Inclusive but Concentrated: The Impacts of Voluntary Inclusionary Zoning Evidence from Los Angeles, California

Ignacio Aravena-Gonzalez*

July 23, 2026

Click here for the latest version

Abstract. Voluntary inclusionary zoning (VIZ) is widely used to address housing affordability, yet its market effects remain uncertain. This paper studies Los Angeles's Transit Oriented Communities (TOC) program, a tiered voluntary bonus that exchanges added capacity and streamlining for on-site affordable units. Exploiting sharp eligibility boundaries, I combine a boundary difference-in-differences (DiD) with a boundary kink (BK) design. Within ±500 m of the boundary, TOC increases the annual probability that a block group receives any affordable project by about 1.1–1.3 percentage points (roughly a 90% rise over the pre-policy mean of 1.3%). Project counts and affordable units show similar patterns. Effects reflect a positive post-policy interior gradient inside TOC rather than a discrete jump at the edge. Guided by a partial-equilibrium model, the heterogeneity analysis shows that gains concentrate in below-median-rent, higher-tier locations, with weak responses in affluent single-family areas. As mechanisms, nearby price effects are near zero in dense, lower-rent areas but small and localized (about 6–8% declines) around single-family sales in above-median neighborhoods; moreover, mixed-income developments in these areas exhibit negative rent premia. At wider radii (±4 km), propensity-score matching (PSM) shows increases in market-rate permits and apartment units; within the boundary window, there is no contemporaneous change in the housing stock, consistent with permit timing. Overall, TOC expands affordable supply where incentives bind without reducing nearby market-rate output, but benefits are spatially concentrated, implying limited citywide inclusion without complementary tools in high-rent areas.

Keywords: Affordable housing; Land use regulation; Housing policy.
JEL Codes: R31; R38; R52; H71; R21.

*London School of Economics and Political Science. i.a.aravena-gonzalez@lse.ac.uk

Acknowledgments: I gratefully acknowledge ANID (Agencia Nacional de Investigación y Desarrollo, Chile) — Scholarship ID 72210405. I am grateful to my supervisors Christian Hilber and Nancy Holman for their guidance. I also thank Felipe Carozzi, Olmo Silva, Gabriel Ahlfeldt, and seminar participants for comments.

1 Introduction

Housing affordability has worsened in many large metropolitan areas as demand has risen against long-run supply constraints, with land-use regulation playing a central role (Hilber & Schöni, 2022). The burden falls disproportionately on low-income renters, who face high rent-to-income ratios and elevated eviction risk (Desmond, 2018). A prominent local policy response is inclusionary zoning, which trades regulatory relief such as added development capacity, waivers, or faster approvals for on-site income-restricted units within market-rate projects.

Despite its prevalence, evidence on inclusionary zoning remains mixed, in part because many programs are voluntary and hinge on whether developers choose to participate. This paper focuses on the adoption margin. When developers can opt in, production should scale up only where the value of regulatory relief exceeds the revenue loss from reserving some units at regulated affordable rents. A key implication is that take-up should be strongest where the gap between market rents and regulated affordable rents is moderate and where baseline density makes additional capacity valuable.

I study Los Angeles's Transit Oriented Communities (TOC) program, adopted in late 2017. TOC offers density and floor-area increases, waivers, and streamlined approvals in eligible areas near transit, in exchange for on-site set-asides that vary by affordability tier. The program provides a sharp test of voluntary inclusionary incentives because eligibility is defined by clear geographic boundaries and incentive generosity varies across tiers.

I assemble a block-group panel from 2014 to 2022 that links geocoded deed restrictions for affordable projects and multifamily permits to TOC eligibility. My core design compares block groups just inside and just outside TOC boundaries within a narrow window and estimates a boundary difference-in-differences. I then examine how effects vary with local market conditions and incentive generosity, and I complement the boundary results with broader-radius analyses that rely on matching and distance patterns.

Three findings emerge. First, within 500 meters of the boundary, TOC increases the annual probability that a block group receives any affordable project by about 1.1 to 1.3 percentage points relative to a pre-policy mean near 1.3 percent, with similar increases in project counts and affordable units. Second, effects are highly uneven across space. Impacts are concentrated in moderate-rent, multifamily contexts and in areas with more generous incentives, while effects are weak in high-rent, single-family areas. Third, the evidence points to composition rather than scale. I find no systematic increase in overall new multifamily permitting or in the growth rate of total or multifamily stock over the study horizon. A placebo using the California Density Bonus program shows modest declines inside TOC areas after 2018, consistent with substitution toward TOC rather than a broader permitting surge.

These patterns align with a simple adoption calculus. Where the market-to-regulated rent gap is large and development costs are high, reserving units at regulated rents is more likely to deter participation. Where the gap is smaller and additional capacity and regulatory savings are valuable, voluntary incentives translate into delivered affordable units.

This paper contributes to three strands of work. First, it adds causal evidence on voluntary inclusionary incentives and affordable housing production. Existing evidence is often hard to compare across jurisdictions because programs differ in generosity, eligibility, and enforcement, and because participation is endogenous under opt-in designs (Schuetz & Meltzer, 2012; Wang & Fu, 2022). By leveraging sharp eligibility boundaries in Los Angeles and tracking geocoded deed restrictions and permits, I identify where voluntary incentives translate into delivered income-restricted projects and units.

Second, it clarifies how policy design interacts with local market conditions. A natural prediction under opt-in inclusionary tools is that take-up depends on the gap between market rents and regulated affordable rents, and on how valuable added capacity and regulatory relief are in a given neighborhood. While prior work emphasizes feasibility and provides simulations or descriptive patterns (Phillips, 2024; Zhu et al., 2021), causal evidence on this interaction is limited. I show that more generous incentives amplify take-up where the market-to-regulated rent gap is modest and baseline density makes additional capacity valuable, helping organize why similar tools perform unevenly across space.

Third, it speaks to broader work on land-use regulation, housing supply, and market outcomes in constrained metros (Hilber & Schöni, 2022; Saiz, 2010). I find that the program shifts the composition of development activity and can substitute away from alternative bonus channels, while I do not detect a systematic increase in the growth rate of total or multifamily stock over the study horizon. This evidence helps reconcile why inclusionary tools may raise affordable output locally without generating large average changes in market-wide permitting.

The remainder of the paper is organized as follows. Section 2 describes TOC and presents a simple framework for opt-in adoption. Section 3 introduces the data. Section 4 presents the empirical strategy and main boundary estimates. Section 5 examines supporting evidence on channels, including localized price responses. Section 6 extends the analysis beyond the boundary. Section 7 reports robustness checks, and Section 8 concludes.

2 Policy Background and Theoretical Predictions

2.1 Policy background on Los Angeles TOC

Los Angeles adopted the Transit Oriented Communities (TOC) Affordable Housing Incentive Program after voter approval of Measure JJJ in November 2016. The program took effect on September 22, 2017 and was revised on February 26, 2018. TOC is voluntary and bundled. Eligible projects receive additional development capacity and regulatory relief in exchange for on-site income-restricted units.

Eligibility is geographic. TOC "Incentive Areas" are defined within one-half mile (2,640 feet) of a qualifying Major Transit Stop, either a rail station or the intersection of qualifying high-frequency bus routes as specified in the Guidelines. Parcels within an Incentive Area are assigned to tiers from 1 to 4 based on the shortest distance from the parcel to the qualifying stop, where higher tiers correspond to closer proximity. Figure 1 shows the spatial distribution of eligible areas and tiers across the city.

Conditional on meeting the on-site set-aside requirement, projects receive base incentives that expand feasible capacity, including tier-specific increases in residential density and floor-area ratio. The Guidelines also allow additional incentives and procedural provisions that can reduce approval timelines. Throughout the paper, I treat TOC as a policy bundle that shifts feasibility through both capacity and regulatory channels.

Table 1 summarizes the minimum on-site set-aside shares and the main base incentives by tier and income category. The tiered design generates systematic variation in incentive generosity within the eligible geography, which I use to study heterogeneous take-up and impacts by local market conditions and treatment intensity.

Table 1: TOC set-aside requirements and base incentives by tier

	Minimum on-site set-aside share			Base incentives	
Tier	ELI	VLI	LI	Max density increase	Max FAR increase
1	0.08	0.11	0.20	0.50	0.40
2	0.09	0.12	0.21	0.60	0.45
3	0.10	0.14	0.23	0.70	0.50
4	0.11	0.15	0.25	0.80	0.55

Notes. Set-aside shares are defined as a fraction of total units and vary by tier and affordability target. ELI denotes Extremely Low Income, VLI denotes Very Low Income, and LI denotes Lower Income as referenced in the TOC Guidelines and California Health & Safety Code. Base incentives report the maximum increases in residential density and floor-area ratio allowed by tier under the Guidelines.

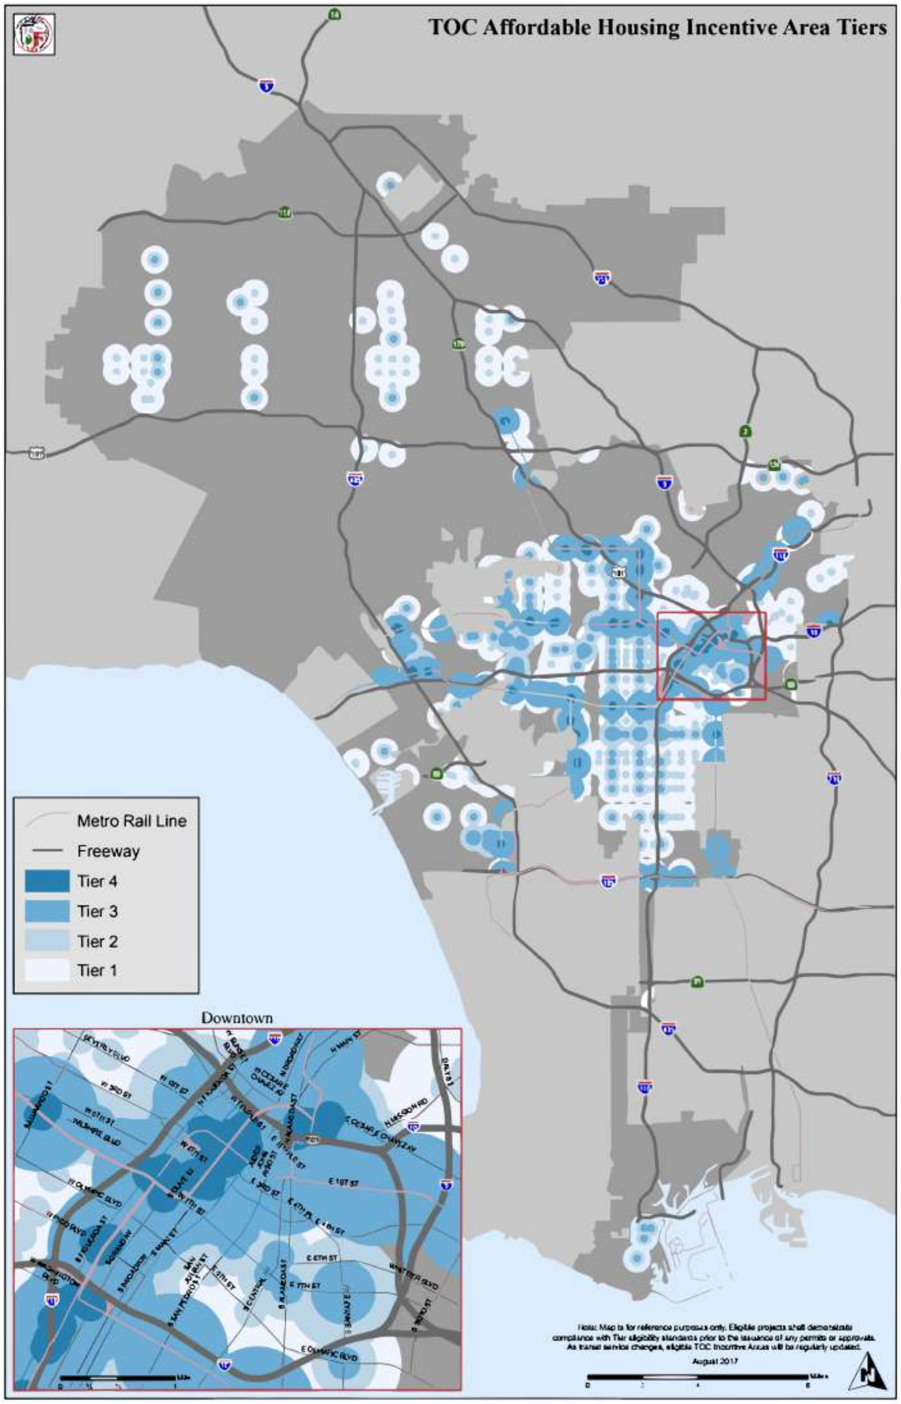

Figure 1: TOC Incentive Areas. Source is Los Angeles City Planning.

2.2 A simple opt-in framework and predictions

This subsection presents a simple opt-in framework for take-up under voluntary inclusionary incentives. The objective is to clarify the participation margin and motivate heterogeneity tests, not to model equilibrium sorting or citywide supply.

Consider a neighborhood i. Let p_i^M denote the local market rent (or price) per unit and p_{ig}^A the regulated affordable rent for target group g. Define the market-to-regulated rent gap

$$\Delta p_{i,g} \equiv p_i^M - p_{ig}^A \ \geq 0.$$

Let q_i^0 be feasible units under the base code and $q_{i\tau}^1 \geq q_i^0$ feasible units under TOC tier $\tau \in \{1, 2, 3, 4\}$. When the set-aside binds, the affordable share is $\alpha_{\tau,g} \in (0, 1)$, so affordable units are $q_a = \alpha_{\tau,g} q_{i\tau}^1$ and market-rate units are $q_m = (1 - \alpha_{\tau,g}) q_{i\tau}^1$.

Costs are $C_i(q; k_\tau)$, increasing and convex in q, and decreasing in concessions k_τ. Let $S_{i\tau}$ capture monetized regulatory savings under tier τ. Allow for local capitalization on market-rate revenue via λ_i, so market-rate revenue is multiplied by $(1 + \lambda_i)$.

Baseline profit is $\pi_i^0 = p_i^M q_i^0 - C_i(q_i^0; 0)$. Under TOC tier τ,

$$\pi_{i\tau}^1 = (1 + \lambda_i) p_i^M (1 - \alpha_{\tau,g}) q_{i\tau}^1 + p_{ig}^A \alpha_{\tau,g} q_{i\tau}^1 - C_i(q_{i\tau}^1; k_\tau) + S_{i\tau}.$$

The incremental profitability of opting in is

$$\Delta\Pi_{i\tau g} \equiv \pi_{i\tau}^1 - \pi_i^0 = \underbrace{p_i^M (q_{i\tau}^1 - q_i^0) + S_{i\tau} - \Big[C_i(q_{i\tau}^1; k_\tau) - C_i(q_i^0; 0)\Big] + \lambda_i p_i^M (1 - \alpha_{\tau,g}) q_{i\tau}^1}_{\text{Bonus}_{i\tau g}} - \alpha_{\tau,g} \Delta p_{i,g}\, q_{i\tau}^1. \tag{2.1}$$

A developer opts in when

$$\Delta\Pi_{i\tau g} \geq 0 \quad \Longleftrightarrow \quad \text{Bonus}_{i\tau g} \geq \alpha_{\tau,g} \Delta p_{i,g}\, q_{i\tau}^1. \tag{2.2}$$

Prediction 1. Take-up is less likely when the market-to-regulated rent gap is larger and more likely when the policy package delivers larger effective benefits through added capacity, regulatory savings, and concessions. If capitalization is negative in higher-income settings, take-up is further discouraged there.

Prediction 2. When the quota binds and TOC increases total feasible units, affordable unit counts rise mechanically because the requirement is defined as a share of total units.

Market-rate unit counts rise as well provided that

$$(1-\alpha_{\tau,g})\, q_{i\tau}^{1} > q_i^{0} \quad \Longleftrightarrow \quad \frac{q_{i\tau}^{1}}{q_i^{0}} > \frac{1}{1-\alpha_{\tau,g}}. \tag{2.3}$$

This implication is mechanical at the project margin and speaks to local displacement within the narrow boundary window, while it does not rule out broader reallocation across the city.

Appendix 8 reports derivations and comparative statics.

3 Data

3.1 Data sources

I built a block–group per year panel by merging several geocoded datasets. The core data comes from the Los Angeles City Controller's registry of affordable housing covenants (2014–2022). The file includes project address, ZIP, council district, covenant year, affordable and total units, the certificate of occupancy (CofO) date and status, and program type (TOC, Density Bonus (DB), other). In the raw data there are 1,405 projects. In the pre-policy period (2010–2017), density bonus programs accounted for about 84% of affordable developments, with the remaining 16% from other ordinances; in the post-policy period (2018–2022), TOC projects accounted for roughly 51% of developments, compared with 29% from density bonus programs and 20% from other ordinances. I focus on privately developed projects and exclude purely publicly financed projects (e.g., land trusts or non–profit acquisitions) that can be delivered through programs unrelated to TOC; furthermore, I also exclude "other ordinances" as they are mostly parking reductions without adding new units.[1]

Developments are geocoded and spatially joined to 2010 Census block groups (BG) and to official TOC buffers/tier polygons from Los Angeles City Planning. I compute the signed distance d_{it} (meters) from each BG to the nearest TOC boundary and define side $S_i = 1\{d_i \geq 0\}$ so that positive distances are inside the policy boundary. I split polygons when the boundary intersects a BG polygon, creating two segments by treatment status; these segments receive distinct unit fixed effects in all panel specifications.[2]

Socioeconomic covariates are drawn from the 2014 ACS 5–year files at the BG level. I assemble data on population, race/ethnicity shares, female–headed households, low educational attainment, poverty, renter share, and median gross rent. These covariates are time–invariant

[1]These exclusions follow the classification in the Controller's registry and program documentation.
[2]As a robustness check, I drop bisected BGs entirely; results are unaffected.

in the main designs and enter only through balancing diagnostics, matching, or subgroup definitions.

