Municipal Zoning and The Geography of Mobile Homes*

Colin Williams†

July 31, 2026
Click here for latest version

Abstract

Local land use regulations shape not only how much housing is built but also its composition. At municipal boundaries, the mobile home share of detached homes more than triples from roughly 2% to over 8%, suggesting that regulations within cities restrict mobile homes. Consistent with this interpretation, quality-adjusted prices of mobile homes decline by roughly $7,700 at the same boundaries even as site-built home prices evolve continuously, further evidence that municipal regulation burdens mobile homes relative to site-built housing. I develop a sufficient-statistics framework to assess these distortions in the composition of the housing stock. The framework leverages a CES demand model over differentiated housing varieties in which municipal zoning enters as a differential implicit tax on mobile homes. I map the two reduced-form statistics to two structural parameters: the implicit regulatory tax and the elasticity of substitution between mobile and site-built homes. The model implies a high elasticity of substitution between the two types conditional on size, vintage, and lot size. The high substitutability means that a modest regulatory tax reduces the share of mobile homes in cities by more than 60% relative to its counterfactual level, suggesting that municipal land use regulations, not household sorting, drive both the ruralization and small market share of mobile homes.

Keywords: mobile homes, zoning, land use regulation, housing supply, boundary discontinuity

*I am grateful to my advisors, Leora Friedberg, Lee Lockwood, and Kerem Coşar, for their guidance and unflagging support on this project. I also thank participants at the University of Virginia public economics workshop, the Institute for Humane Studies papers workshop at UCSD, Kyle Butts, and W. Ben McCartney for their helpful comments. I acknowledge data access from The Household and Urban Finance Lab at UVA. All errors are my own.

†University of Virginia. chv7bg@virginia.edu

1 Introduction

Housing affordability has become a pervasive concern: even in rural areas, growth in home prices and rents has exceeded income growth since 2000, driven largely by increasing construction costs (Baum-Snow and Duranton, 2025; Glaeser and Gyourko, 2025). Mobile homes cost roughly 30% less to build than comparable site-built homes, yet they comprise only about 10% of the US housing stock and are largely confined to rural and exurban areas (Herbert et al., 2023).[1] If factory-built housing offers substantial cost savings, why do so few households choose it, and why are those who do concentrated outside cities?

One explanation for the concentration of manufactured homes in rural areas emphasizes demand.[2] Manufactured homes (MH) may be lower quality along many dimensions, including more rapid depreciation, greater exposure to extreme weather, and social stigma, leading wealthier suburban households to avoid them even when they appear cheaper on a per-square-foot basis. Supply-side factors may also matter. Cities restrict the supply of mobile homes through zoning regulations and building codes, eroding their cost advantage and pushing mobile homes into unincorporated and less regulated areas.

In this paper, I document that municipal boundaries generate a spatial discontinuity in the share of mobile homes. Among single-family homes within 2 km of a municipal boundary, the MH share inside cities is roughly 2.1%, while outside the boundary it jumps to 8.3%. The same pattern holds when comparing homes near the same city and within the same school district, conditioning on each home's physical characteristics, and flexibly controlling for distance to the boundary line. Quality-adjusted mobile home prices fall by about $7,700 relative to site-built homes at these same boundaries (s.e. $2,500), the price counterpart of the quantity gap. Site-built prices themselves show no discontinuity, suggesting that sorting on demand for municipal amenities is limited and unlikely to explain the sharp change in market shares. I support this interpretation with a placebo test showing that the mobile home share does not change at Census Designated Place boundaries, population centers with no formal governments or zoning authority.

I then develop a sufficient-statistics framework to assess distortions in the composition of the housing stock. The framework uses a CES demand model over differentiated housing varieties to relate discontinuities in mobile home shares and relative prices at municipal boundaries to two economic fundamentals: the elasticity of substitution between mobile and site-built homes, and the implicit differential regulatory tax on mobile homes imposed by cities. I use a double-difference structure, comparing MH to site-built homes and incorporated to unincorporated areas, to isolate regulatory costs from common land and construction costs that affect both housing types equally. Using a boundary RDD in relative prices, I first estimate the implicit tax on mobile homes; then I recover the elasticity of substitution from the share discontinuity using the model's equilibrium condition. The resulting elasticity is long-run and conditional on living area, vintage, lot size, and location. Identification requires that *relative* housing demand — households' demand for mobile homes relative to otherwise-similar site-built homes — is continuous at boundaries. Amenities that change at the boundary and shift demand for both housing types equally do not threaten the design;

[1] The cost of structures accounts for an average of 60% of US house values, so rising construction costs matter greatly everywhere except a handful of the most expensive cities (Davis et al., 2021).

