ordinances on Market Rents

(b) Effect of IZ Ordinances on % Change on Low-Skill Residents

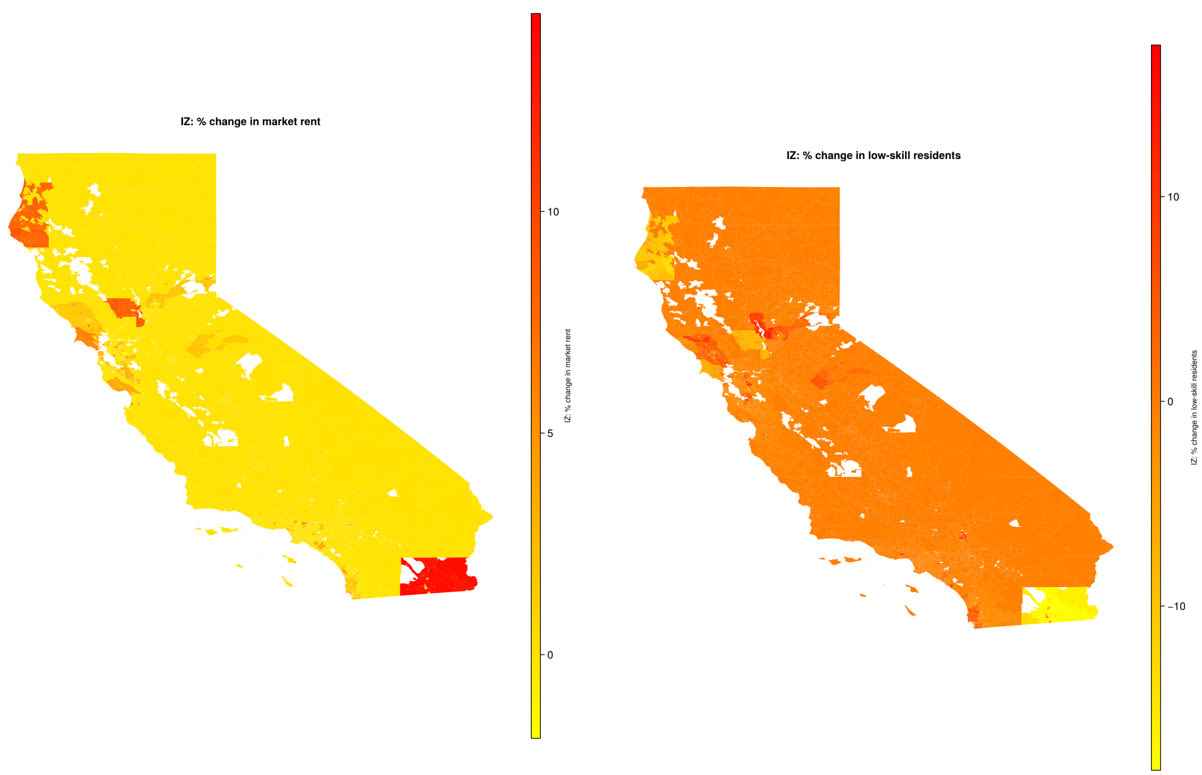

Table 1: Policy Parameters, Neighborhood Level

IZ Policy Parameter	Mean	St. Dev.	Max
All IZ Policy Options $\times$ Neighborhoods (N = 1,216)			
Extremely Low Income Share ($\eta_{.3}$)	0.510	2.148	10
Very Low Income Share ($\eta_{.5}$)	2.475	4.956	25
Low Income Share ($\eta_{.8}$)	8.726	7.123	30
Moderate Income Share ($\eta_{1.25}$)	3.004	6.227	25
IZ Stringency (τ_p)	0.054	0.037	0.240
Stringency Minimizing Policy Option in Each Neighborhood (N = 614)			
Extremely Low Income Share ($\eta_{.3}$)	0.244	1.480	10
Very Low Income Share ($\eta_{.5}$)	2.095	4.829	25
Low Income Share ($\eta_{.8}$)	7.445	5.973	25
Moderate Income Share ($\eta_{1.25}$)	4.987	7.541	25
IZ Stringency (τ_p)	0.042	0.030	0.135

Summary Statistics of IZ policy parameters at the neighborhood level are unweighted and conditional on a neighborhood falling in a place with an Inclusionary Zoning ordinance.