I also add geocoded building permits (BP) from LADBS to construct a panel between 2014–2022. I drop revoked or expired permits and retain permits whose use types map into four categories: additions, alterations, demolitions, and new construction. I then classify projects into single-family, apartment, and commercial units. I also use Assessor Parcel Data (APD) spanning 2010-2024 to track the stock of units and developments over time. I use BPs to capture developers' responses as their dates are close to the covenant, while APDs allow me to test for changes in the overall supply.

Lastly, I use data on prices, rents, and costs to study potential mechanisms for λ_i and $C_i(q^1_{i\tau}; k_\tau)$. I use (i) the Secured Basic File (Assessor) with parcel–level attributes such as construction class/quality, unit construction cost, bedrooms, bathrooms, units, size, and the last sale on each parcel, and (2) a CoStar project–level cross-section on build–to–rent developments with information on unit counts, floors, average rent/size, address, submarket, and tenure/affordability subgroups.

3.2 Summary statistics and stylized facts

According to Los Angeles City Planning's dashboard, since the program launch, there have been about 28,954 approved TOC units, of which roughly 22% are income-restricted; the affordable share is consistently split between Extremely Low and Low Income targets (roughly 40–45% each). Annual approvals stabilized near 6,000–7,000 units through 2022 (with a smaller first year in 2018).

Panel A in Table 2 shows that demographic and socioeconomic covariates show similar trends pre treatment in the±500,m window[3]. These variables are relevant as many of them can help to predict the location of affordable units within the sample[4]. Panel B also shows similar levels of affordable units per block group and the percentage of affordable units within the Comptroller's data. By contrast, Panels C and D document systematically higher permitting activity and multifamily stock on the treated side, consistent with TOC-eligible areas being denser and more development-intensive at baseline. As expected, citywide comparisons shows much larger and more systematic differences between treated and control locations; these results are reported in Table A.2.

[3]Densities between groups are also similar for the variables (Figure A.2)

[4]To illustrate, Table A.1 shows that, citywide, affordable projects are more likely in areas with larger populations and higher renter shares, and slightly more common in places with greater poverty status, to name a few.

About the distribution of affordable units, there is a shift from DB to TOC. Before 2018, 447 developments (20,086 units) used DB citywide; after implementation, 446 developments (21,088 units) used TOC, and 250 (13,561 units) used DB. Figure 2a suggests a similar situation, as kernel density plots show movements of affordable projects toward the interior of the policy area. Moreover, since high-intensity tiers tend to cluster farther from the boundary than lower-tier treated areas, within ±500,m the difference between pre- and post-period bin means show fitted lines in Figure 2b, which suggest that the change in slopes is statistically significant and larger farther from the boundary in the treated area.[5]

[5]This relationship is descriptive; Section 4 formalizes it in the empirical strategy.

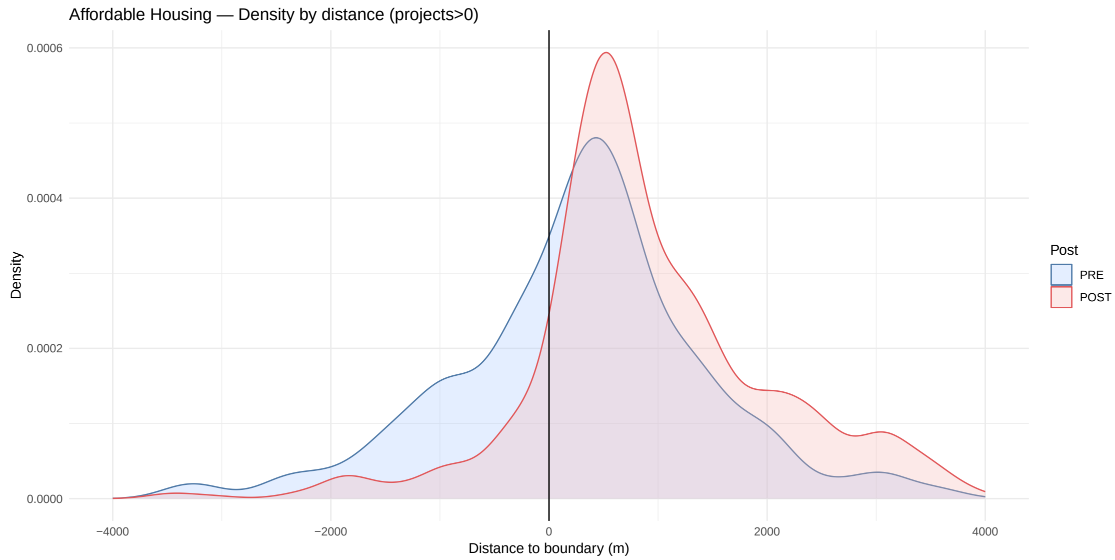

(a) Affordable housing—density of project by distance to boundary (pre vs. post)

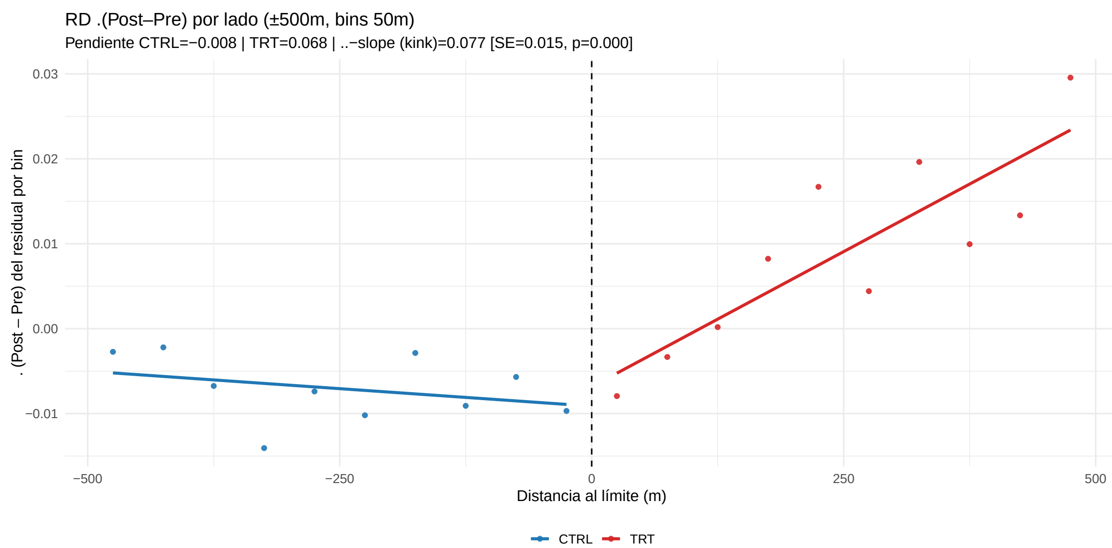

(b) Pre vs. post fitted boundary slopes (bin means; $\pm 500\,\mathrm{m}$)

Figure 2: Stylized facts around TOC eligibility.

Table 2: Pre-policy balance within ±500 m of the boundary (block groups)

	CTRL mean	TRT mean	Diff (TRT−CTRL)	p-value
Panel A. Sociodemographics (2014)				
Log(Pop)	7.220	7.220	−0.004	0.865
Pct. Asian	0.092	0.099	0.007	0.185
Pct. Black	0.145	0.140	−0.005	0.596
Pct. White	0.531	0.536	0.006	0.589
Pct. Hispanic	0.461	0.439	−0.022	0.110
Low education	0.634	0.614	−0.019*	0.069
Female HH	0.176	0.169	−0.007	0.146
Pct. renters	0.574	0.569	−0.005	0.678
Pct. poverty	0.201	0.198	−0.003	0.657
Median rent (USD)	1,137	1,150	12.8	0.601
Panel B. Affordable units per block group (mean 2010–2017)				
Aff. units per BG	0.055	0.052	−0.003	0.879
Pct. affordable	0.129	0.115	−0.014	0.608
Panel C. Building permits (annual mean 2010–2017)				
Additions (count)	0.448	0.506	0.058***	0.005
Alterations (count)	2.363	2.998	0.635***	0.000
Demolitions (count)	0.138	0.206	0.068***	0.000
New construction	0.222	0.280	0.058***	0.001
1–2 fam. units	0.075	0.089	0.014	0.388
Apartment units	0.296	0.496	0.200*	0.099
Panel D. Housing stock (pre-policy, 2016)				
Units, total	228.0	291.4	63.4***	0.000
Units, single-family	122.0	126.1	4.1	0.496
Units, multifamily	106.0	165.4	59.4***	0.000
Parcels, total	131.1	138.7	7.6	0.215
Parcels, single-family	114.3	113.5	−0.8	0.885
Parcels, multifamily	16.8	25.2	8.5***	0.000

Notes: Means are pre-policy. Panel A reports 2014 sociodemographic covariates; Panel B reports pre-policy (2010–2017) averages of affordable units; Panel C reports annual averages of building permits over 2010–2017; Panel D reports 2016 housing stock. The sample restricts to block groups within ±500 m of the TOC boundary. Stars reflect two-sided tests of equal means. Sociodemographic and affordable-housing covariates are well balanced (standardized mean differences below 0.1), whereas building-permit and housing-stock variables show larger differences, consistent with treated areas being denser and more development-intensive. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4 Empirical Strategy

This section estimates the local effects of TOC on affordable housing by exploiting the spatial discontinuity at the program boundary. I begin with a Diff-in-Diff specification that treats eligibility as binary within a ±500 m window. Then, I study heterogeneity across market conditions and policy tiers and, finally, complement the design with a boundary kink–in–DiD that allows for post–policy slope changes inside and outside the boundary as a function of distance to the boundary to account for any unobservables that might be endogenous to the location of affordable units and the policy.

Section 2.2 motivates these choices. Because eligibility is tied to transit proximity, development choices may be endogenous; therefore, I contrast treated and adjacent locations within a narrow window and rely on pre–treatment similarity across the boundary. Within ±500 m the panel is balanced across sides, supporting a conventional boundary DiD with unit and year fixed effects. However, Section 7 reports extensive diagnostics such as covariate balance, pre–trend tests, donut windows, bandwidth sensitivity, spatial SEs, and placebos.

Then, guided by the theoretical predictions in Section 2.2, I then examine heterogeneity across rent gaps and treatment intensity (tiers) and incorporate a BK–DiD to capture non-constant treatment intensity within the boundary, given the geographical nature of higher tiers concentrated in the inner rings of TOC.

4.1 The Impact of TOC on Affordable Housing

The baseline DiD for block group i in year t is:

$$Y_{it} \;=\; \beta\left(\mathrm{TOC}_i \times \mathrm{Post}_t\right) \;+\; \mu_i \;+\; \gamma_t \;+\; \varepsilon_{it}, \tag{4.1}$$

where Y_{it} is (i) the linear probability of any affordable project, (ii) the number of affordable projects, or (iii) the number of affordable units. TOC_i indicates being inside the policy boundary, and $\mathrm{Post}_t = 1$ for $t \geq 2018$. μ_i and γ_t are block–group and year fixed effects. Standard errors are clustered at the block–group level.

Table 3 reports local effects within 0–500 m. For "Probability," the estimate of 0.011 equals roughly one percentage point over a pre–policy mean of 0.012—about a 90% increase. For "Projects," 0.013 compares to a pre–TOC mean of 0.014 projects per block–group–year. For "Units," the estimate implies 0.15 more affordable units per block–group–year (pre–policy mean 0.053), roughly 2.7 times the baseline. These magnitudes are local to the boundary window and estimated with unit and year fixed effects and block–group–clustered SE.

Table 3: Impact of TOC on Affordable Housing within 0–500 m

	Probability	Projects	Units
TOC×Post ($t \geq 2018$)	0.011***	0.013***	0.146***
	(0.004)	(0.004)	(0.048)
Observations	17,091	17,091	17,091
R^2 (within)	0.197	0.206	0.139
Pre-treatment mean (2014–2017)	0.012	0.014	0.053

Notes: Block–group×year panel restricted to 0–500 m on either side of the TOC boundary. Unit and year fixed effects; standard errors clustered by block group. I implement different robustness checks in Section 7.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

To assess parallel trends, I estimate a dynamic specification with leads and lags of $\text{TOC}_i \times \text{Post}_t$, as shown in Figure A.3. The pre–treatment coefficients are small relative to the effects after 2018, and a series of pre–trend diagnostics—including joint tests of pre–treatment leads, linear pre–period slope differences, and a placebo DiD with a false policy date in 2016—does not reveal systematic pre–policy differences between sides. Post–2018 coefficients rise gradually, consistent with permitting and construction lags in Los Angeles.

Although zoning relaxation can mechanically increase covenants, the boundary design and dynamics suggest a local supply response rather than a displacement effect that merely reallocates projects instead of expanding the supply of affordable units. Section 7 formally tests this assumption, showing robust coefficients when excluding observations near the boundary, when changing the boundary, using placebos, and when accounting for potential spatial correlation (Conley standard errors). Taken together, these checks suggest the effects are not driven by mechanical reallocation.

4.2 Heterogeneity analysis: the location of affordable housing

A follow-up question concerns the heterogeneous effects of TOC. Motivated by Section 2.2, the impacts should be larger where the market–subsidy rent gap is small and developers can exploit the incentive most effectively. I therefore study heterogeneity along two dimensions. First, I split tracts into below- and above-median market rents, which serve as proxies for the rent gap. Second, I construct functional-area clusters using a machine-learning algorithm that groups blocks by baseline density, land-use intensity, and socioeconomic characteristics. In each case, I interact with them $\text{TOC}_i \times \text{Post}_t$, distinguishing higher tiers (3–4) from lower tiers (1–2). If the mechanism operates as in the model, responses should be strongest in dense, centrally located, below-median-rent areas and at higher TOC tiers.

4.2.1 Median rents

Let Below_i indicate below-median rent tracts and HighTier_i indicate eligibility for TOC Tiers 3–4. Define $\text{DiD}_{it} \equiv \text{TRT}_i \times \text{Post}_t$ with block-group and year fixed effects $\{\alpha_i, \gamma_t\}$ and standard errors clustered by block group. I estimate

$$\begin{aligned} Y_{it} = \alpha_i + \gamma_t + \beta_D \, \text{DiD}_{it} + \beta_{DB} \, (\text{DiD}_{it} \times \text{Below}_i) + \beta_{DH} \, (\text{DiD}_{it} \times \text{HighTier}_i) \\ + \beta_{DBH} \, (\text{DiD}_{it} \times \text{Below}_i \times \text{HighTier}_i) + \varepsilon_{it}, \end{aligned} \tag{4.2}$$

First, I only interact by rents, showing larger effects below the median (Table A.3); additional specifications that split the rent distribution into quartiles (Tables A.4) and k-means bins[6] (Table A.5) reveal a similar monotone pattern where effects decline as pre-policy rents rise, and with significant responses concentrated below the median (around \$1,100–1,200 in monthly rents). Then, the double interaction between Below_i and HighTier_i effects in the form of linear combinations [7]; are shown in Table 4.

Consistent with the model, effects are concentrated in areas with smaller rent wedges (below-median-rent tracts) and stronger incentives (high tiers). In Panel A (Probability), the effect is close to zero in above-median/low-tier cells (AL), becomes economically and statistically significant in below-median/low-tier cells (BL), and is largest in below-median/high-tier cells (BH). The count and units outcomes (Panels B–C) mirror this pattern. Furthermore, the Wald tests[8] suggest that more generous tiers amplify take-up precisely where market and subsidized rents are closer together. These gradients align with developers exploiting the incentive menu (density/FAR/height) in locations with looser feasibility constraints, where rents differences are lower.

[6] Details in Section 8

[7] Formally, $\tau_{AL} = \beta_D$, $\tau_{BL} = \beta_D + \beta_{DB}$, $\tau_{AH} = \beta_D + \beta_{DH}$, and $\tau_{BH} = \beta_D + \beta_{DB} + \beta_{DH} + \beta_{DBH}$.

[8] Wald summary: Probability—BH>AL ($p = 0.049$); Count—BH>AL ($p = 0.077$); Units—BL>AL ($p = 0.0056$), BH>AL marginal ($p = 0.067$). Within–stratum high-tier increments are not significant (BH–BL and AH–AL $p > 0.15$). Joint tests (DiD×Below×HighTier): global $DB{=}DH{=}DBH{=}0$ significant only for Units ($p = 0.013$); "Only Below" $DB{=}DBH{=}0$ significant for Units ($p = 0.0106$); "Only High tier" $DH{=}DBH{=}0$ not significant.

Table 4: Heterogeneity: Above/Below Median Rents × TOC Tiers (AL/BL/AH/BH; 0–500 m)

	AL	BL	AH	BH
Panel A. Probability of any AH				
ATT	0.005	0.014***	0.014	0.034**
	(0.005)	(0.005)	(0.022)	(0.014)
Panel B. Count of AH projects				
ATT	0.006	0.017***	0.026	0.032**
	(0.006)	(0.006)	(0.025)	(0.014)
Panel C. Number of AH units				
ATT	0.032	0.212***	0.098	0.520*
	(0.040)	(0.068)	(0.238)	(0.268)

Notes: AL = Above-median × Low tier (1–2); BL = Below-median × Low tier; AH = Above-median × High tier (3–4); BH = Below-median × High tier. Entries are ATTs (linear combinations of (4.2)); SEs in parentheses via the delta method with clustering by block group. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

4.2.2 Functional areas and land-use intensity

Next, I test whether the pattern sharpens along baseline conditions of the city. Motivated by the model, I allow the post-treatment effect of TOC to vary jointly across large-scale functional areas generated by a SKATER algorithm detailed in Appendix 8. These clusters differ systematically in zoning and built form like single- and multifamily shares, density, and in socioeconomic composition (racial mix, poverty, rents), as summarized in Table A.30. I use these clusters as a compact way to capture the spatial attributes of different city areas that might not be captured in my previous heterogeneity analysis.