[2] "Manufactured homes" are the preferred industry term and refer specifically to mobile homes that satisfy the 1976 Department of Housing and Urban Development (HUD) Code. As the vast majority of mobile homes produced since 1976 satisfy the code, I use the terms "manufactured" and "mobile" interchangeably. About 40% of mobile homes are located in trailer parks, where residents may rent the lot, the unit, or both; I focus on the 60% sited on private land (Genz, 2001).

only a discontinuity in the relative valuation of the two types would. I support this assumption with a variety of evidence.

Empirically, municipal regulation, and not household sorting, accounts for most of the "ruralization" of mobile homes. The RDD estimates imply that municipal zoning regulations reduce the MH share inside cities from roughly 5.8% to just over 2.1% — eliminating about 63% of the mobile homes that would otherwise be sited there. The price wedge behind this large quantity response is modest: a differential implicit tax of 5.9% of the mobile home price, several times smaller than the 41% cost advantage of factory-built construction.[3]

I estimate an imprecise long-run elasticity of substitution between mobile and site-built homes of roughly 18, conditional on living area, vintage, lot size, and side of the boundary. Relative quantities shift substantially at municipal boundaries; relative prices by much less, suggesting that households are close to indifferent between the two housing types in their own consumption choices. A modest tax can therefore generate a large quantity response: a 5.9% wedge is consistent with the observed 63% reduction in the mobile home share because substitution is easy.[4] The high substitution elasticity, combined with the small baseline share of MH, limits the aggregate welfare cost of municipal regulations, which I estimate to be 0.13% of housing expenditures. This loss may understate the true welfare consequences because it does not account for distortions to the overall quantity of housing, resorting across locations, blanket bans on mobile homes, or potential economies of scale in the production of mobile homes (Gann, 1996).

The differential tax on MH is largest where the institutional logic of local control predicts. Exploiting boundary-level variation in the estimated price effects, I find that the implicit tax is larger where the municipality is geographically smaller — so that each potential MH parcel represents a larger share of the housing stock that incumbent residents see — and in states that grant cities broad autonomy over land use through Home Rule rather than Dillon's Rule charters. This heterogeneity ties the average boundary effect to the scope of local control: more compact and more autonomous municipalities impose the greatest relative restrictions on MH.

My results contribute to a large literature on the role of local regulations in shaping housing markets that is reviewed in Gyourko and Molloy (2015). A smaller literature examines the specific consequences for mobile homes, finding that zoning regulations are a significant barrier (Dawkins, 2011, 2025). My results highlight the consequences of land use regulations for housing characteristics, even in relatively small cities, and emphasize the importance of studying MH at the outskirts of cities, where they are most likely to be a realistic substitute for detached site-built homes.

I also contribute to a nascent methodological literature that exploits variation at regulatory boundaries to identify household preferences for housing attributes (Gyourko and McCulloch, 2024; Kulka et al., 2026; Anagol et al., 2021; Rollet and Weiwu, 2025). I extend this literature by interpreting estimates from a difference-in-discontinuities design in a sufficient statistics framework, clarifying the interpretation and economic implications of the reduced-form estimates. My empirical strategy extends the traditional difference-in-discontinuity design, which generally relies on temporal variation, to a static setting with variation across product types (Grembi et al., 2016; Butts,

[3]I focus on the average differential effect of municipal regulation on mobile homes relative to the adjacent unincorporated areas. Municipalities have numerous tools to restrict mobile homes, either directly through bans or zoning regulations or indirectly through fees, architectural standards, fire codes, and so on (Dawkins, 2011). The cumulative effect of such regulations may be large even if any single regulation has little effect. Because the estimation sample includes only boundaries with mobile homes observed on both sides, the estimates capture this intensive margin of regulation and exclude cities that ban mobile homes outright.

[4]The prevalence of anti-MH regulation itself suggests that households are far less indifferent about the housing types chosen by their *neighbors*.

2023). In so doing, I contribute to a fast-growing literature on housing demand that emphasizes the gains from housing variety (Calder-Wang, 2021; Ma, 2024; Zhang, 2022).