Table 2: Place Attributes by IZ Policy Stringency

	No Policy	0-2.5%	2.5-5%	5-7.5%	over 7.5%
ACS 2013-2017					
Avg. H.H. Income (thou. $)	73.23	82.82	88.88	94.73	107.49
	(38.21)	(28.86)	(31.16)	(29.70)	(27.46)
Avg. Rent ($)	1,366	1,484	1,683	1,748	1,879
	(519)	(399)	(391)	(359)	(306)
Avg. O.O. Unit Value (thou. $)	473	591	660	712	935
	(400)	(332)	(337)	(358)	(394)
Share of Units Multi-Family (%)	22.44	28.52	29.26	34.38	30.10
	(13.21)	(17.15)	(19.81)	(16.04)	(14.95)
Share of Units Renter Occupied (%)	40.98	42.92	41.53	44.82	40.36
	(14.70)	(14.62)	(17.89)	(10.82)	(11.55)
Avg. Year Home Built	1974	1969	1974	1972	1972
	(11)	(11)	(11)	(10)	(10)
4-year Degree Attainment (%)	27.79	39.18	38.71	44.71	57.79
	(19.24)	(20.23)	(18.39)	(17.77)	(12.48)
Share non-Hisp. White (%)	43.23	49.43	50.97	46.02	68.99
	(25.40)	(26.15)	(24.35)	(20.74)	(11.81)
Share Hispanic (%)	39.05	32.13	32.02	25.17	14.19
	(26.38)	(26.81)	(22.50)	(11.95)	(6.21)
Share non-hisp. Asian (%)	10.10	11.73	10.98	18.81	10.71
	(12.85)	(15.05)	(13.49)	(15.09)	(8.22)
Share non-Hisp. Black (%)	3.70	2.61	2.35	4.50	1.80
	(5.04)	(2.21)	(3.17)	(5.30)	(1.34)
Share of HH. Moved in Past Year (%)	12.74	11.79	12.05	13.43	15.26
	(4.80)	(2.16)	(2.99)	(5.40)	(5.75)
Population (thou.)	69.23	46.53	174.13	102.77	41.48
	(211.30)	(41.73)	(313.13)	(199.66)	(31.94)
Building Permit Survey 2013-2017					
Units Permitted/Year	201	122	595	333	76
	(745)	(153)	(1,147)	(601)	(89)
Assessed Value/Unit (thou. $)	311	309	381	335	421
	(298)	(151)	(477)	(239)	(240)
Multi-Family Units Permitted/Year	77	60	458	216	41
	(571)	(79)	(1,048)	(416)	(56)
Building Permit Survey 2018-2024					
Units Permitted/Year	158	105	582	436	111
	(672)	(160)	(1,359)	(1,000)	(160)
Assessed Value/Unit (thou. $)	311	275	377	308	398
	(331)	(121)	(374)	(229)	(185)
Multi-Family Units Permitted/Year	69	60	469	314	56
	(564)	(114)	(1,201)	(748)	(112)
WRLURI-18	0.61	0.77	1.30	1.14	1.63
	(0.92)	(0.59)	(1.06)	(0.62)	(0.38)
2016 Democratic Vote Share	0.52	0.61	0.62	0.71	0.70
	(0.14)	(0.15)	(0.10)	(0.09)	(0.10)
Housing Unit Supply Elasticity	0.22	0.14	0.15	0.08	0.10
	(0.23)	(0.21)	(0.20)	(0.11)	(0.21)
Place Count	416	26	24	24	18

ACS summary statistics aggregated from block groups for remainder of counties. Wharton Residential Land Use Restriction Index (WRLURI-18) missing for most places. Housing Unit Supply Elasticity aggregated from tract level estimates in Baum-Snow and Han (2024)

Table 3: Headline Results: IZ and Units Permitted

Dependent Variable:	log(Permitted Units)			
Model:	(1)	(2)	(3)	(4)
$Post_t \times \tau_p$	-7.576***	-9.108**	-9.103**	-7.409***
	(2.112)	(4.390)	(4.435)	(2.496)
$Post_t \times \mathbb{1}(IZ\ Policy_i)$		0.0877		
		(0.2380)		
Place F.E.	Yes	Yes	Yes	Yes
Year F.E.	Yes	Yes	Yes	Yes
Sample Restriction			Treated Only	Considering
Places	484	484	88	220
Observations	4,919	4,919	916	2,230

"Considering" sample restricts to places which reported they were considering adoption of IZ in their 4th cycle housing element.
Weighted by pre-treatment (2013-2017) building permit counts
Place clustered standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 4: Robustness to Discrete Treatment

Dependent Variable:	log(Permitted Units)			
Model:	(1)	(2)	(3)	(4)
$Post_t \times \mathbb{1}(IZ\ Policy_i)$	-0.3278**			
	(0.1268)			
$Post_t \times \mathbb{1}(\tau_p < 2.5\%)$		-0.1578		
		(0.2485)		
$Post_t \times \mathbb{1}(2.5\% \leq \tau_p < 5\%)$		-0.1554		
		(0.1988)		
$Post_t \times \mathbb{1}(5\% \leq \tau_p < 7.5\%)$		-0.5534***		
		(0.1294)		
$Post_t \times \mathbb{1}(7.5\% \leq \tau_p)$		-0.6255***		
		(0.2297)		
$Post_t \times \mathbb{1}(\tau_p < \bar{\tau})$			-0.0233	
			(0.1705)	
$Post_t \times \mathbb{1}(\bar{\tau} \leq \tau_p)$			-0.5861***	-0.5627***
			(0.1068)	(0.1864)
Place F.E.	Yes	Yes	Yes	Yes
Year F.E.	Yes	Yes	Yes	Yes
Sample Restriction				Treated Only
Places	484	484	484	88
Observations	4,919	4,919	4,919	916