Let $D_{it} \equiv \text{TOC}_i \times \text{Post}_t$ denote the baseline DiD indicator, and let $G_i \in \{1, 2, 3\}$ index the spatial clusters (with G_1 as the omitted group).

$$Y_{it} \;=\; \alpha_i + \gamma_t \;+\; \beta_D \, D_{it} \;+\; \beta_{DG2} \Big(D_{it} \times \mathbf{1}\{G_i{=}2\}\Big) \;+\; \beta_{DG3} \Big(D_{it} \times \mathbf{1}\{G_i{=}3\}\Big) \;+\; \varepsilon_{it}, \quad (4.3)$$

I estimate the equation with block–group and year fixed effects (α_i, γ_t) and standard errors clustered at the block–group level. Group-specific post-treatment effects are linear combinations of Eq.4.3[9]

Table 5 reports group-level ATTs within each functional area, showing uneven results across them. In Group 1 (affluent, predominantly single-family), the ATTs are small and non-significant across all outcomes. By contrast, Groups 2 and 3 exhibit positive impacts. In probability terms (Panel A), the likelihood of any affordable project rises by about 2.7 pp in Group 2 and 1.6 pp in Group 3, compared to roughly 0.3 pp in Group 1. Project counts (Panel B) follow the same pattern: treatment adds around 0.03 projects per block–group–year in Group 2 and 0.02 in Group 3, versus essentially zero in Group 1. For units (Panel C), the estimated effects reach roughly 0.31 and 0.30 affordable units per block–group–year in Groups 2 and 3, with much smaller and imprecise effects in Group 1.

Table 5: Heterogeneity: functional areas (G1–G3)

	Group 1	Group 2	Group 3
Panel A. Probability of any AH			
ATT	0.003	0.027***	0.016**
	(0.005)	(0.006)	(0.007)
Panel B. Count of AH projects			
ATT	0.005	0.029***	0.018**
	(0.006)	(0.006)	(0.009)
Panel C. Number of AH units			
ATT	0.043	0.314***	0.304**
	(0.041)	(0.106)	(0.152)

Notes: Block-group×year panel within 0–500 m of the TOC boundary; block-group and year fixed effects; standard errors clustered by block group. Entries are group-specific ATTs for each functional area, obtained as linear combinations of the coefficients in (4.3). Group 1 (G1) corresponds to affluent, predominantly single-family areas; Group 2 (G2) to central, dense/multifamily areas; and Group 3 (G3) to more mixed, lower-rent areas. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[9] $\tau_{G_1} = \beta_D, \tau_{G_2} = \beta_D + \beta_{DG2}, \tau_{G_3} = \beta_D + \beta_{DG3}$, and a Wald test of no geographic heterogeneity is $H_0 : \beta_{DG2} = \beta_{DG3} = 0$

These results align with expectations. Group 1 combines higher single-family shares, higher median rents, and a larger share of White and highly educated residents. Group 2 is denser and more multifamily, with higher renter shares and poverty, and the lowest median rents, while Group 3 is more mixed, with a high share of Black residents, intermediate density, and moderate rents. The concentration of TOC effects in Groups 2 and 3 is therefore consistent with the model as mixed-income feasibility improves where the market–subsidy rent wedge is smaller and density/FAR bonuses are more valuable, whereas affluent single-family areas remain largely unresponsive.

Additional interactions with high-intensity eligibility, reported in Appendix Table A.6, do not change this pattern. Wald tests confirm strong geographic heterogeneity across functional areas. At the same time, the incremental effect of high-intensity tiers is small and not statistically distinguishable from zero after accounting for group differences. In other words, the dominant margin of heterogeneity is spatial (Groups 2/3 vs. Group 1), rather than a uniform boost from high-intensity tiers. Section 5.2 explores mechanisms—prices, rents, and costs—to clarify why effects are muted in Group 1 and stronger in central districts.

4.3 Place-invariant gradients: Kink–DiD at the policy boundary

The baseline DiD summarizes the average post-TOC effect inside the 0–500 m window. To unpack where that effect concentrates spatially, I estimate a boundary "kink–DiD" that (i) allows for a discrete jump at the edge and (ii) recovers the post-policy slope with respect to signed distance on each side of the boundary. The design asks whether the policy's effects grow inside TOC areas and whether any part of the effect is an immediate discontinuity at the edge. This helps separate a boundary-level discontinuity from interior spatial gradients driven by different tiers and their endogenous location.

Let d_i be the signed distance to the boundary in kilometers (positive inside). Define the nonnegative components $D_{R,i} \in \{d_i, 0\}$ and $D_{L,i} \in \{-d_i, 0\}$, and $S_i = 1\{d_i \geq 0\}$. Within a symmetric window of ± 0.5 km, I estimate the linear specification

$$y_{it} = \alpha_i + \gamma_t \ + \ \tau\,(\text{Post}_t \times S_i) \ + \ \beta_R\,(\text{Post}_t \times D_{R,i}) \ + \ \beta_L\,(\text{Post}_t \times D_{L,i}) \ + \ \varepsilon_{it}, \qquad (4.4)$$

with block-group and year fixed effects (α_i, γ_t) and standard errors clustered by block group. The post-treatment jump at the boundary is τ. The post-treatment kink (change in slope at $d = 0$) is $\kappa \equiv \beta_R - \beta_L$. Equation (4.4) nests the baseline DiD as a special case; by averaging over a symmetric $[-b, b]$ window and differencing treated ($d_i \geq 0$) from control ($d_i < 0$) sides, one obtains the approximation $\beta^{DiD} \ \approx \ \tau \ + \ \frac{b}{2}\,\kappa$, because $\overline{D_R} = \overline{D_L} = b/2$ inside a symmetric window. Thus, the DiD can be interpreted as the sum of a boundary-level jump plus the average interior gradient implied by the kink.

Table 6 reports the jump and kink estimates for the three main outcomes within ±0.5 km. Jumps at the boundary are small and statistically insignificant. By contrast, the post-treatment kink is positive: slopes inside TOC become steeper than on the control side. The kink is statistically significant for the probability of any affordable project, counts, and affordable units. Pre-period diagnostics do not detect pre-policy breaks or differential pre-trends.

Table 6: Kink–DiD at the boundary (linear, ±0.5 km)

	Jump ($\hat{\tau}$)	Kink ($\hat{\kappa}$)
Probability of any AH	−0.004	0.069***
	(0.004)	(0.026)
Count of AH projects	−0.003	0.073**
	(0.005)	(0.030)
Number of AH units	0.000	0.634**
	(0.053)	(0.269)

Notes: Linear specification in (4.4) estimated within ±0.5 km; standard errors clustered by block group. A 100 m donut (excluding the ±0.1 km band around the boundary) delivers similar signs with somewhat wider confidence intervals (results available upon request). Pre-level RD and panel pre-trend tests (2014–2017) do not reject, and a 2017 placebo shows no jump and a unit kink consistent with end-2017 implementation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

I also estimate DiD specifications that interact treatment with high-intensity eligibility hi_any_i to obtain group-specific average effects. For the probability of any affordable project, the effect is significant only in BH (below median and high-tier block groups), consistent with more substantial take-up in areas where the TOC menu is more generous (see Appendix Table A.8), as in the previous subsections.

Pre-treatment diagnostics in Table A.7 support the validity of the kink–DiD design. A pre-policy RD in levels at the boundary yields small, statistically insignificant jumps across all outcomes. The coefficient on the interaction between the pre-period time trend and the treated side is also close to zero and imprecisely estimated, indicating no evidence of differential linear pre-trends between treated and control sides. Finally, the interactions between the pre-period time trend and the distance components on each side are individually insignificant, suggesting that the spatial gradient does not evolve differently across the boundary over time. Taken together, these results do not indicate any sharp pre-policy break or systematic differential pre-trends at the boundary, consistent with the identifying assumptions of the kink–DiD specification.

The kink design shows that TOC's effects scale with distance inside treated areas, as the average DiD effect is the sum of a tiny boundary jump and a positive interior gradient. This spatial pattern is strongest where the program's incentives are more valuable (high tiers) and feasibility constraints are looser (below-median rents), consistent with the model's rent-wedge and density channels.

5 Mechanisms

The heterogeneity results above show that TOC's supply response concentrates in block groups with lower market rents and greater zoning premiums, while effects are muted in affluent, single-family areas. This pattern is consistent with the theoretical propositions. In this section, I examine two mechanisms. First, I test whether nearby home prices capitalize mixed-income projects in ways that deter take-up in high-rent, single-family neighborhoods. Second, I assess construction costs across Los Angeles and rents in mixed-income projects to determine whether local rent gaps depress net operating income in affluent areas.

5.1 Mechanism I: "Not in my backyard" (local price capitalization)

To understand why mixed-income projects concentrate in the below-median and high Tier locations, I estimate a near–far DiD around each new affordable covenant:

$$\log P_{it} \;=\; \alpha_i \;+\; \gamma_t \;+\; \beta\,(\text{Near}_i \times \text{Post}_t) \;+\; X_{it}'\delta \;+\; \varepsilon_{it}, \tag{5.1}$$

where Near_i=1 for sales within 0–500 m and 0 for 500–1,000 m. The vector X_{it} includes standard hedonic controls, such as lot/building square footage, number of bedrooms, number of bathrooms, construction type, and quality. I include block-group and year fixed effects (α_i, γ_t) and two-way clustered standard errors by block group and year. I then allow the near–post effect to vary with neighborhood income (Below vs. Above) and structure type (single-family vs. multifamily):

$$\log P_{it} = \alpha_i + \gamma_t + \sum_{g \in \{\text{Below}, \text{SF}\}} \beta_g \,(\text{Near} \times \text{Post} \times g) + \beta_0 \,(\text{Near} \times \text{Post}) + X_{it}'\delta + \varepsilon_{it}.$$

Overall, prices within 500 m of a covenant fall by about 6.1%, with sharp heterogeneity by income and structure. Table 7 shows that the effect is concentrated in Above×SF (–8.1%), whereas other groups are small and/or imprecise. Event-study paths indicate slow-moving declines that materialize 2–3 years after the covenant—consistent with permitting-to-completion lags (Figure 3).

Table 7: Nearby price effects (Near×Post) by structure and rent bin

	Multifamily (MF)		**Single-family (SF)**	
	Below	**Above**	**Below**	**Above**
Near×Post	−0.031	0.041	−0.016	-0.081^{***}
	(0.049)	(0.056)	(0.016)	(0.025)

Notes: Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

I also study dynamics using a distance-decay design that compares concentric 200 m rings to a common 800–1,000 m baseline. Ring estimates show steep, short-ranged gradients in above-median areas (about –14% at 0–200 m and –8% at 200–400 m), with near-null effects in below-median tracts (Figure 3). Interacting by program type (DB vs. TOC) and intensity (tiers 3–4 vs. 1–2) yields similarly negative Near×Post estimates, with differences statistically indistinguishable from zero (Appendix Table A.9).

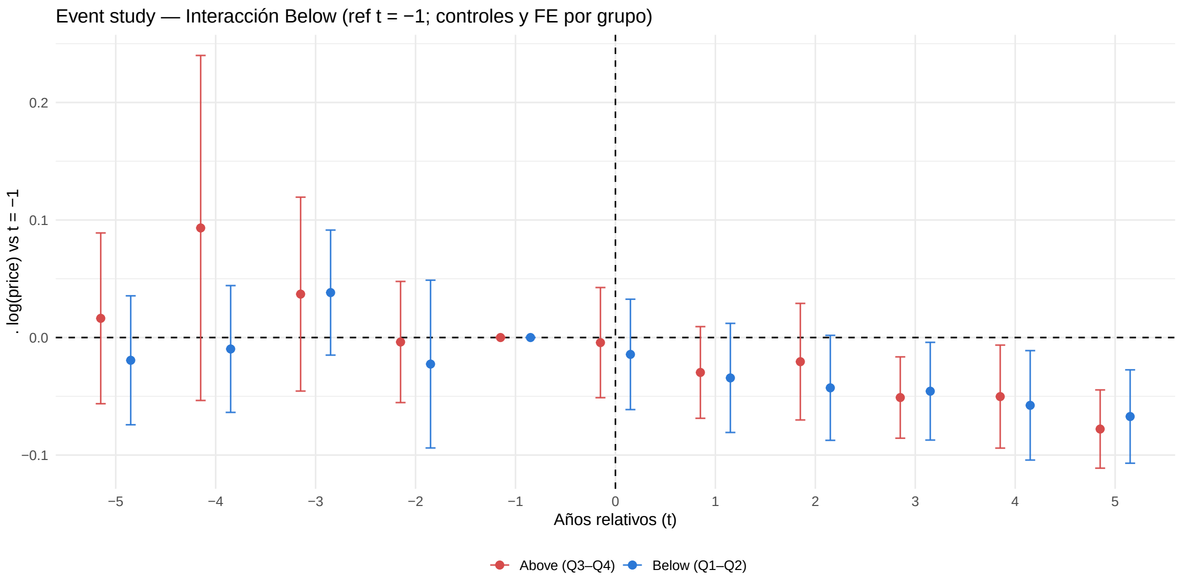

(a) Event study by income (ref. $t = -1$).

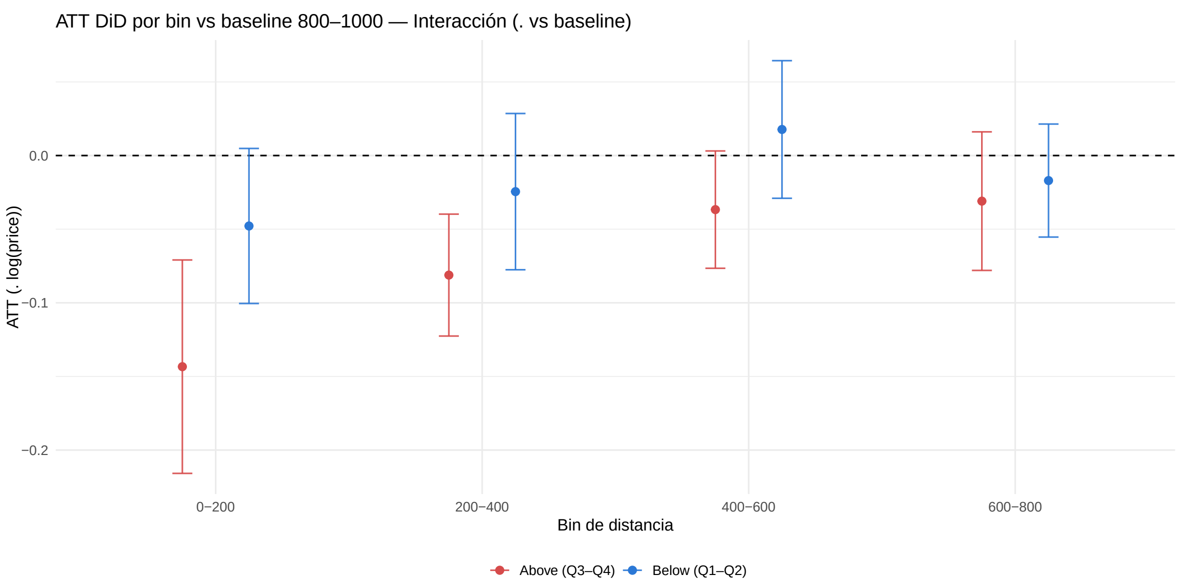

(b) Distance decay by income (vs. 800–1,000 m baseline).

Figure 3: Local sales price responses around new affordable projects

Notes: Panel (a) plots $\widehat{\beta}_k$ from the event study version of (5.1) (leads/lags of Near×Year; baseline $t = -1$), separately for Above (Q3–Q4) and Below (Q1–Q2). Panel (b) reports ATTs for 0–200, 200–400, 400–600, 600–800 m rings relative to 800–1,000 m, by income group. All specs include unit and year FEs, rich property controls, and two-way clustered SEs (block & year). Key point estimates: overall ATT -0.061 (SE 0.017); Above: -0.075 (SE 0.025); Below: -0.020 (SE 0.012); SF: -0.067 (SE 0.018); MF: 0.004 (SE 0.043).

Taken together, these results are consistent with localized disamenity capitalization in affluent single-family contexts and near-zero average capitalization elsewhere, aligning with the heterogeneous take-up patterns in the main analysis and mapping closely to functional area G1 in Section 4.2.2.

5.2 Not in my development: rents and construction costs

This subsection turns to the developer utility used in Section 2.2. I first use CoStar cross-sectional data on (project-level average rents to document how mixed-income composition relates to achieved rents; second, I then relate neighborhood rent bins to construction costs using assessor/SBF microdata. Because both sources are cross-sectional, the evidence is descriptive; nevertheless, it helps interpret how rents and costs—and thus profits—vary when developers include affordable set-asides.

Let r^M denote market rent, r^A the regulated rent, and $s \in [0,1]$ the affordable share within a project. The average rent in the building is

$$\bar{r} = (1-s)r^M + sr^A \quad \Rightarrow \quad \Delta \ln \bar{r} \equiv \ln \bar{r} - \ln r^M \ = \ \ln\Big(1 - s(1-\rho)\Big), \quad \rho \equiv r^A/r^M.$$

For small changes, $\Delta \ln \bar{r} \approx -s(1-\rho)$. For example, if r^M=\$3,000, r^A≈\$1,900 (ρ≈0.63), and s=0.20, the implied dilution is about 7.4%.Table A.12 reports $\Delta \ln \bar{r}$ across rent bins and set-aside shares using HUD rent schedules and observed CoStar unit mixes.

Consistent with this identity, project-level regressions of log(average rent) on an Affordable/Mixed indicator—controlling for unit mix, size, class/quality, vintage, vacancy, and fixed effects—show a sizable penalty in above-median tracts (−0.132), while below-median tracts are near zero (Table A.11). Of course, these are correlations, not causal estimates; but they help to understand how rents vary across neighborhoods, which goes in line with previous research (Davidoff et al., 2022; Diamond & Mcquade, 2019). Similarly, the assessor/SBF microdata further indicate that both normalized quality and per-unit costs rise monotonically with neighborhood rents. Relative to Quartile 1, coefficients for Quartiles 2–4 are (0.225, 0.518, 0.732) in quality units and (\$185, \$421, \$945) per unit, respectively (Table A.13).