The welfare framework applies to other settings where local regulation distorts the composition of housing varieties. Examples of such distortions include multi-family apartments, which may be discouraged through affordable unit requirements; "family-friendly" apartment layouts, which may be discouraged through minimum parking-per-bedroom regulations; and small single-family homes, which may be disproportionately affected by minimum lot size regulations (Mei, 2022). My framework extends the "implicit zoning tax," introduced by Glaeser and Gyourko (2002), to settings where land use regulations vary across housing types.

Finally, I show that the geographic scope of land-use authority shapes housing markets. A long literature argues that highly local control over land use produces more restrictive zoning overall, since the costs of new development are borne by a larger share of the incumbent electorate while the diffuse benefits flow to relatively more non-resident (and non-voting) households (Hankinson, 2018). Recent evidence shows that transitions from "at-large" to ward-based elections reduce new housing production (Mast, 2024; Hankinson and Magazinnik, 2023). A direct extension of this logic suggests that incumbent residents may also prefer and enact relatively greater compositional restrictions, such as architectural standards or outright bans, if they have type-specific objections to nearby development. The boundary-level heterogeneity documented above — larger differential taxes in smaller and Home Rule municipalities — is consistent with this collective-action interpretation and points back to the same institutional channel, the scope of local control, that explains why the average effect arises at the boundary in the first place.

This paper focuses on the costs of anti-MH zoning. Such rules may also generate benefits if manufactured homes create negative externalities, either directly through their appearance or indirectly through resident behavior (Munneke and Slawson, 1999). Such externalities would be consistent with how single-family homeowners view another common type of low-income housing: dense, multi-family apartments (Chiumenti et al., 2022). Separately, residents of mobile homes may generate fiscal externalities if they pay lower property taxes than site-built homeowners but impose similar demands on public services and schools, a concern that is especially relevant given that the property tax differentially burdens more durable forms of capital and MH tend to depreciate relatively more quickly than other housing (Hamilton, 1975; Arnott and Young, 1979; Zhou, 2013). I leave the assessment of any such benefits to future work.

2 Institutions and Data

This section provides the institutional background for the empirical design and then describes the data. Section 2.1 explains how cities regulate mobile homes, why regulatory stringency changes discretely at municipal boundaries, and where theory predicts the city–county regulatory gap to be largest. Section 2.2 describes the parcel-level data on housing characteristics, sale prices, and neighborhood demographics, along with the construction of the analysis sample.

2.1 Mobile Homes and Municipal Zoning

Manufactured homes are assembled in factories and transported to their final location, where they are installed on a permanent foundation. About 40% of mobile homes are located in trailer parks, where residents may rent the lot, the unit, or both. I focus on the 60% of mobile homes that are located on private land (Genz, 2001).[5] MH on private land share the advantages of land ownership with conventional single-family homes — potential price appreciation, privacy, and no risk of eviction — and therefore represent a realistic substitute, differing primarily in the structure itself.

The federal Department of Housing and Urban Development (HUD) code specifies minimum standards for manufactured homes and preempts local building codes. However, local governments retain broad authority to regulate the location, density, and architectural features of mobile homes through zoning ordinances. They may also impose general land use regulations, such as minimum lot sizes or fire safety requirements, which may have disproportionate effects on mobile homes (Dawkins, 2011).

My empirical strategy leverages spatial variation in regulatory stringency at municipal boundaries, where the zoning authority of an incorporated city gives way to that of the surrounding county. Cities and counties differ in both their geographic scope and their legal authority over land use, so theory predicts that mobile home regulation will be systematically stricter just inside the municipal boundary — and that the regulatory gap will be widest where these institutional differences are most pronounced, a prediction I return to in Section 6.

The zoning ordinance of Lumberton, NC, a small city 90 minutes south of Raleigh, provides a typical example of this pattern. Appendix Section G reproduces excerpts from Lumberton's ordinance that impose architectural standards on mobile homes above the federal HUD code, including on the roof pitch and exterior materials, as well as limits on their placement in residential zones. The neighboring Robeson County zoning ordinance (which applies in the unincorporated land adjoining Lumberton) imposes fewer restrictions on mobile homes, allowing their placement in residential zones with minimal design requirements.[6]

Unincorporated areas generally impose less restrictive land use regulations than incorporated cities do. The rapid growth of the Sunbelt cities, for instance, is partly the result of large swaths of unincorporated land available for subdivision; in the Northeast, by contrast, all land was already managed by incorporated municipalities with the structure to enact tighter zoning rules (Glaeser and Gyourko, 2025). County authority over land use varies somewhat across states. In Texas, most counties are prohibited from enacting any zoning regulations at all; their authority is limited to subdivision regulations and building codes for site-built homes, not mobile homes (Reid Wilson, 2025). In other states, such as Virginia, Florida, and California, comprehensive county zoning is allowed, though enforcement focuses on specific complaints and may be limited by county resources (Maantay, 2001).