Weighted by pre-treatment (2013-2017) building permit counts
Place clustered standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 5: Treatment Models for Inverse Propensity

	Dependent variable:		
	IZ Stringency	IZ Stringency >0	IZ Stringency Bin
	OLS	*Logistic*	*Ordered Logistic*
	(1)	(2)	(3)
log(pop)	0.001	0.237^{*}	0.265^{**}
	(0.001)	(0.131)	(0.124)
Share non- Hisp. White	0.0002^{**}	0.021	0.019^{*}
	(0.0001)	(0.013)	(0.011)
Share Hispanic	0.00005	0.016	0.012
	(0.0001)	(0.015)	(0.015)
Share with <30 min. Commute	0.0002	0.033	0.042^{*}
	(0.0001)	(0.023)	(0.022)
Share with $\geq$ 30 min. Commute	0.00005	0.016	0.021
	(0.0001)	(0.019)	(0.018)
4-year Degree Attainment	0.0002	−0.003	−0.00005
	(0.0002)	(0.024)	(0.020)
log(Avg. H.H. Income)	0.004	2.141	1.712^{*}
	(0.008)	(1.401)	(0.970)
Share of Units Renter Occupied	0.0003^{*}	0.023	0.036
	(0.0002)	(0.027)	(0.025)
Share of Units Multi-Family	0.0001	0.026	0.007
	(0.0001)	(0.022)	(0.020)
Share of Units Build 2000-2010s	0.0001	0.003	0.005
	(0.0001)	(0.021)	(0.020)
Share of Units Build 1980s-1990s	0.0003^{***}	0.036^{***}	0.045^{***}
	(0.0001)	(0.014)	(0.014)
Share of Units Build 1960s-1970s	0.0001	0.019	0.025^{*}
	(0.0001)	(0.015)	(0.014)
log(Avg. Rent)	0.005	0.282	0.091
	(0.008)	(1.305)	(1.267)
log(O.O. Unit Value)	−0.002	−0.457	−0.084
	(0.005)	(0.868)	(0.817)
2016 Democratic Vote Share	0.057^{***}	8.212^{***}	8.426^{***}
	(0.011)	(1.663)	(1.547)
Housing Unit Supply Elasticity	0.006	0.194	−0.194
	(0.006)	(1.014)	(1.004)
Constant	-0.132^{*}	-35.286^{***}	
	(0.074)	(11.260)	
Observations	444	444	444

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 6: Robustness to Inverse Propensity Weighting

Dependent Variable:	log(Permitted Units)			
Model:	(1)	(2)	(3)	(4)
$Post_t \times \mathbb{1}(\tau_p < 2.5\%)$	-0.0695			
	(0.2521)			
$Post_t \times \mathbb{1}(2.5\% \leq \tau_p < 5\%)$		-0.0376		
		(0.2030)		
$Post_t \times \mathbb{1}(5\% \leq \tau_p < 7.5\%)$			-0.4040***	
			(0.1328)	
$Post_t \times \mathbb{1}(7.5\% \leq \tau_p)$				-0.4560**
				(0.2302)
Places	375	370	373	366
Observations	3,864	3,820	3,841	3,779

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Locality and Year fixed effects included in all models.
Weighted by pre-treatment (2013-2017) building permit counts.
Controls reweighted estimated propensity to adopt IZ at treated bin.
Propensity to adopt IZ estimated by locality characteristics in Table 5.
Sample restricted to localities treated at relevant bin and control localities.
Locality clustered standard-errors in parentheses.

Table 7: Robustness to Alternative Weighting and Estimation Strategies

Dependent Vars:	log(Permitted Units)		Permitted Units		log(Permitted Units)	
Model:	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	OLS	Poisson	Neg. Bin.	OLS	Callaway
$Post_t \times \tau_p$	-5.212**	-1.861	-6.690***	-7.358***	-7.840***	
	(2.226)	(1.488)	(1.838)	(1.854)	(2.127)	
ATT						-0.159
						(0.0129)
ACRT						-6.101***
						(0.591)
Over-dispersion				0.311***		
				(0.0394)		
Place F.E.	Yes	Yes	Yes	Yes	Yes	
Year F.E.	Yes	Yes	Yes	Yes	Yes	
Sample Restriction					Balanced	Balanced
Weighting	Population	None	Pre-treatment (2013-2017) building permit counts			
Places	484	484	487	487	358	358
Observations	4,919	4,919	5,344	5,344	3,938	3,938