To connect magnitudes to investment returns, I simulate a pro forma development following Zhu et al. (2021) and Phillips (2024). Figure 4 plots ΔIRR ≡ IRR(TOC) − IRR(baseline) over two grids. On the left, with no depreciation of market rents inside the mixed project, ΔIRR > 0 at monthly rents below \$2,000; introducing a 5% depreciation of market-rent units inside the mixed building shifts the threshold down to roughly \$1,500. On the right, increasing the affordable share s lowers ΔIRR at higher rents but raises it at lower rents when bonus market-rate units are added, mirroring the empirical pattern that mixed-income take-up is concentrated where market and regulated rents are closer.

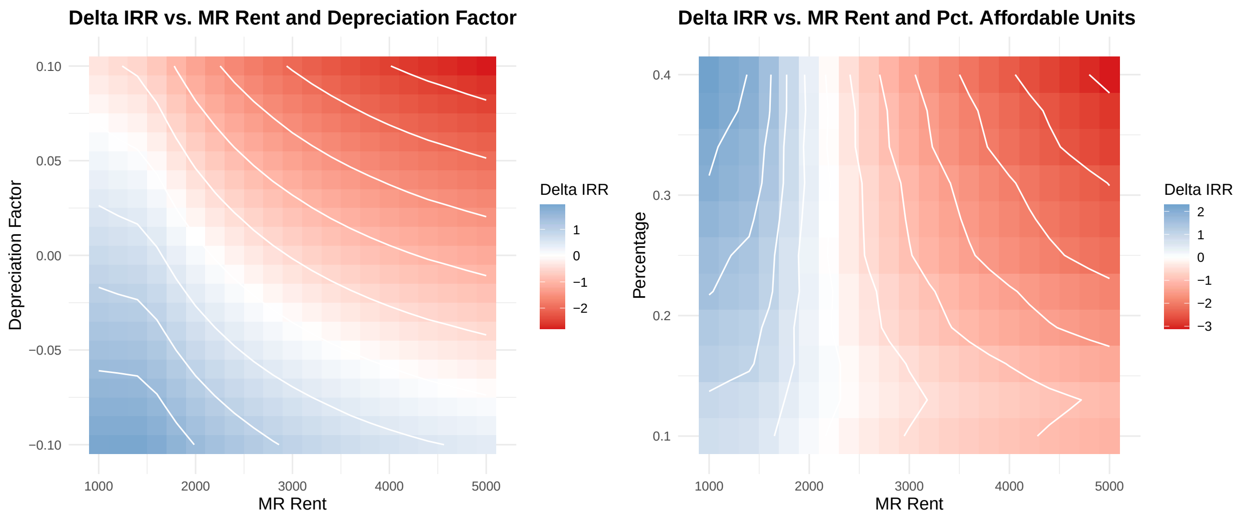

Figure 4: **ΔIRR sensitivity.** Each panel reports $\Delta\text{IRR} = \text{IRR(TOC)} - \text{IRR(baseline)}$ from a Monte Carlo pro forma calibrated to Zhu et al. (2021) and Phillips (2024). *Left:* grid over market rent (r^M) and a depreciation factor applied to market-rate units inside the mixed project; contours at 1 pp intervals. *Right:* grid over r^M and affordable share s (with bonus market-rate units). Positive values (blue) indicate higher IRR under TOC than baseline. Additional settings and draws in Appendix §2.9.2.

Together with the localized price capitalization in Section 5.1, these facts rationalize why TOC's positive permitting effects concentrate in below-median tracts. Where r^M is closer to r^A, the set-aside dilution is smaller and input costs are lower; in high-rent, single-family contexts, a larger rent wedge and higher costs—potentially coupled with localized opposition[10] —tighten feasibility. Developers therefore adopt TOC where the rent gap is smaller and bonuses translate more directly into delivered units.

6 Market-wide Building Responses

In this section, I examine whether TOC's positive effects reshape the built environment. To do so, I study building permits and total housing stock within the 500 m bandwidth and citywide. Because the decision to file building permits (BP)—new builds, demolitions,

[10]There have been controversies around TOC as some neighbors have argued about its negative effects on the places they live in. Predominantly single-family neighborhoods have organized against "tall and dense" TOC developments. One example is the lawsuit against the Santa Monica project, which, according to local inhabitants, is expected to reach 79 feet instead of the 57 that would be allowed under the existing zoning ordinance. More details here, as they can change the neighborhood characteristics. These perceptions should negatively impact housing sales near NIMBY neighbors.

additions, and alterations—can be endogenous to both affordable housing (AH) activity and TOC exposure, I use propensity-score matching (PSM) to control for baseline pre-treatment characteristics.

6.1 Market-wide building responses beyond AH: PSM + DiD within 500 m

Because AH responds to TOC (Section 4.1), a naïve treated–control comparison could confound direct TOC effects with selection into construction. To sharpen the counterfactual, I construct a matched control group within the same ±500 m band using PSM based on pre-policy characteristics (socio-demographics, baseline permitting mix, and distance bins), and then re-estimate DiD on the matched panel. Figure A.4 documents covariate balance: standardized mean differences are uniformly below 0.10 across covariates.

Propensity–score matching DiD corroborates the baseline pattern. Within the ±500 m window (Table 8), effects are small or imprecise in AL and AH and concentrate in BL, with the largest effects in BH, for probability, project counts, and units. Expanding the window to ±4 km preserves these gradients with similar magnitudes (Table A.15). Taken together, the results suggest the effects are not driven by observable composition and are robust to the spatial bandwidth used.

For building permits, the PSM–DiD indicates increases in additions and demolitions in BL/BH, with no systematic change in new developments. By unit type, apartment units rise while 1–2 family and commercial units do not (Table A.14); in the broader window, demolitions and additions follow the same pattern, with no increase in apartments (Table A.16). This points to a reconfiguration margin of the multifamily stock—rather than entirely new buildings—in the locations where TOC is most feasible.

Table 8: PSM–DiD (±500 m): efectos por AL/BL/AH/BH en resultados principales

	AL	BL	AH	BH
Panel A. Probability of any AH				
ATT	0.006	0.014***	0.015	0.037**
	(0.005)	(0.006)	(0.024)	(0.015)
Panel B. Count of AH projects				
ATT	0.007	0.017***	0.028	0.034**
	(0.006)	(0.006)	(0.027)	(0.015)
Panel C. Number of AH units				
ATT	0.036	0.222***	0.109	0.561*
	(0.041)	(0.071)	(0.257)	(0.287)

Notes: Propensity-score matching dentro de ±500 m y DiD con efectos fijos de block group y año; SE agrupados por block group. AL=Above×Low (tiers 1–2); BL=Below×Low; AH=Above×High (tiers 3–4); BH=Below×High. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.2 Housing supply

I study whether TOC changes the housing stock by estimating PSM–DiD on block–group per year aggregates of units. Outcomes are first differences of log counts (growth rates) for total units, single–family (SF), and multifamily (MF). I report results within the ±500 m window and, for context, at ±4 km. Specifications include block–group and year fixed effects; standard errors are clustered by block group. These are area–level aggregates and also rely on PSM.

Across both bandwidths, the estimates for $\Delta \ln$ total units and $\Delta \ln$ multifamily are close to zero and statistically indistinguishable from zero. This is consistent with the permit evidence: TOC coincides with more additions and demolitions but not with systematic increases in new ground–up activity, suggesting reconfiguration or replacement of existing stock rather than net expansion over the sample horizon. For single–family, the above–median stratum shows a small, borderline negative growth effect, while the below–median stratum does not. Given the magnitude, these estimates should be interpreted cautiously; they may reflect higher baseline SF shares in above–median areas or timing/compositional dynamics, rather than a sizeable change in SF supply.

Table 9: Housing supply (growth rates), PSM–DiD within ±500 m

	Above median	Below median
Panel A. $\Delta \ln$ total units		
ATT	-0.000	-0.000
	(0.001)	(0.002)
Panel B. $\Delta \ln$ SF units		
ATT	-0.001^{**}	0.001
	(0.001)	(0.001)
Panel C. $\Delta \ln$ MF units		
ATT	-0.000	0.001
	(0.002)	(0.003)

Notes: Outcomes are first differences of logs of block–group–level unit counts (total, SF, MF). PSM builds a matched control within the same band on pre–policy covariates; DiD includes block–group and year fixed effects; SEs clustered by block group. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7 Robustness checks

This section shows that the estimates are not an artifact of bandwidth choice, local spillovers, or spatial correlation in the errors. I (i) sweep bandwidths from 100 m to 1 km and run "donut" DiD that removes the innermost observations; (ii) test for spatial gradients and internal displacement; (iii) use shifted (fake) boundaries as placebos; (iv) compute Conley (spatial HAC) standard errors; (v) validate the kink–DiD with pre-trend diagnostics, placebo timing, and by excluding potential anticipation years; and (vi) test for effects on affordable housing when isolating DB activity.

7.1 Bandwidth sweeps and donut DiD

Table A.18 reports ATTs for the main outcomes when varying the symmetric window around the boundary from 100 m to 1,000 m, separately for below- and above-median areas. Effects grow with the window and are consistently larger in below-median neighborhoods. Table A.19 shows "donut" estimates that exclude observations within $d \in \{25, 50, 100, 150\}$ m of the boundary; results are stable or slightly larger.

7.2 Spatial gradient and (lack of) displacement

To test whether the policy simply relocates activity right across the line, I (i) estimate ATTs by concentric rings and (ii) run "internal displacement" tests that compare near vs. far rings on each side. Table A.20 shows a smooth interior gradient: effects inside TOC persist and grow with distance from the boundary up to 2 km. Wald tests fail to reject that the sum of ring ATTs is zero when comparing near vs. far on each side—evidence against short-run displacement across the boundary. I also shift the boundary by ± 1 km and re-estimate the DiD. Placebo ATTs are small and imprecise (Table A.21), supporting the identification.

7.3 Spatial HAC (Conley) standard errors and heterogeneity

Table A.22 reports ± 500 m DiD with Conley correction. Effects remain highly significant. Below-median areas respond much more (Wald test for Below$-$Above is large and significant).

7.4 Kink–DiD validation: pre-trends, placebos, and timing

I validate the boundary kink–DiD in four steps. First, pre-period level RD at the boundary is near zero and pre slopes are statistically similar across sides (Table A.23). Second, a 2017 placebo finds no "jump" or "kink" before TOC (Table A.24). Third, excluding 2016–2017 (potential anticipation) strengthens the post boundary kink for all outcomes (Table A.25).

Finally, the kink–DiD with Conley errors delivers a precise post-slope difference (kink) of 0.056 (SE = 0.021) for *prob_aff*, consistent with the bin–means overlay in Figure 2b.

7.5 DB-only placebo: are main effects driven by DB?

A concern is that the boundary effects might be driven by contemporaneous use of the pre-existing Density Bonus (DB) rather than by TOC. To probe this, I re-define outcomes to isolate *DB-only* activity and re-estimate the baseline DiD within ± 500 m using block–group and year fixed effects, clustering at the block–group level. If TOC is the active margin inside the boundary, we should see no positive DB effect (and possibly a small decline consistent with program substitution).

Figure 5 reports an event study for the three DB-only outcomes. Pre-trends are flat and centered at zero; after 2017, coefficients remain small and statistically indistinguishable from zero, with no systematic post jump.

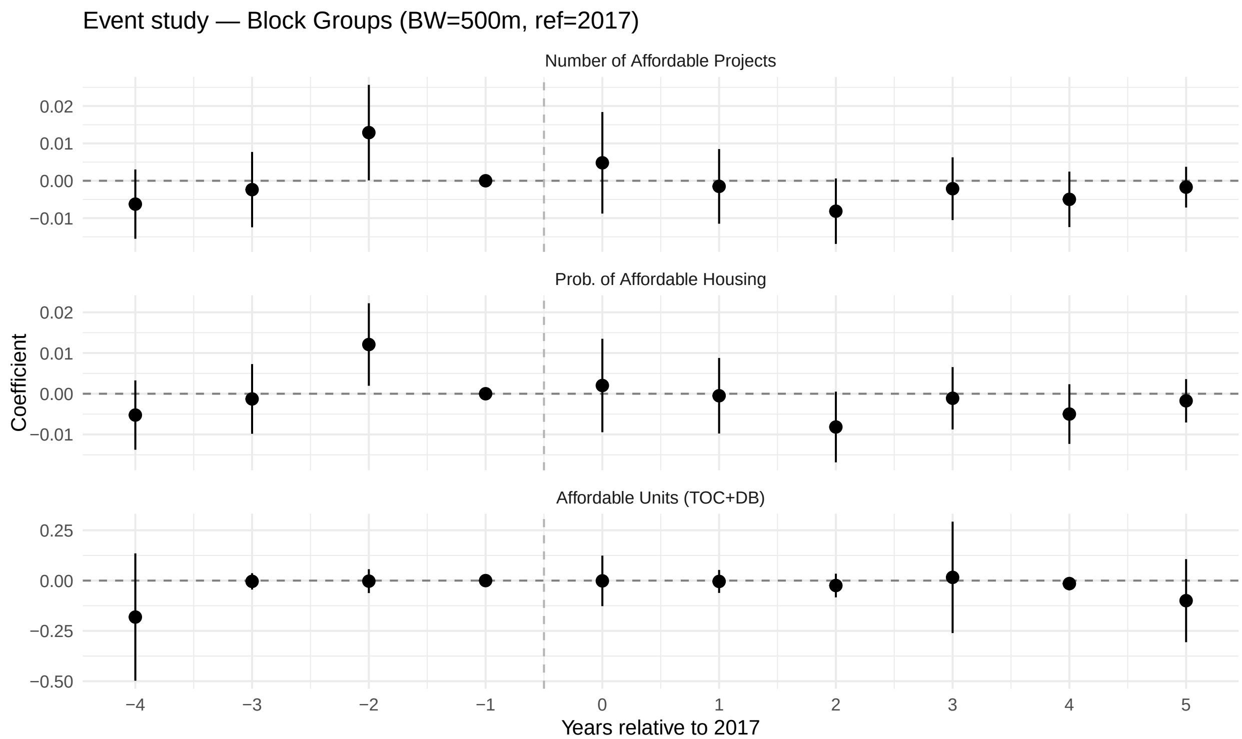

Figure 5: Event study (DB-only outcomes), ±500 m around the boundary; reference year 2017.

Notes: Block–group and year fixed effects; standard errors clustered by block group. Outcomes isolate DB activity as defined in the text.

Table A.26 summarizes the DiD coefficients. Inside the boundary, post–TOC, both the probability of a DB project and the number of DB permits *decline* by roughly 0.9–1.0 p.p. relative to the control side; DB units show no significant change. Given pre-policy means of 1.3% (probability), 1.4% (permits), and 5.8% (units), these magnitudes are small but consistent with developers substituting away from DB toward TOC in the eligible area. This placebo thus reinforces that the main positive effects I document are *TOC-driven* rather than an artifact of DB.

7.6 Aggregation checks (blocks and tracts)

The DiD is robust to the spatial unit of aggregation. Effects remain positive at the block level (Table A.27) and at the tract level (Table A.28), while above-median tract estimates are smaller and less precise, mirroring the baseline gradients.

8 Conclusion and Policy Implications

This paper exploits sharp eligibility boundaries and granular microdata to evaluate Los Angeles's Transit Oriented Communities (TOC) program. Within ± 500 m of the boundary, TOC raises the annual probability that a block group receives any affordable project by roughly 1.0–1.3 percentage points over a pre-policy mean near 1.3%, with parallel gains in project counts and affordable units. A boundary kink–in–DiD shows that these averages are not driven by a discrete jump at the edge but by a positive interior gradient on the treated side, consistent with the tiered geography of incentives.

Disaggregated estimates reveal that take-up concentrates in below-median-rent, multifamily settings and in higher-intensity tiers (3–4), whereas affluent, single-family areas respond weakly. The pattern fits a simple threshold: adoption occurs when a monetized effective bonus—capacity, concessions, and streamlining—exceeds the "inclusion tax" implied by the local market–subsidy rent wedge. Importantly, incentive size operates as an amplifier rather than a substitute: larger tiers magnify effects where wedges are modest and feasibility is tight, but do little where wedges remain large. This interaction between rent gaps and bonus intensity is a design margin largely under-emphasized in the IZ literature and helps reconcile heterogeneous findings across places and programs.

Two mechanisms rationalize these gradients. First, localized capitalization dampens take-up in high-rent single-family contexts: near–far DiD around covenant timing indicates small, short-ranged discounts of roughly 6–8% within 500 m for nearby single-family sales in above-median areas, effects that attenuate with distance and time, while capitalization is near zero elsewhere. Second, within mixed-income projects, the arithmetic of mixing lowers average achievable rents more when the wedge is large, and assessor-based costs/quality rise with neighborhood rents, tightening feasibility where wedges are greatest. These channels align with developer-feasibility perspectives Phillips (2024) and Zhu et al. (2021) and with evidence on localized capitalization around affordable housing (Diamond & Mcquade, 2019).

On market-wide activity, matched-panel DiD within ± 500 m and ± 4 km points to composition rather than scale. Additions and demolitions rise in below-median, high-tier cells, but there is no systematic increase in new ground-up permitting and no detectable change in the growth rate of total or multifamily stock over the study horizon. A DB-only placebo shows small post-2018 declines in DB use inside eligibility, consistent with substitution toward TOC rather than a broader permitting surge. In this sense, TOC directly raises the quantity of affordable production without shifting the near-term stock trajectory; at the same time, siting near transit improves quality along the access margin, plausibly lowering generalized transport costs even if neighborhood socioeconomic composition does not change.