Two institutional features of city government shape how aggressively cities can exclude MH. The first is the geographic scope of the municipality itself. Cities are far smaller than the counties they border, both in area and population, so any individual development represents a larger share of the housing stock that incumbents see. This concentrates the perceived cost of MH on a smaller pool of voters and lowers the collective-action threshold for organized opposition, while the diffuse benefits — affordable housing for prospective residents — fall largely on non-voters (Hankinson, 2018; Mast, 2024). Counties, as larger political units, are plausibly more likely to internalize the

[5] The private-land share is even higher among recent placements: between 2007 and 2013, over 70% of new manufactured homes were placed on private property. "Cost & Size Comparisons: New Manufactured Homes and New Single-Family Site-Built Homes (2007 - 2014)." U.S. Census Bureau. Accessed May 12, 2026. https://www2.census.gov/programs-surveys/mhs/tables/time-series/sitebuiltvsmh.pdf.

[6] Robeson County Zoning Ordinance, Section 5.9. "Requirements for Mobile/Manufactured Homes." Available at https://www.rockcreeklandcompany.com/wp-content/uploads/2025/04/Zoning-Robeson-County-.pdf.

broader benefits of a varied housing stock and less likely to respond to a small number of highly vocal incumbents (Favilukis and Song, 2023). By the same logic, the smallest cities should impose the largest differential burden on MH.

The second is the state-level legal regime governing local authority. Under Dillon's Rule, local governments possess only those powers expressly granted by the state legislature, and their ordinances are construed narrowly; under Home Rule, cities possess broad residual authority over local matters, including land use, subject only to explicit state preemption (Richardson et al., 2003). The legal distinction has measurable correlates, with Home Rule cities being relatively more likely to annex land and to avoid large drops in tax revenue (Zhang et al., 2022; Zhang and Nguyen-Hoang, 2023). Separately, a number of states preempt local exclusion of MH outright or limit the design and placement restrictions cities may impose; Dawkins (2011) catalog the most protective of these regimes. The legal regime therefore caps what cities can do at the boundary even when their incumbents would prefer to do more, and the differential MH tax should be larger in Home Rule states and smaller where state MH protections bind.

Land use regulations are not the only factors that change discretely at municipal boundaries. Changes in other services, such as public schools, street lighting, road maintenance, utilities, and policing, as well as in property tax rates, may also cause a discontinuous change in relative demand for mobile homes at the boundary. I address these concerns in three ways. First, I restrict the analysis sample to boundaries where parcels on both sides are located in the same school district, avoiding the threat that demand-side sorting on school quality — by far the largest component of local public spending — rather than regulation, drives the change in MH shares. Second, I include flexible controls for distance to the boundary to capture local amenities — such as environmental quality and non-excludable public goods like public parks — that vary continuously at the boundary. Finally, I present evidence in Table 3 that site-built home prices do not change discontinuously at municipal boundaries, which rules out the bid-rent channel through which sorting on relative MH demand would show up in the data.

2.2 Data

I construct a sample of single-family homes across the US which contains detailed property characteristics, sale prices, and neighborhood demographics.

Property characteristics and values. I obtain housing characteristics and sale prices from CoreLogic, a vendor of real estate data compiled from county assessors' offices. CoreLogic provides a snapshot of housing data from 2023 that includes detailed characteristics: the type of structure (manufactured or site-built), lot size, living area, year built, address, and assessed value. CoreLogic also provides a database of property transaction histories that include sale price and date. I filter these data to the most recent arm's-length transactions, excluding sales with trivial prices ($<$ \$5,000), sales involving multiple or split parcels, and all cases where a parcel transacted multiple times on the same day. I deflate sales prices to real 2023 dollars using the FHFA House Price Index® for single-family homes based on the year of sale.

Demographic characteristics. The housing data do not report demographic characteristics for the occupants of individual homes, so I instead measure neighborhood demographics using granular data from the Census Bureau's Environmental Impact Frame (EIF). These data are constructed using administrative tax records and provide population estimates by race, age, and income at a fixed 0.01-degree grid cell level, or roughly 1 square kilometer.[7] I match each parcel in CoreLogic to the nearest EIF grid cell and construct demographic characteristics based on the 2023 population estimates.