Place clustered standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 8: Compositional Effects

Dependent Vars:	$\log\left(\frac{\text{Assessed Value}}{\text{Permitted Unit}}\right)$	$\frac{\text{Multi-Family Units}}{\text{Permitted Units}}$	$\frac{\text{5+ Unit Multi-Family}}{\text{Permitted Units}}$	Costar Units Built	
Model:	(1)	(2)	(3)	(4)	(5)
	OLS	OLS	OLS	Poisson	Poisson
$Post_t \times \tau_p$	0.1203	-0.0945	-0.0283		
	(0.7118)	(0.5889)	(0.5676)		
$Post_t \times \tau_n$				-9.245***	-24.80**
				(3.057)	(9.741)
Place F.E.	Yes	Yes	Yes		
Year F.E.	Yes	Yes	Yes	Yes	Yes
Neighborhood F.E.				Yes	Yes
Place × Year F.E.					Yes
Clustering	Place	Place	Place	Neighborhood	Neighborhood
Weighting	Pre-treatment (2013-2017) building permit counts			None	None
Neighborhoods				880	650
Places	484	484	494	370	142
Observations	4,919	4,919	4,919	9,680	5,218

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 9: Evaluating Potential Spillover Effects of IZ

Dependent Variable:	log(Permitted Units)			
Model:	(1)	(2)	(3)	(4)
$Post_t \times \sum_j w_{ij}^{adj} \tau_j$	-4.634		-3.992	
	(5.838)		(5.117)	
$Post_t \times \sum_j w_{ij}^{com} \tau_j$		-4.621		-3.986
		(7.450)		(7.177)
$Post_t \times \tau_p$			-8.580***	-5.967
			(3.305)	(5.968)
$Post_t \times \sum_j w_{ij}^{adj} \tau_j \times \tau_p$			106.4	
			(140.1)	
$Post_t \times \sum_j w_{ij}^{com} \tau_j \times \tau_p$				-94.65
				(380.9)
Place F.E.	Yes	Yes	Yes	Yes
Year F.E.	Yes	Yes	Yes	Yes
Sample Restriction	Untreated Only			
Places	398	401	479	484
Observations	4,020	4,053	4,873	4,919

Samples for models using adjacency weights are smaller due to 5 localities with no adjacent non-missing localities. In all cases, these are communities with unincorporated LA county (currently excluded) as their only neighbor.
Weighted by pre-treatment (2013-2017) building permit counts
Place clustered standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 10: Model Parameters

Parameter	Symbol	Value	Source
Production (Cobb-Douglas, CRS)			
Low-skill labor share	a	0.541	ACS PUMS 2017 earnings by skill
High-skill labor share	$b = 1 - a$	0.459	
Preferences			
Low-skill housing share	α^L	0.28	ACS 2017 gross rent as a fraction of household income, by skill
High-skill housing share	α^H	0.20	
Low-skill Fréchet shape	ε^L	7.94	Estimated by minimizing the squared gap between model-implied and observed variance of PWPUMA-level mean earnings by skill
High-skill Fréchet shape	ε^H	7.15	
Commuting			
Low-skill commuting decay	$\kappa^L\varepsilon^L$	0.07302	Gravity regression on LODES 2017 inferred commuting flows by skill and r5r median travel times
High-skill commuting decay	$\kappa^H\varepsilon^H$	0.07301	
Housing supply			
Supply elasticity	η_j	Varies by j	Baum-Snow and Han (2024)
Inclusionary Zoning Policy			
Set-aside rate	z_j	Varies by j	California Housing Elements

Notes. High skill workers are those with a bachelors degree who are at least 30 years old. All other workers are considered low skill. The Fréchet shape ε^s and the reduced-form commuting decay $\kappa^s\varepsilon^s$ are estimated separately for each skill group. The exogenous fundamentals of productivity A_j, amenity B_i, and the housing-supply shifter ψ_i are recovered by structural inversion to match observed workplace employment, residence employment, and market rents. IZ lottery offer rate λ is set endogenously so that the equilibrium share of winners choosing IZ exactly fills the available IZ units.

Table 11: Welfare Effects of IZ no-IZ Baseline

Group	Current IZ Ordinances	15% Statewide IZ
Low Skill Welfare	+0.15% (11.4 million)	-0.42% (11.4 million)
IZ Winners Welfare	+2.7% (1.1 million)	+1.7% (2.3 million)
IZ Non-winners Welfare	-0.16% (10.2 million)	-1.97% (9.1 million)
High Skill Welfare	-0.11% (3.7 million)	+0.37% (3.7 million)

Group populations in parenthesis.
Estimated in the quantitative spatial model described in section 5.