These results resonate with a wider body of work showing that supply effects from relaxing land-use constraints emerge most strongly where density is already viable and constraints bind (Dong, 2021; Gabbe, 2018; Greenaway-McGrevy & Phillips, 2023). They complement the mixed inclusionary-zoning (IZ) literature by providing boundary-based causal estimates that highlight market dependence and program design (Bento et al., 2009; Freeman & Schuetz, 2017; Schuetz et al., 2009; Wang & Fu, 2022), and they are consistent with the broader view that stringent land-use controls elevate prices and dampen supply in constrained metros (Glaeser & Gyourko, 2003, 2018; Hilber & Schöni, 2022; Quigley et al., 2005; Saks, 2008). On incentive magnitude, cross-jurisdictional evidence associates density bonuses with a higher likelihood of producing any units but offers mixed links to pushing policies into the top productivity tier; the boundary design here reconciles this by showing that larger effective bonuses steepen interior responses only where rent wedges are not too large and the built form is multifamily (Wang & Fu, 2022). In short, bonus size is an amplifier—not a substitute—for alignment with local feasibility.

The policy interpretation is therefore specific. TOC behaves like a calibrated, transit-proximate instrument: where rent gaps are modest and density is viable, tiered bonuses translate into delivered affordable units; where wedges and costs are high, identical set-asides and envelopes underperform. Designing around this threshold—by calibrating effective capacity and process relief to local wedges, and by pairing regulatory menus with finance-side complements where stacking is restricted—can raise delivery without broad fiscal outlays. In high-wedge, single-family contexts, meaningful spatial integration likely requires either larger effective bonuses or targeted fiscal complements; absent those, stronger by-right pathways, participation, and design standards may reduce frictions and delay but will not overturn the arithmetic of the wedge. Because the program concentrates production in well-connected, multifamily areas, transparent monitoring and periodic recalibration of tiers and concessions are natural next steps.

Policy trade-offs follow. On the "quantity" margin, TOC directly delivers affordable units where feasible, but because the program does not reconfigure high-opportunity single-family enclaves, it does not substantially alter citywide stock composition over the horizon studied. On the "quality" margin, TOC concentrates affordable housing in transit-proximate, amenity-rich areas, likely reducing commuting costs and improving access—gains that matter even absent large compositional change. Where localized price fears fuel opposition, pairing IZ with streetscape/parking management, design standards, and tenant protections/anti-displacement tools can mitigate salient concerns and sustain approvals. Finally, the non-stacking constraint argues for coordination with finance-side tools (gap financing, site acquisition) to substitute for foregone regulatory stacking when wedges are large.

Several limitations qualify these conclusions. Estimates are local to the boundary and to 2018–2022. TOC is a bundle—capacity, waivers, and streamlining—so component-level causal weights are not separately identified. Donut windows and ring tests show no short-run displacement within $\pm 500\,\mathrm{m}$, but broader reallocation cannot be ruled out, and general-equilibrium feedbacks on land values, migration, and labor markets lie outside the design. Within these bounds, isolating the interior gradient and the bonus-as-amplifier mechanism provides a boundary-based account of voluntary IZ that links market conditions to program design in ways that are actionable for jurisdictions contemplating transit-proximate, tiered incentives.

References

Bento, A., Lowe, S. E., Knaap, G.-J., & Chakraborty, A. (2009). *Housing Market Effects of Inclusionary Zoning* (tech. rep.).

Davidoff, T., Pavlov, A., & Somerville, T. (2022). Not in my neighbour's back yard? Laneway homes and neighbours' property values. *Journal of Urban Economics*, *128*, 103405. https://doi.org/10.1016/J.JUE.2021.103405

Desmond, M. (2018). Heavy is the House: Rent Burden among the American Urban Poor. *International Journal of Urban and Regional Research*, *42*(1), 160–170. https://doi.org/10.1111/1468-2427.12529

Diamond, R., & Mcquade, T. (2019). *Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development* (tech. rep.). https://web.archive.org/web

Dong, H. (2021). Exploring the Impacts of Zoning and Upzoning on Housing Development: A Quasi-experimental Analysis at the Parcel Level. *Journal of Planning Education and Research*. https://doi.org/10.1177/0739456X21990728/FORMAT/EPUB

Freeman, L., & Schuetz, J. (2017). *Producing Affordable Housing in Rising Markets: What Works?* (Tech. rep.).

Gabbe, C. J. (2018). How Do Developers Respond to Land Use Regulations? An Analysis of New Housing in Los Angeles. *Housing Policy Debate*, *28*(3), 411–427. https://doi.org/10.1080/10511482.2017.1368031

Glaeser, E., & Gyourko, J. (2003). The impact of building restrictions on housing affordability. *Economic Policy Review*, (Jun), 21–39. https://ideas.repec.org/a/fip/fednep/y2003ijunp21-39nv.9no.2.html%20https://ideas.repec.org//a/fip/fednep/y2003ijunp21-39nv.9no.2.html

Glaeser, E., & Gyourko, J. (2018). The economic implications of housing supply. *Journal of Economic Perspectives*, *32*(1), 3–29. https://doi.org/10.1257/jep.32.1.3

Greenaway-McGrevy, R., & Phillips, P. C. (2023). The impact of upzoning on housing construction in Auckland. *Journal of Urban Economics*, *136*, 103555. https://doi.org/10.1016/J.JUE.2023.103555

Hilber, C., & Schöni, O. (2022). Housing Policy and Affordable Housing. *Oxford Research Encyclopedia of Economics and Finance*. https://doi.org/10.1093/ACREFORE/9780190625979.013.829

Phillips, S. (2024). *Modeling Inclusionary Zoning's Impact on Housing Production in Los Angeles: Tradeoffs and Policy Implications* (tech. rep.). www.ternercenter.berkeley.edu

Quigley, J. M., Raphael, S., Ulsen, E., Mayer, C., & Schill, M. (2005). Regulation and the High Cost of Housing in California. *American Economic Review*, *95*(2), 323–328. https://doi.org/10.1257/000282805774670293

Saiz, A. (2010). *THE GEOGRAPHIC DETERMINANTS OF HOUSING SUPPLY* (tech. rep.). http://qje.oxfordjournals.org/

Saks, R. E. (2008). Job creation and housing construction: Constraints on metropolitan area employment growth. *Journal of Urban Economics*, *64*(1), 178–195. https://doi.org/10.1016/J.JUE.2007.12.003

Schuetz, J., & Meltzer, R. (2012). Inclusionary Zoning to Support Affordable Housing.

Schuetz, J., Meltzer, R., & Been, V. (2009). 31 flavors of inclusionary zoning: Comparing policies from San Francisco, Washington, DC, and suburban Boston. *Journal of the American Planning Association*, *75*(4), 441–456. https://doi.org/10.1080/01944360903146806

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. https://doi.org/10.1080/01944363.2022.2027263

Zhu, L., Burinskiy, E., De La Roca, J., Green, R. K., & Boarnet, M. G. (2021). *Cityscape: A Journal of Policy Development and Research - Regulatory Reform and Affordable Housing - Los Angeles' Housing Crisis and Local Planning Responses: An Evaluation of Inclusionary Zoning and the Transit-Oriented Communities Plan as Policy Solutions in Los Angeles* (tech. rep.). https://www.census.gov/housing/hvs/data/histtabs.html

Appendix

Appendix: Derivations and comparative statics

Assume the affordable quota binds under TOC, so $q_a = \alpha_{\tau,g} q$ and $q_m = (1 - \alpha_{\tau,g})q$, with $\alpha_{\tau,g} \in (0, 1)$. Define $\Delta p_{i,g} \equiv p_i^M - p_{ig}^A \geq 0$.

The baseline value is

$$\Pi_i^{\text{base}} = \max_{0 \leq q \leq q_i^0} \{\, p_i^M q - C_i(q; 0) \,\}.$$

The TOC value at tier τ is

$$\Pi_{i\tau}^{\text{toc}} = \max_{0 \leq q \leq q_{i\tau}^1} \left\{ \left[p_i^M - \alpha_{\tau,g} \Delta p_{i,g} \right] q - C_i(q; k_\tau) + S_{i\tau} + \lambda_i p_i^M (1 - \alpha_{\tau,g}) q \right\}.$$

Let $q_{i\tau}^*$ and q_i^{0*} denote the optimal choices under TOC and baseline. The value difference can be written as

$$\Delta\Pi_{i\tau g} \equiv \Pi_{i\tau}^{\text{toc}} - \Pi_i^{\text{base}} = \underbrace{p_i^M (q_{i\tau}^* - q_i^{0*}) + S_{i\tau} - \left[C_i(q_{i\tau}^*; k_\tau) - C_i(q_i^{0*}; 0) \right] + \lambda_i p_i^M (1 - \alpha_{\tau,g}) q_{i\tau}^*}_{\text{Bonus}_{i\tau g}} - \alpha_{\tau,g} \Delta p_{i,g}\, q_{i\tau}^*,$$

which implies the adoption threshold

$$\text{Bonus}_{i\tau g} \geq \alpha_{\tau,g} \Delta p_{i,g}\, q_{i\tau}^*.$$

Holding $q_{i\tau}^*$ fixed, envelope arguments yield

$$\begin{aligned}
\frac{\partial \Delta\Pi_{i\tau g}}{\partial \Delta p_{i,g}} &= -\,\alpha_{\tau,g}\, q_{i\tau}^* \;<\; 0, \\
\frac{\partial \Delta\Pi_{i\tau g}}{\partial k_\tau} &= -\frac{\partial C_i(q_{i\tau}^*; k_\tau)}{\partial k_\tau} \;>\; 0 \quad \text{if } \frac{\partial C_i}{\partial k_\tau} < 0, \\
\frac{\partial \Delta\Pi_{i\tau g}}{\partial \lambda_i} &= p_i^M (1 - \alpha_{\tau,g})\, q_{i\tau}^* \;>\; 0.
\end{aligned}$$

A tier can raise or lower value depending on how capacity relief and regulatory savings compare with the larger set-aside burden that may come with higher tiers. One convenient way to express this is

$$\frac{\partial \Delta\Pi_{i\tau g}}{\partial \tau} = S_{i\tau}' - \left(\frac{\partial C_i}{\partial k_\tau} \right) k_\tau' + \mu_{i\tau} (q_{i\tau}^1)' - q_{i\tau}^*\, \alpha_{\tau,g}' \left[\Delta p_{i,g} + \lambda_i p_i^M \right],$$

where $\mu_{i\tau} \geq 0$ is the multiplier on the capacity constraint $q \leq q^1_{i\tau}$. Higher tiers increase adoption incentives when marginal capacity relief, concessions, and regulatory savings dominate the marginal increase in the set-aside burden scaled by local gaps and capitalization.

Finally, when the quota binds, TOC expands both market-rate and affordable units relative to the baseline whenever

$$q_m - q_i^0 = (1 - \alpha_{\tau,g})\, q^1_{i\tau} - q_i^0 > 0 \quad \Longleftrightarrow \quad \frac{q^1_{i\tau}}{q_i^0} > \frac{1}{1 - \alpha_{\tau,g}}.$$

Appendix: Additional institutional details on TOC

This appendix records additional features of the TOC Guidelines that are not required to follow the identification strategy.

Eligibility and tier assignment. TOC Incentive Areas are defined within one-half mile of a qualifying Major Transit Stop. A Major Transit Stop is either a rail station or the intersection of two or more qualifying bus routes meeting peak-period frequency thresholds. Parcels are assigned to tiers (1 through 4) based on the shortest distance from the parcel to the qualifying stop, as detailed in the Guidelines.

Incentive menu. Projects that meet the on-site affordability requirement are eligible for base incentives that increase residential density and floor-area ratio. Parking minimums are reduced substantially, with the largest reductions in Tier 4. Projects may also request additional incentives such as yard and setback reductions, open-space relief, and height increases. Projects meeting specified labor standards may request a larger number of additional incentives. Fully affordable projects may qualify for a tier bump, as defined by the Guidelines.

Affordability requirements and compliance. Minimum set-aside shares vary by tier and by income target group. Rents for each income category are tied to area median income thresholds specified in the Guidelines. Compliance is secured through a recorded covenant prior to the issuance of the building permit, requiring affordability for the duration specified in the program rules.

Administrative process. TOC provides procedural benefits and may reduce approval timelines relative to comparable projects outside the program. I treat these administrative provisions as part of the bundled treatment in the empirical analysis.

Program activity. For descriptive context, the City Planning dashboard reports aggregate counts of approved TOC units and the affordable share over time. These figures are not used for identification and are included only as background.

Tables

Table A.1: Logit: Any Affordable Housing by Block Group (citywide)

	Dependent variable: Any affordable project (prob.)
Log(pop.)	0.563^{***}
	(0.085)
Share Asian	0.327
	(0.333)
Share Black	0.337
	(0.318)
Share Hispanic	0.618^{**}
	(0.300)
Educational attainment	-1.951^{***}
	(0.402)
Female-headed HH	-1.290^{**}
	(0.511)
Share renters	2.911^{***}
	(0.199)
Share poverty	0.687^{*}
	(0.380)
Median rent	-0.00004
	(0.0001)
Constant	-6.637^{***}
	(0.614)
Observations	4,963

Notes: Logit coefficients; robust (HC1) standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.2: Pre-policy balance: citywide (block groups)

	CTRL mean	TRT mean	Diff (TRT−CTRL)	p-value
Panel A. Sociodemographics (2014)				
Log(Pop)	7.220	7.220	0.004	0.781
Pct. Asian	0.112	0.106	−0.005	0.167
Pct. Black	0.099	0.137	0.038***	0.000
Pct. White	0.575	0.481	−0.094***	0.000
Pct. Hispanic	0.455	0.481	0.027**	0.003
Low education	0.622	0.650	0.028***	0.000
Female HH	0.164	0.179	0.015***	0.000
Pct. renters	0.506	0.661	0.155***	0.000
Pct. poverty	0.171	0.237	0.066***	0.000
Median rent (USD)	1,138	1,105	−33.1**	0.035
Panel B. Affordable units per block group (mean 2010–2017)				
Aff. units per BG	0.033	0.114	0.081***	0.000
Pct. affordable	0.118	0.149	0.031*	0.093
Panel C. Building permits (annual mean 2010–2017)				
Additions (count)	0.739	0.514	−0.225***	0.000
Alterations (count)	3.440	4.266	0.826***	0.000
Demolitions (count)	0.179	0.302	0.122***	0.000
New construction	0.366	0.379	0.013	0.436
1–2 fam. units	0.112	0.105	−0.007	0.648
Apartment units	0.332	1.619	1.288***	0.000
Panel D. Housing stock (pre-policy, 2016)				
Units, total	846.0	411.0	−435.0***	0.001
Units, single-family	573.0	142.0	−431.0***	0.000
Units, multifamily	273.0	269.0	−4.2	0.897
Parcels, total	583.0	162.0	−420.0***	0.000
Parcels, single-family	539.0	123.0	−416.0***	0.000
Parcels, multifamily	44.1	39.3	−4.8	0.350

Notes: Means are pre-policy. Panel A reports 2014 sociodemographic covariates; Panel B reports pre-policy (2010–2017) averages of affordable units; Panel C reports annual averages of building permits over 2010–2017; Panel D reports 2016 housing stock. The sample includes all block groups in Los Angeles. Sociodemographic and affordable-housing variables are much closer across treatment and control than building-permit and housing-stock variables, which show sizable differences consistent with treated block groups being located in denser, more multifamily neighborhoods. Stars reflect two-sided tests of equal means. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.3: Heterogeneity: Above vs. Below median rents (0–500 m)

	Above median	Below median
Panel A. Probability of any AH		
ATT	0.006	0.017***
	(0.005)	(0.005)
Observations	17,091	17,091
R^2 (within)	0.197	0.197
Panel B. Count of AH projects		
ATT	0.008	0.019***
	(0.006)	(0.006)
Observations	17,091	17,091
R^2 (within)	0.207	0.207
Panel C. Number of AH units		
ATT	0.039	0.258***
	(0.045)	(0.072)
Observations	17,091	17,091
R^2 (within)	0.140	0.140

Notes: Block–group×year panel within 0–500 m; block–group and year fixed effects; SEs clustered by block group. Aggregates over tiers within each rent stratum. Wald summary: Probability—BH>AL ($p = 0.049$); Count—BH>AL ($p = 0.077$); Units—BL>AL ($p = 0.0056$). Joint tests (DiD×Below×HighTier): global $DB{=}DH{=}DBH{=}0$ significant only for Units ($p = 0.013$); "Only Below" $DB{=}DBH{=}0$ significant for Units ($p = 0.0106$); "Only High tier" $DH{=}DBH{=}0$ not significant. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4: Heterogeneity: Rent Quartiles × TOC Eligibility (0–500 m)

	Q1	Q2	Q3	Q4
Panel A. Probability of any AH				
ATT	0.015**	0.019**	0.006	0.005
	(0.006)	(0.008)	(0.006)	(0.007)
Panel B. Count of AH projects				
ATT	0.017***	0.021**	0.008	0.008
	(0.006)	(0.009)	(0.007)	(0.008)
Panel C. Number of AH units				
ATT	0.203**	0.310***	0.009	0.070
	(0.082)	(0.109)	(0.058)	(0.054)

Notes: Q1–Q4 index block groups by quartiles of pre-policy median rent (Q1 lowest, Q4 highest). Entries are ATTs from difference-in-differences specifications; SEs in parentheses clustered by block group. Wald summary: For Panels A and B, pairwise differences in ATTs across quartiles are not statistically significant ($p > 0.17$). For Panel C, Q3 has a significantly lower ATT than Q1 and Q2 ($p = 0.031$ and $p = 0.009$), and Q4 < Q2 ($p = 0.035$); other pairwise differences are not significant. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Heterogeneity: Rent k-means Bins × TOC Eligibility (0–500 m)

	K1	K2	K3	K4
Panel A. Probability of any AH				
ATT	0.014***	0.011	0.003	0.004
	(0.005)	(0.008)	(0.010)	(0.008)
Panel B. Count of AH projects				
ATT	0.015***	0.014	0.004	0.005
	(0.005)	(0.010)	(0.010)	(0.008)
Panel C. Number of AH units				
ATT	0.196***	0.099	0.053	-0.015
	(0.063)	(0.065)	(0.061)	(0.040)

Notes: K1–K4 index block groups by k-means clusters of pre-policy median rent: K1 $\leq$ 1,200, K2 1,200–1,600, K3 1,600–2,200, K4 2,200–3,600. Entries are ATTs from difference-in-differences specifications; SEs in parentheses clustered by block group. Wald summary: For Panels A and B, pairwise differences in ATTs across bins are not statistically significant ($p > 0.15$). For Panel C, K4 has a significantly lower ATT than K1 ($p < 0.001$), while K3 < K1 and K4 < K2 are marginally significant ($p \approx 0.06$); other pairwise differences are not significant. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.6: Heterogeneity: functional areas (G1–G3) × high-intensity eligibility (*hi_any*)

	G1: hi=0	G1: hi=1	G2: hi=0	G2: hi=1	G3: hi=0	G3: hi=1
Panel A. Probability of any AH						
ATT	0.002	0.018	0.025***	0.035**	0.015**	0.021
	(0.005)	(0.021)	(0.006)	(0.018)	(0.006)	(0.025)
Observations	17,091	17,091	17,091	17,091	17,091	17,091
R^2 (within)	0.198	0.198	0.198	0.198	0.198	0.198
Panel B. Count of AH projects						
ATT	0.003	0.025	0.028***	0.036**	0.017**	0.022
	(0.006)	(0.023)	(0.007)	(0.018)	(0.009)	(0.025)
Observations	17,091	17,091	17,091	17,091	17,091	17,091
R^2 (within)	0.207	0.207	0.207	0.207	0.207	0.207
Panel C. Number of AH units						
ATT	0.029	0.174	0.263***	0.544	0.306*	0.295
	(0.039)	(0.159)	(0.091)	(0.402)	(0.172)	(0.302)
Observations	17,091	17,091	17,091	17,091	17,091	17,091
R^2 (within)	0.140	0.140	0.140	0.140	0.140	0.140

Notes: Block-group×year panel within 0–500 m of the TOC boundary; block-group and year fixed effects; standard errors clustered by block group. Cells report post-treatment ATTs by functional area $g \in \{1, 2, 3\}$ and tier intensity $hi\_any \in \{0, 1\}$, obtained as linear combinations of coefficients in (**??**). For interpretation: G1 (affluent, predominantly single-family), G2 (central, dense/multifamily), G3 (more mixed, lower-rent).

Wald tests (geographic heterogeneity): joint test $\beta_{DG2} = \beta_{DG3} = 0$ gives $\chi^2(2) = 10.66$ ($p = 0.0049$) for Probability, 9.37 ($p = 0.0093$) for Projects, and 9.44 ($p = 0.0089$) for Units. For $hi\_any = 0$, a joint test of $\beta_{DG2} = \beta_{DG3} = 0$ yields $\chi^2(2) = 10.54$ ($p = 0.0052$) for Probability, with Group 2 differing from Group 1 at the 1% level across outcomes (e.g., Probability: $p = 0.0015$; Projects: $p = 0.0023$; Units: $p = 0.0079$), while differences involving Group 3 are smaller and less precisely estimated. For $hi\_any = 1$, equality across areas cannot be rejected (e.g., Probability: $p = 0.81$).

Wald tests (high-intensity tiers): a joint test of $DH = 0$, $DH{+}DHG2 = 0$, and $DH{+}DHG3 = 0$ yields $p = 0.987$ for Probability, $p = 0.980$ for Projects, and $p = 0.971$ for Units, indicating no common incremental effect of high-intensity tiers once group differences are controlled for.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Pre-treatment diagnostics for kink–DiD (± 0.5 km, donut 0 m)

	Probability of any AH	Count of AH projects	Number of AH units
Panel A. RD pre (level at boundary, S)			
Estimate	0.003	0.003	-0.010
	(0.005)	(0.006)	(0.031)
Panel B. Pre-period slope diagnostics			
Pre slope diff (trend$\times S$)	-0.002	-0.002	-0.029
	(0.003)	(0.003)	(0.021)
Pre slope right (trend$\times R1$)	0.010	0.011	0.097
	(0.009)	(0.011)	(0.092)
Pre slope left (trend$\times L1$)	-0.003	0.000	-0.005
	(0.012)	(0.012)	(0.077)

Notes: Coefficients from pre-treatment regressions within a ± 0.5 km window around the TOC boundary; block-group and year fixed effects; standard errors in parentheses clustered by block group. "RD pre (level) – S" is the pre-policy discontinuity at the boundary. "Pre slope diff – trend$\times S$" captures differential linear time trends between treated and control sides. "Pre slope right/left – trend$\times R1$/trend$\times L1$" measure how the spatial gradient in distance evolves over time on each side. These objects arise from separate specifications and are not mechanically linked (e.g., "Pre slope diff" is not the difference between the right and left slopes). None of the coefficients is statistically significant at conventional levels, indicating no clear pre-policy breaks or differential trends at the boundary.

Table A.8: Kink–DiD by income × high-tier (manual linear combinations) — Probability of any AH

Cell	Jump	Right slope	Left slope	Kink (Right−Left)
AL	−0.002	0.031	−0.016	0.047
	(0.006)	(0.028)	(0.018)	(0.033)
BL	0.009	0.018	0.012	0.006
	(0.006)	(0.039)	(0.017)	(0.043)
AH	−0.016	0.105	−0.016	0.121
	(0.018)	(0.089)	(0.018)	(0.091)
BH	−0.016*	0.168***	0.012	0.155***
	(0.008)	(0.057)	(0.017)	(0.060)

Notes: Linear combinations of coefficients from the interacted version of (4.4). SEs by delta method with block-group clustering. Stars for the Kink column reflect the reported p-values (e.g., BH kink p=0.009). Endpoints implied by these fits at ±0.5 km match the plotted lines: for example, $y_{right}(BH) \approx -0.016 + 0.168 \times 0.5 = 0.068$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.9: Nearby price effects by project type and incentive intensity

	Project type		**Intensity (hi_any)**	
	DB	**TOC**	**Low tier**	**High tier**
Near×Post	−0.062***	−0.068***	−0.059***	−0.070***
	(0.021)	(0.017)	(0.021)	(0.022)
Diff (right − left)	−0.006		−0.011	
	(0.028)		(0.030)	

Notes: Same DiD on $\log P_{it}$ as in Table 7. "Project type" splits the nearby affordable project into Density Bonus (DB) vs. TOC; post-filter counts: DB n=75,542, TOC n=38,069. "Intensity" splits by high-tier eligibility (FINALTIER $\in \{3, 4\}$). Unit and year FE; controls include size, bedrooms, baths, construction type and quality. SEs two-way clustered by block and year. Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.10: No differential price effects by program type or tier (DiD near × post)

	ATT (log price)	SE
Panel A. TOC vs. DB		
DB (TOC=0)	−0.062	(0.021)
TOC (TOC=1)	−0.068	(0.017)
Diff (TOC − DB)	−0.006	(0.028)
Panel B. Density tier (hi_any: 1 if Tier 3–4)		
Low tier (hi_any=0)	−0.059	(0.021)
High tier (hi_any=1)	−0.070	(0.022)
Diff (High − Low)	−0.011	(0.030)

Notes: Parcel-level DiD on log sale price with unit and year fixed effects and rich housing controls (size, bedrooms, baths, construction type, quality). Two-way clustered SEs (block and year). Differences are statistically indistinguishable from zero (TOC–DB $p = 0.841$; High–Low $p = 0.706$).

Table A.11: Cross-sectional association between affordable/mixed projects and average rents

	Dependent variable: log(average rent)		
	(All)	(Below median)	(Above median)
Affordable / Mixed project	-0.058^{***}	−0.027	-0.132^{***}
	(0.022)	(0.025)	(0.033)
Observations	18,113	9,359	8,606
R^2	0.941	0.949	0.935
Panel B: By functional area (SKATER)			
	(1)	(2)	(3)
Affordable / Mixed project	-0.097^{***}	−0.002	0.181
	(0.024)	(0.043)	(0.226)
Observations	11,224	5,475	1,266
R^2	0.939	0.952	0.937

Notes: CoStar project-level data. Each column includes controls for unit mix (share of 0–3+ bedrooms), total units, average unit size, vacancy, project class and rating, style/segment, rent type and subsidy type; block (CTCB10) fixed effects. Standard errors clustered at the block level. * $p<.10$, ** $p<.05$, *** $p<.01$.

Table A.12: Predicted decay in average project rent from mixing affordable units, by rent bin

Average rent ()	**Pct. affordable within project**			
	20%	**15%**	**10%**	**5%**
2,000	1.6%	1.2%	0.8%	0.4%
2,500	5.3%	4.0%	2.6%	1.3%
3,000	7.7%	5.8%	3.9%	1.9%
3,500	9.5%	7.1%	4.7%	2.4%
4,000	10.8%	8.1%	5.4%	2.7%
4,500	11.8%	8.9%	5.9%	3.0%
5,000	12.6%	9.5%	6.3%	3.2%
Average	**8.5%**	**6.4%**	**4.2%**	**2.1%**

Notes: Entries report $\Delta \ln \bar{r} \approx -s(1-\rho)$ using observed affordable shares $s \in \{0.05, 0.10, 0.15, 0.20\}$ and a rent ratio $\rho = r^A/r^M$ calibrated from HUD schedules for Los Angeles; see Figure **??** for the empirical distribution of s by rent bin. Values are percent changes relative to the market-only average rent.

Table A.13: Construction quality and unit costs by neighborhood rent quartile

	Quality (normalized)	*Unit cost ($)*
Quartile 2	0.225***	185.119***
	(0.051)	(71.302)
Quartile 3	0.518***	420.549***
	(0.047)	(98.422)
Quartile 4	0.732***	945.279***
	(0.055)	(116.497)
Observations	301,442	301,442
R^2	0.106	0.038

Notes: Assessor (SBF) microdata for residential properties within 500m of TOC boundaries, pre-2017 builds; excludes extreme qualities and very small units. Regressions include controls for year built and construction type; specifications with year fixed effects yield similar patterns. Standard errors clustered by census tract. * p<.10, ** p<.05, *** p<.01.

Table A.14: PSM–DiD (±500 m): building permits — conteos y unidades

	AL	BL	AH	BH
A. Conteos de permisos				
n_addition (ATT)	-0.154^{***}	0.061	-0.052	0.219^{***}
	(0.048)	(0.052)	(0.111)	(0.077)
n_alter (ATT)	0.024	0.295	1.181^{*}	0.454
	(*n/r*)	(*n/r*)	(*n/r*)	(*n/r*)
n_demo (ATT)	-0.033	0.086^{**}	-0.019	0.164^{**}
	(*n/r*)	(*n/r*)	(*n/r*)	(*n/r*)
n_new (ATT)	-0.003	0.001	-0.008	0.112
	(0.052)	(0.052)	(0.100)	(0.111)
B. Unidades por tipo				
u_12fam (ATT)	0.070	-0.035	0.080	0.030
	(0.057)	(0.054)	(0.098)	(0.124)
u_apt (ATT)	-0.091	0.445^{*}	2.154	0.940^{**}
	(0.380)	(0.240)	(3.808)	(0.465)
u_comm (ATT)	-0.252	0.067	0.334	-0.141
	(0.527)	(0.131)	(0.402)	(0.090)

Notes: Block-group×year panel within 0–500 m of the TOC boundary; unit and year fixed effects; standard errors clustered by block group. Estimates are obtained on the propensity-score matched sample and reported separately for four rent×tier cells: AL, BL, AH, BH. "A"/"B" denote above- and below-median pre-policy rents, while "L"/"H" denote low- and high-intensity TOC eligibility (*hi_any*). Coefficients are post-treatment average treatment effects on the treated (ATTs) by cell, constructed as linear combinations of the interaction coefficients in the triple-differences specification..

Table A.15: PSM–DiD (±4,000 m): efectos por AL/BL/AH/BH en resultados principales

	AL	BL	AH	BH
Panel A. Probability of any AH				
ATT	0.006	0.015**	−0.003	0.036**
	(0.005)	(0.006)	(0.023)	(0.016)
Panel B. Count of AH projects				
ATT	0.007	0.017**	0.006	0.037**
	(0.006)	(0.007)	(0.026)	(0.017)
Panel C. Number of AH units				
ATT	0.030	0.231***	−0.018	0.604*
	(0.040)	(0.078)	(0.259)	(0.333)

Notes: Block-group×year panel within 0–500 m of the TOC boundary; unit and year fixed effects; standard errors clustered by block group. Estimates are obtained on the propensity-score matched sample and reported separately for four rent×tier cells: AL, BL, AH, BH. "A"/"B" denote above- and below-median pre-policy rents, while "L"/"H" denote low- and high-intensity TOC eligibility (*hi_any*). Coefficients are post-treatment average treatment effects on the treated (ATTs) by cell, constructed as linear combinations of the interaction coefficients in the triple-differences specification..

Table A.16: PSM–DiD (±4,000 m): building permits — conteos y unidades

	AL	BL	AH	BH
A. Conteos de permisos				
n_addition (ATT)	-0.177^{***}	0.064	-0.053	0.270^{***}
	(0.052)	(0.059)	(0.118)	(0.086)
n_alter (ATT)	0.022	0.437^{**}	0.616	-0.017
	(0.214)	(0.198)	(0.545)	(0.575)
n_demo (ATT)	-0.028	0.107^{***}	-0.041	0.161^{**}
	(0.035)	(0.036)	(0.121)	(0.072)
n_new (ATT)	0.001	-0.001	-0.004	0.047
	(0.052)	(0.057)	(0.110)	(0.126)
B. Unidades por tipo				
u_12fam (ATT)	0.059	-0.044	0.100	-0.013
	(0.056)	(0.059)	(0.099)	(0.135)
u_apt (ATT)	0.104	0.459^{*}	2.120	0.491
	(0.322)	(0.244)	(3.930)	(0.307)
u_comm (ATT)	-0.284	0.130	-0.051	-0.147
	(0.506)	(0.155)	(0.153)	(0.094)

Notes: Block-group×year panel within 0–500 m of the TOC boundary; unit and year fixed effects; standard errors clustered by block group. Estimates are obtained on the propensity-score matched sample and reported separately for four rent×tier cells: AL, BL, AH, BH. "A"/"B" denote above- and below-median pre-policy rents, while "L"/"H" denote low- and high-intensity TOC eligibility (*hi_any*). Coefficients are post-treatment average treatment effects on the treated (ATTs) by cell, constructed as linear combinations of the interaction coefficients in the triple-differences specification..

Table A.17: Housing supply (growth rates), PSM–DiD within ± 4 km

	Above median	Below median
Panel A. $\Delta \ln$ total units		
ATT	-0.000	-0.000
	(0.001)	(0.001)
N	42,504	42,504
R^2 (within)	0.091	0.091
Panel B. $\Delta \ln$ SF units		
ATT	-0.002^{*}	0.000
	(0.001)	(0.001)
N	42,504	42,504
R^2 (within)	0.108	0.108
Panel C. $\Delta \ln$ MF units		
ATT	-0.001	0.000
	(0.002)	(0.002)
N	42,504	42,504
R^2 (within)	0.086	0.086

Notes:* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.18: Bandwidth sweep ($\pm b$ meters), DiD at the boundary

Band b (m)	**prob_aff (Any AH project)** Above ATT	Above SE	Below ATT	Below SE	**count (AH projects)** Above ATT	Above SE	Below ATT	Below SE
100	−0.002	0.0037	0.0027	0.0027	−0.001	0.0060	0.0028	0.0029
200	−0.000	0.0051	0.0057	0.0039	0.0027	0.0061	0.0073	0.0047
300	0.0016	0.0049	0.0104**	0.0045	0.0028	0.0058	0.0120**	0.0049
400	0.0031	0.0050	0.0142***	0.0044	0.0052	0.0058	0.0157***	0.0047
500	**0.0057**	**0.0048**	**0.0154*****	**0.0047**	**0.0077**	**0.0055**	**0.0173*****	**0.0056**
600	0.0114**	0.0047	0.0172***	0.0045	0.0137**	0.0055	0.0196***	0.0054
700	0.0142***	0.0047	0.0181***	0.0043	0.0188***	0.0058	0.0213***	0.0051
800	0.0138***	0.0046	0.0218***	0.0043	0.0174***	0.0057	0.0263***	0.0052
900	0.0146***	0.0047	0.0224***	0.0042	0.0176***	0.0059	0.0272***	0.0051
1000	0.0155***	0.0046	0.0235***	0.0042	0.0181***	0.0058	0.0278***	0.0051

Notes: ATT and SE in level units (probabilities on 0–1 scale). "Above/Below" refer to above/below-median rent areas. All models include block-group and year fixed effects. $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.19: Donut DiD at ± 500 m (excluding inner ring)

Donut (m)	prob_aff ATT	prob_aff SE	count ATT	count SE	TOC+DB units ATT	TOC+DB units SE
0	0.0195***	0.0033	0.0230***	0.0042	0.262***	0.056
25	0.0209***	0.0036	0.0248***	0.0046	0.294***	0.060
50	0.0220***	0.0039	0.0261***	0.0049	0.311***	0.063
100	0.0244***	0.0043	0.0289***	0.0054	0.332***	0.070
150	0.0264***	0.0045	0.0306***	0.0057	0.361***	0.072

Notes: Symmetric ± 500 m window with an excluded inner ring of indicated radius.

Table A.20: ATT by distance ring (post vs. pre)

Ring (m)	prob_aff		count		TOC+DB units	
	ATT	SE	ATT	SE	ATT	SE
0–500	0.0103***	0.0036	0.0123***	0.0043	0.158***	0.050
500–1000	0.0400***	0.0072	0.0472***	0.0097	0.498***	0.142
1000–1500	0.0500***	0.0123	0.0650***	0.0145	0.817***	0.232
1500–2000	0.0412***	0.0132	0.0515***	0.0154	0.543***	0.174

Notes: Rings measured from the boundary into TOC. Estimates use Conley (spatial HAC) standard errors.

Table A.21: Placebo shifted boundaries (DiD at $\pm 500\,\text{m}$)

Shift (km)	ATT (prob_aff)	SE	ATT (count)	SE	n
−1	0.00345	0.00543	0.00305	0.00709	5,472
+1	0.00658	0.01324	0.01480	0.01602	6,003

Table A.22: Conley-corrected DiD at $\pm 500\,\text{m}$

Outcome	Est	SE	t	p	Heterogeneity (Below vs. Above)	
					Diff	Wald p
prob_aff	0.0103***	0.00224	4.60	0.0000	0.0097	1.3×10^{-15}
count	0.0123***	0.00279	4.40	0.0000	0.0060	< 0.01
TOC+DB units	0.158***	0.0369	4.28	0.0000	—	—

Notes: Last two columns report Below−Above and the Wald test p-value for *prob_aff* (full table available in the Appendix).

Table A.23: Pre-period diagnostics ($\pm 500\,\text{m}$, donut 0)

Outcome & term	Est	SE	z	p
prob_aff: RD pre level (S)	0.003	0.005	0.659	0.510
prob_aff: Pre slope diff (trend×S)	−0.001	0.002	−0.289	0.773
count: RD pre level (S)	0.001	0.006	0.224	0.823
count: Pre slope diff (trend×S)	0.000	0.003	−0.034	0.973
TOC+DB: RD pre level (S)	−0.016	0.033	−0.472	0.637
TOC+DB: Pre slope diff (trend×S)	−0.030	0.022	−1.356	0.175

Table A.24: Placebo timing: 2017 pseudo-reform (±500 m)

Outcome & term	Est	SE	z	p
prob_aff: Placebo jump ($\text{Post}^P \times$S)	−0.001	0.004	−0.176	0.860
prob_aff: Placebo kink ($\text{Post}^P \times$R1 − $\text{Post}^P \times$L1)	0.032	0.027	1.196	0.232
count: Placebo jump ($\text{Post}^P \times$S)	0.000	0.005	0.028	0.977
count: Placebo kink ($\text{Post}^P \times$R1 − $\text{Post}^P \times$L1)	0.039	0.032	1.214	0.225
TOC+DB: Placebo jump ($\text{Post}^P \times$S)	−0.090	0.065	−1.390	0.165
TOC+DB: Placebo kink ($\text{Post}^P \times$R1 − $\text{Post}^P \times$L1)	1.042	0.352	2.963	0.003

Table A.25: Kink–DiD excluding 2016–2017 (±500 m)

Outcome & term	Est	SE	z	p
prob_aff: Kink (Post×R1 − Post×L1)	0.074**	0.031	2.413	0.016
count: Kink (Post×R1 − Post×L1)	0.077**	0.037	2.070	0.038
TOC+DB: Kink (Post×R1 − Post×L1)	1.210***	0.408	2.964	0.003

Table A.26: DB-only DiD near the TOC boundary (±500 m)

	Prob(DB project)	DB permits	DB units
Treated × Post	−0.009***	−0.010***	−0.030
	(0.003)	(0.004)	(0.040)
Observations	18,216	18,216	18,216
R^2 (within)	0.176	0.177	0.121
Pre-treatment mean	0.013	0.014	0.058

Notes: Block–group and year fixed effects; standard errors clustered at the block–group level. Outcomes set non-DB entries to zero. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.27: Aggregate checks at the block level (AL/BL/AH/BH)

	AL	BL	AH	BH
Panel A. Probability of any AH				
ATT	0.0006**	0.0021***	0.0015**	0.0050***
	(0.0003)	(0.0004)	(0.0007)	(0.0006)
Panel B. Count of AH projects				
ATT	0.0007**	0.0024***	0.0017**	0.0056***
	(0.0003)	(0.0004)	(0.0007)	(0.0007)
Panel C. Number of AH units				
ATT	0.0046**	0.0263***	0.0165*	0.0856***
	(0.0018)	(0.0060)	(0.0084)	(0.0175)

Notes: AL = above-median rents × low tiers (1–2); BL = below-median × low tiers; AH = above-median × high tiers (3–4); BH = below-median × high tiers. Coefficients with SE in parentheses; unit and year fixed effects; SEs clustered at the unit level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.28: Aggregate checks at the census-tract level (AL/BL/AH/BH)

	AL	BL	AH	BH
Panel A. Probability of any AH				
ATT	0.022**	0.033***	0.043	0.045
	(0.009)	(0.009)	(0.072)	(0.040)
Panel B. Count of AH projects				
ATT	0.035***	0.036**	0.046	0.048
	(0.011)	(0.013)	(0.106)	(0.040)
Panel C. Number of AH units				
ATT	0.121	0.306***	1.108	0.611
	(0.080)	(0.115)	(0.986)	(0.777)

Notes: Definiciones de celdas como en la tabla de blocks. Coeficientes (SE) con FE de tracto censal y año; errores agrupados al nivel correspondiente. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figures

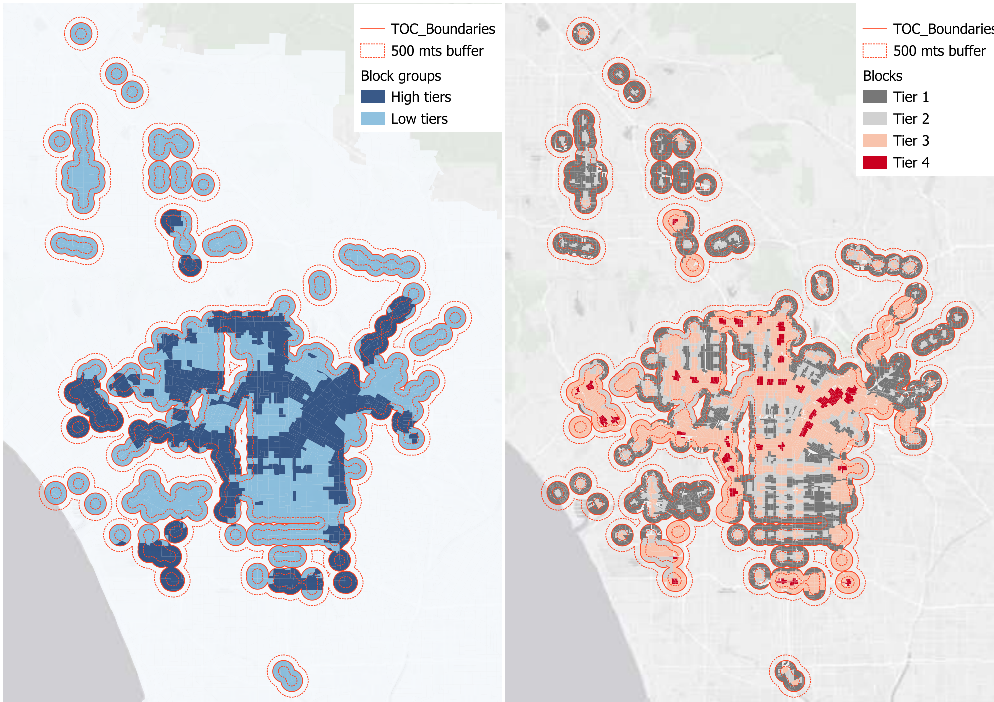

Figure A.1: Tier definitions.

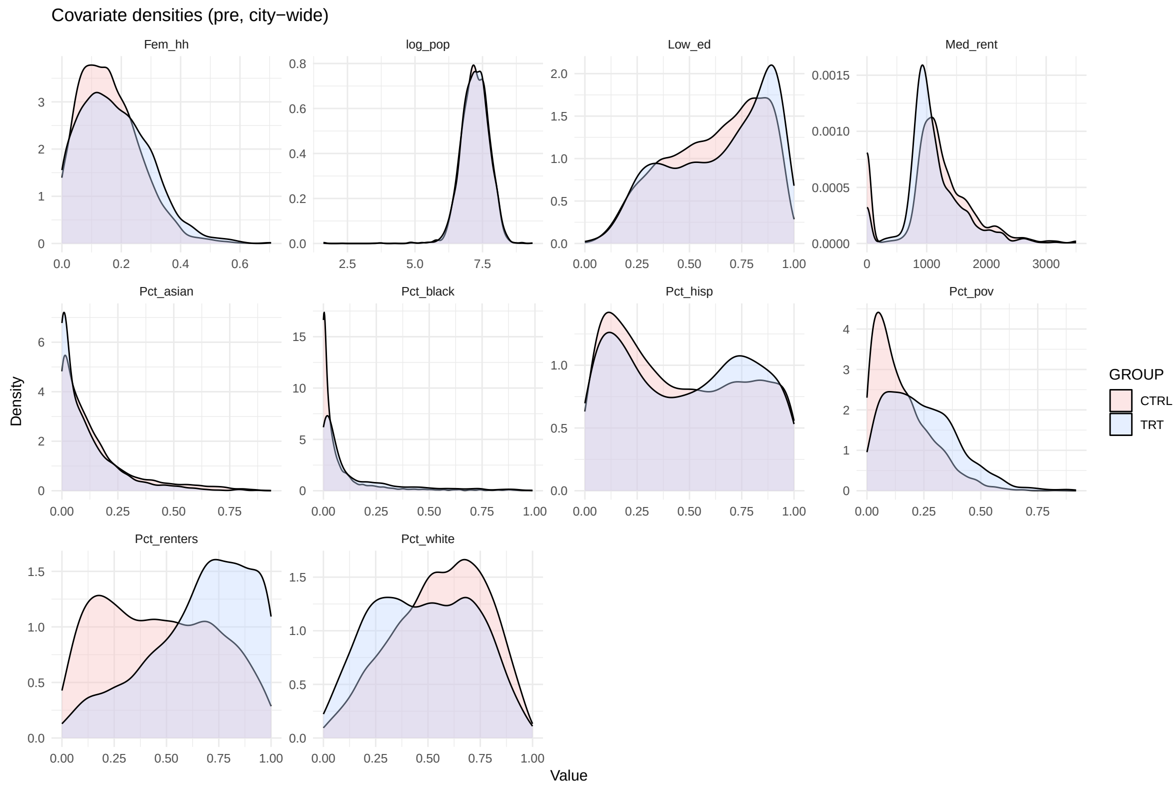

Figure A.2: Density of covariates between TRT and CTRL groups. Bandwith = 500 meters.

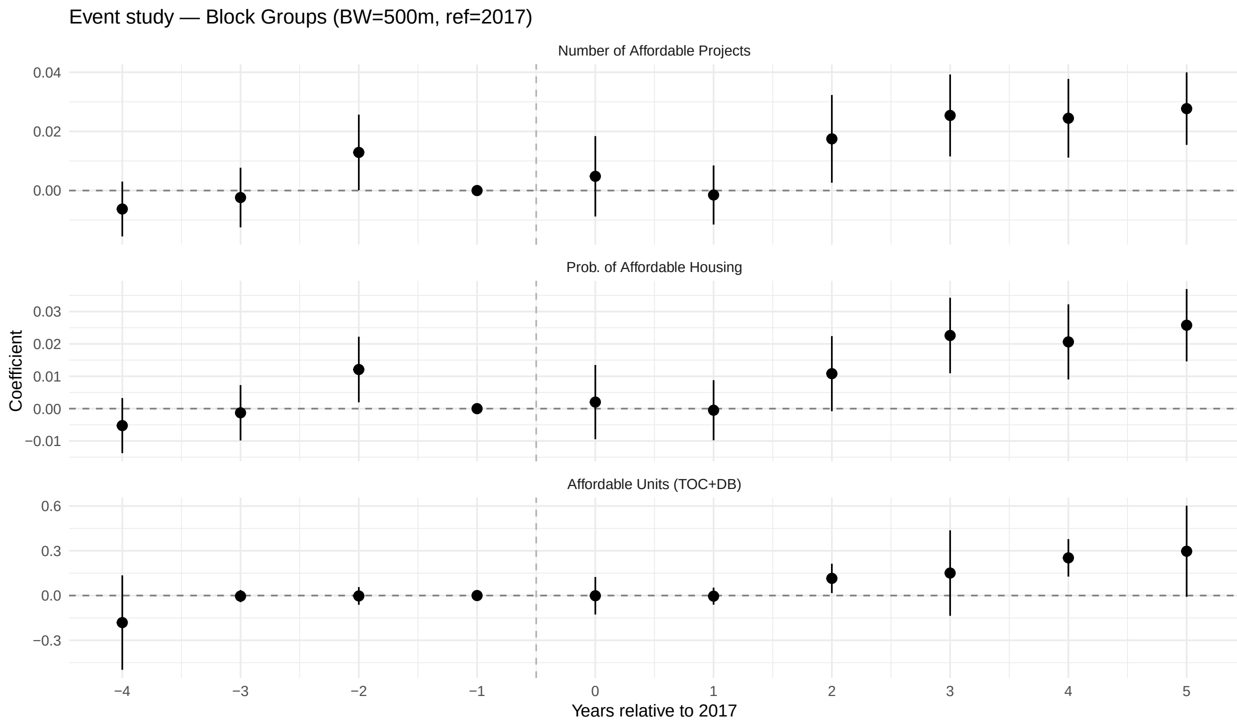

Figure A.3: Event study of permitting outcomes within 0–500 m

Notes: Coefficients $\hat{\beta}_k$ from a dynamic DiD with block–group and year fixed effects; interactions between treatment and relative–year dummies; baseline $k = -1$ omitted; 95% confidence intervals; SEs clustered by block group. *Pre–trend checks*: (i) Linear pre–period slope–difference (2014–2017)—Probability $p = 0.717$; Projects $p = 0.421$; Units $p = 0.852$. (ii) Placebo DiD with false Post=2016—Probability $p = 0.935$; Projects $p = 0.639$; Units $p = 0.665$.

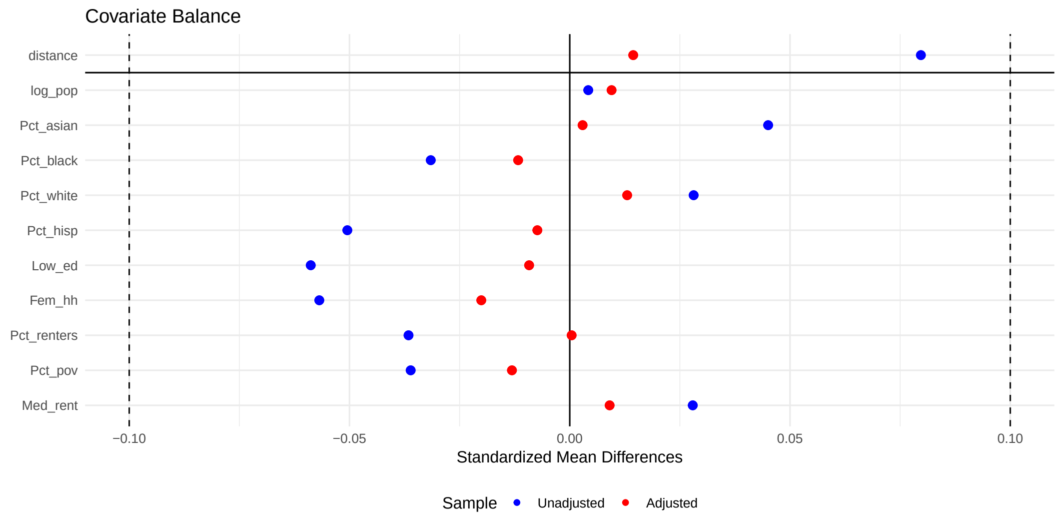

Figure A.4: PSM balance within ±500 m: standardized mean differences

Notes: Each point is a covariate's standardized mean difference (treated vs. matched control) in the pre period. Horizontal bands at 0.10 and 0.05. The matched sample attains SMD<0.10 for all displayed covariates.

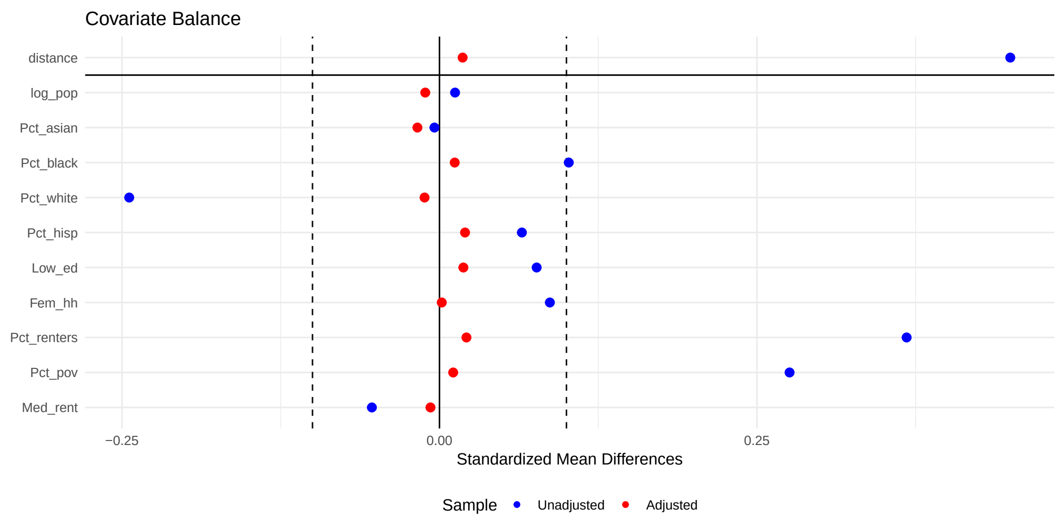

Figure A.5: PSM balance within ±4 km: standardized mean differences after matching

Notes: Love plot of absolute standardized mean differences (SMD). Dashed line at 0.1. The matched treated/control groups are constructed within the ±4 km window around the TOC boundary using pre-policy covariates and pre-trends.

K-means algorithm

I use a K-means clustering algorithm to define clusters for the heterogeneity analysis using the demographic and socioeconomic variables displayed in the determinants of affordable housing described in the data Section. I select the number of clusters, k, based on minimizing the total within-cluster variance. Each observation is assigned to the cluster whose centroid (mean) is closest to it, as measured by the Euclidean distance.

To determine the optimal number of clusters, k, I employed several standard cluster validation techniques:

1. **Elbow Method**: Evaluates the total within-cluster sum of squares (WSS) for different values of k and identifies the point at which the marginal gain from adding more clusters decreases.

2. **Silhouette Analysis**: Measures the average silhouette width, which reflects how similar each point is to its assigned cluster compared to other clusters.

3. **Calinski-Harabasz Index**: Quantifies the ratio of between-cluster dispersion to within-cluster dispersion.

4. **Gap Statistic**: Compares the WSS against that expected under a null reference distribution and identifies the optimal number of clusters based on the largest gap.

Each of these techniques was applied to the dataset with values of k ranging from 2 to 10 to find the most suitable number of clusters. Figure A.6 displays the results from each clustering validation method:

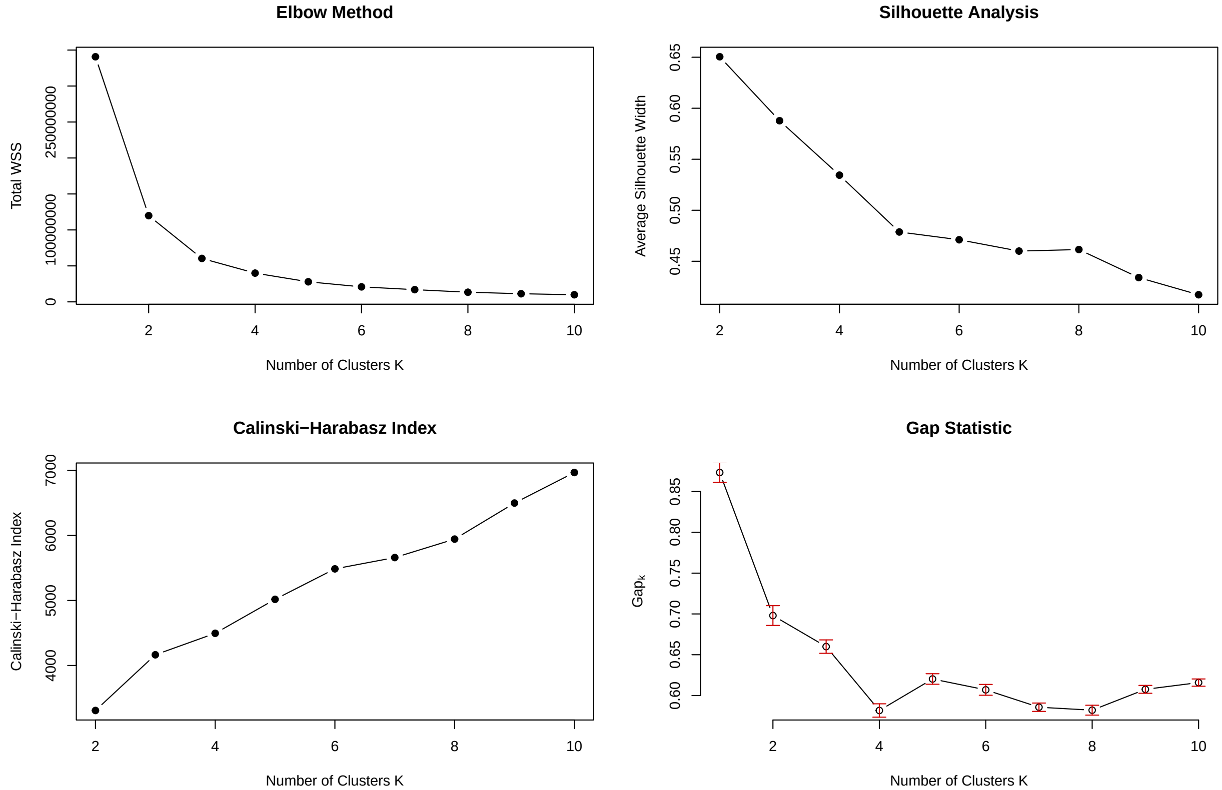

Figure A.6: Clustering validation methods

The elbow plot shows a tipping point at $k = 4$, where the total WSS starts to plateau, suggesting diminishing returns for more clusters. The Calinski-Harabasz Index increases steadily as k grows, with a more gradual increase beyond $k = 4$, suggesting that adding more clusters does not significantly improve the separation of clusters. The Gap Statistic peaks at $k = 4$, making this the strongest indication that $k = 4$ is the optimal number of clusters. Lastly, while the highest silhouette width occurs at $k = 2$, the drop from $k = 3$ to $k = 4$ is minimal, indicating that the clustering remains reasonably well-defined at $k = 4$.

To cross-validate these results, I also use a Principal Component Analysis (PCA). The results in Figure A.7 show a clear distinction between the four groups. Clusters 1 (red) and 2 (blue) the largest. However, Clusters 3 (green) and 4 (purple) exhibit distinguishable differences, too; although they might be more similar between them due to a smaller sub-sample.

Based on the validation techniques, I choose four clusters ($k = 4$). A map illustrating the spatial distribution of the four identified clusters is provided below :

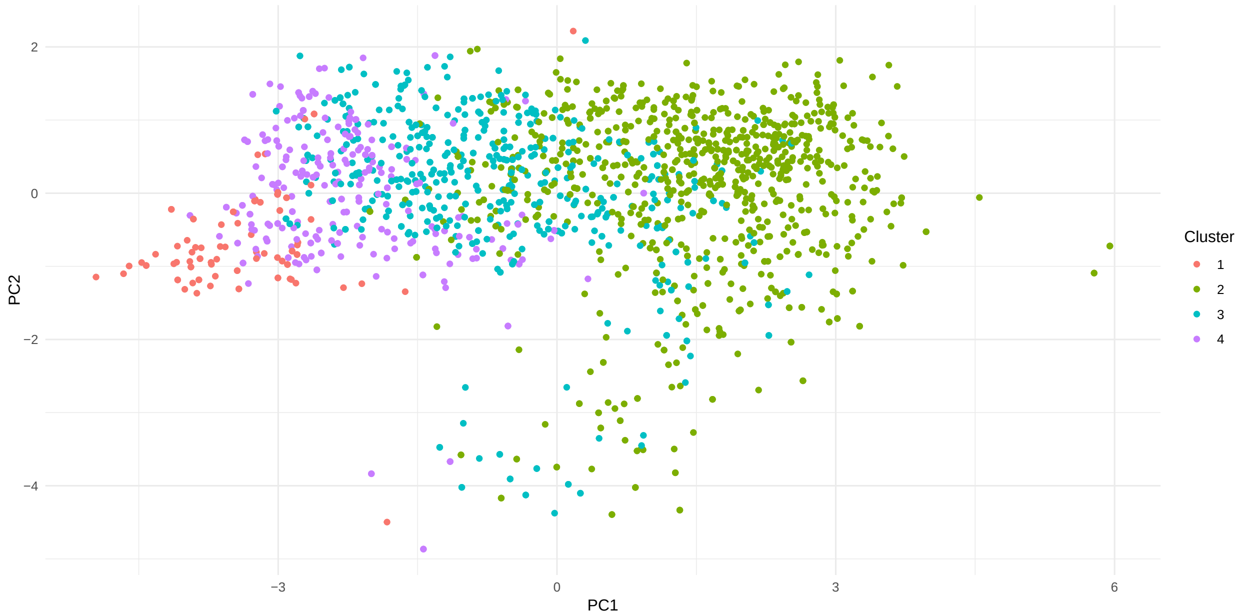

Figure A.7: Principal Component Analysis.

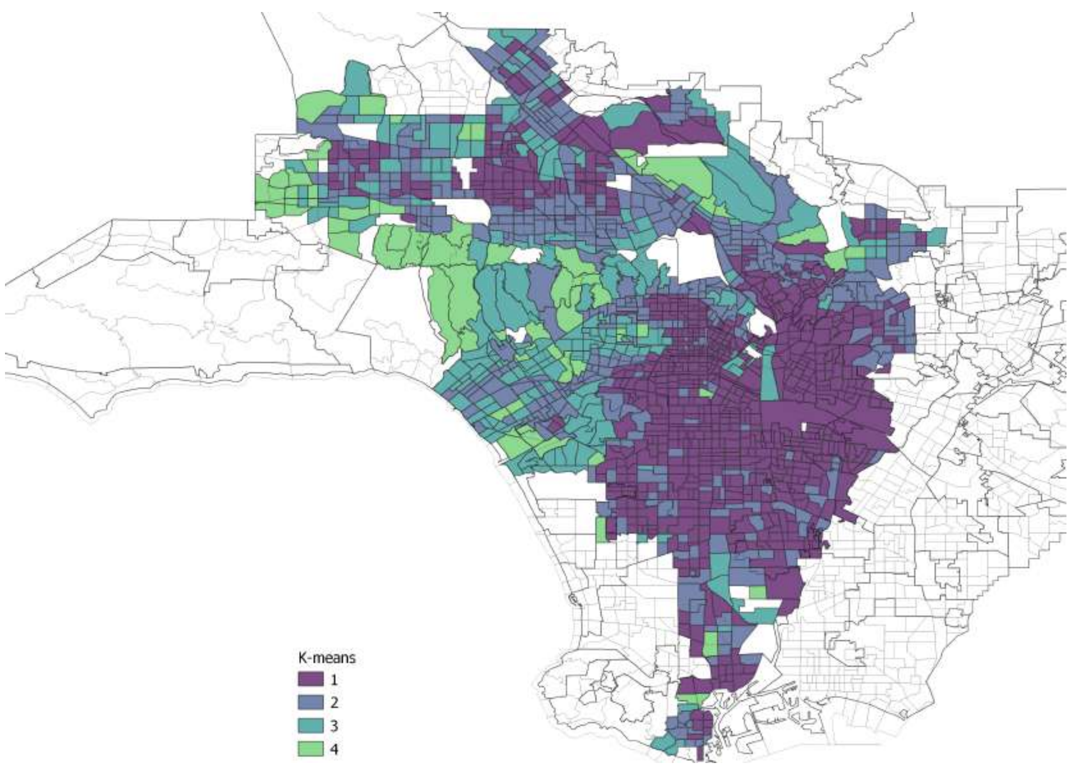

Figure A.8: K-means clustering by median rents

Spatially Constrained Clustering Algorithm

To complement the heterogeneity analysis, I added a second restriction to clustering: spatial contiguity, which accounts for autocorrelation and the influence on proximate observation blocks in this case. For this purpose, I developed a SKATER algorithm that creates clusters by building a minimum spanning tree (MST) that connects geographic areas with the least total dissimilarity. It progressively removes edges from the tree to form clusters while maintaining geographic contiguity. This is particularly advantageous over traditional k-means, which do not account for spatial variables.

For this analysis, geographic areas were clustered based on land use characteristics, density measures, and socio-demographic factors. These included percentages of single-family (SF) and multi-family (MF) housing, commercial and industrial land uses, as well as the socioeconomic variables used across the chapter.

The clustering analysis initially considered 4 clusters to match with the K-means algorithm. However, with more than three clusters, the number of observations in the third and fourth became too small to offer meaningful insights. For instance, introducing a fourth or fifth cluster led to groups with fewer than 50 units, reducing the reliability of the findings. Therefore, I opted for a three-cluster solution, balancing interpretability and statistical robustness.

I used Moran's I statistics to evaluate the clustering solution's performance. Moran's I statistic results indicate that the clusters are geographically consistent, with significant spatial autocorrelation. Hence, the SKATER algorithm successfully captures spatial patterns.

Metric	**Value**
Moran's I Statistic	0.8715
Expectation	-0.00057
Variance	0.00028
Z-Score	51.91
P-Value	< 0.0001

Table A.29: Moran's I Test Results

Table A.30 presents the characteristics of each cluster. Cluster 1 is dominated by single-family housing and low-density areas, with a predominantly White population and higher median rents, suggesting an affluent suburban profile. Cluster 2 shows higher multi-family housing shares and a higher proportion of Hispanic residents, indicating more urban, affordable areas. Cluster 3 stands out with a predominantly Black population and mixed density, reflecting minority neighborhoods with moderate affordability.

Table A.30: Mean values per clusters.

Cluster	1	2	3
Pct. SF	0.63	0.44	0.60
Pct. MF	0.17	0.31	0.28
Pct. Comm	0.06	0.09	0.06
Pct. Ind	0.03	0.05	0.02
Low Density	0.60	0.25	0.53
Mid Density	0.20	0.44	0.29
High Density	0.03	0.07	0.01
Black	4.35	16.80	59.30
Hispanic	41.50	66.50	29.90
White	63.20	43.30	26.10
Poverty	12.90	27.50	14.60
Female_head	13.60	25.30	22.40
Renters	58.60	68.70	55.50
Ed_Att	24.80	10.70	15.30
Median_Ren	1,314	958	1,160

As shown in the following map, the spatial clusters match with density and land uses, capturing Los Angeles's spatial configuration.

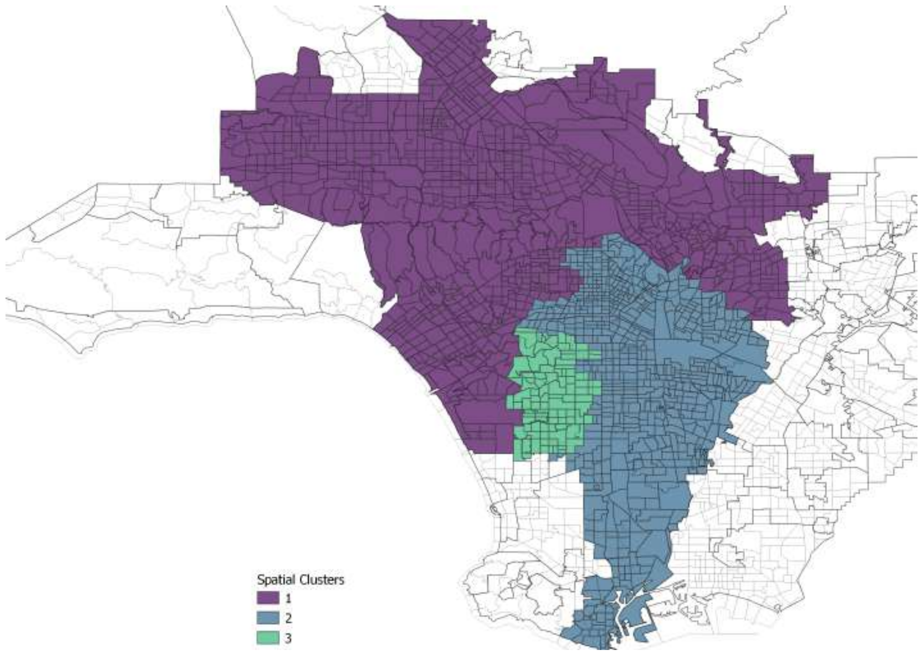

Figure A.9: Spatially constrained clustering

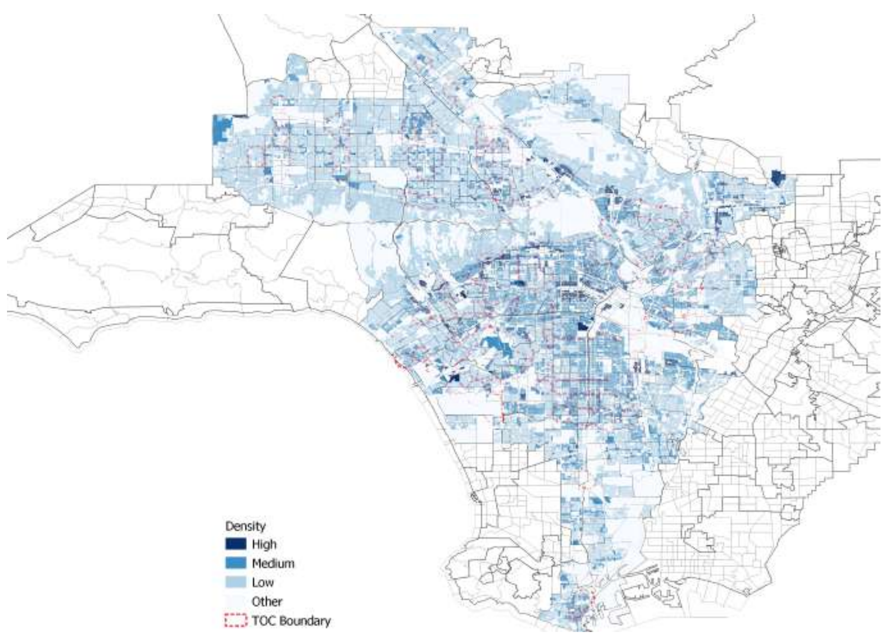

Figure A.10: Residential density per lot

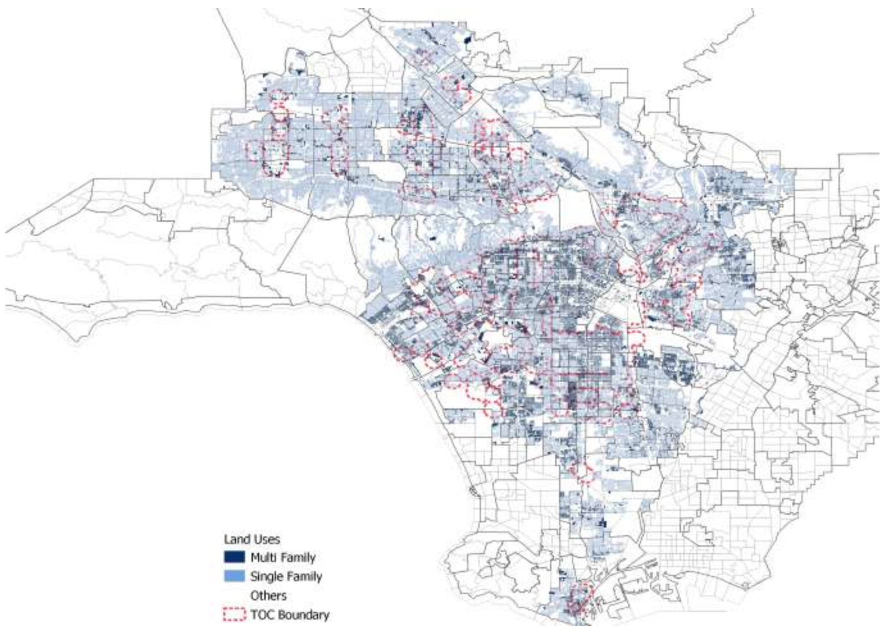

Figure A.11: Residential zoning per lot