Incorporated cities and boundary areas. I determine municipal and school zone boundaries using shapefiles from the US Census Bureau's TIGER/Line database. I focus on incorporated places, which include cities with formal governments that possess taxing and zoning authority. For each parcel, I assign its boundary area based on the nearest incorporated place, using boundaries from the 2011 Census TIGER/Line database to reduce concerns about endogenous boundary changes.

Each boundary area is defined as the interaction of the nearest city and the local school district, so all parcels within a boundary area are located in the same school district. This restriction limits the potential for sorting on unobserved amenities given that school quality is a key driver of household location choices (Black, 1999; Bayer et al., 2007).

Final sample. I study detached single-family housing, the segment in which manufactured and site-built homes compete most directly. I impose several sample restrictions before reporting descriptive statistics. First, I keep homes built since 1980, because modern, federally regulated MH emerged only at the end of the 1970s and earlier vintages have widely variable quality. Second, I drop parcels that report no lot size or living area. This step removes some manufactured homes located in trailer parks, which are inconsistently included in assessor data, typically with no associated lot size. Third, because the price RDD is specified in levels, I winsorize sale prices at the 1st and 99th percentiles within each housing type, which keeps a handful of extreme sales from driving the level price estimate without discarding observations. Finally, I keep parcels within 2 kilometers of a city boundary and restrict to boundary areas with at least 50 observations and at least one manufactured home on both sides of the boundary.[8] Appendix Section D documents how each step affects the sample and shows that the results are robust to alternative sample restrictions.

In the analysis I also condition on physical characteristics directly, using a rich product segmentation based on lot size, living area, and the construction decade. Figure 1 plots MH shares across the living area $\times$ lot size decile grid and shows that MH are most common among smaller homes on the very largest lots.

The sample does not condition on tenure, but detached housing is overwhelmingly owner-occupied: in the ACS tract data covering my sample areas, detached site-built and mobile homes together account for 95.1% of owner-occupied housing and 80% of the total housing stock; by contrast, buildings with 5 or more units account for 11% of all housing, a share comparable to mobile homes. I therefore interpret the elasticity estimate as a conditional substitution elasticity within the detached single-family market — the segment for which it is most relevant, and which is often, though not exclusively, owner-occupied — not as an all-housing elasticity that includes multifamily. Put differently, households priced out of MH by municipal regulation may sometimes substitute into renting an apartment unit, but that margin matters for the elasticity only if the relative attractiveness of apartment units changes discontinuously at the municipal boundary in a

[7]The EIF's fixed grid is much finer than the Census tract, and, because it is built from administrative records, it avoids both the sampling error of ACS estimates and changes in statistical boundaries over time (Voorheis et al., 2023).

[8]Because almost the entire area of New England is incorporated, I exclude this region from the analysis. New England accounts for only 0.65% of all US manufactured homes in the CoreLogic data.

way that also differentially shifts demand for MH relative to detached site-built homes.[9]

The final sample includes almost 6.3 million parcels across 3,627 boundary areas. Appendix Section D traces the sample through each restriction and reports its composition by state. The requirement of at least one manufactured home on both sides of the boundary is necessary to estimate boundary effects on relative prices; it implies that my results reflect intensive margin effects of zoning regulations on MH quantities and prices, rather than an extensive margin effect from a complete ban.

Table 1 reports summary statistics for the final sample, comparing incorporated and unincorporated areas. The sample includes approximately 4.5 million parcels in incorporated areas and 1.8 million parcels in unincorporated areas. Unincorporated areas have a higher share of manufactured homes (8.3% vs 2.1%) and lower average home prices ($410k vs $436k). Homes in unincorporated areas tend to be located on larger lots, with correspondingly lower population density. However, both areas have similar demographics and incomes, and the average home on each side is roughly 2,000 square feet and was built in the year 2000.

2.3 Motivating Evidence

Figure 2 plots the mobile home share around municipal boundaries. The relative quantity of MH jumps at the boundary line, with the share of manufactured homes increasing from about 2.1% inside the city to over 8.3% just outside. Shares are relatively stable away from the boundary, though they do increase slightly with distance from the city.

Figure 3 shows how this spatial pattern manifests in Lumberton, North Carolina, a small city in the southeastern part of the state. The city extends roughly along the major north-south highway, Interstate 95, and site-built housing comprises the majority of its housing stock. However, just outside the city boundary, mobile homes are much more prevalent, with large clusters to the north and east of the city.

These descriptive results suggest that cities differentially restrict the supply of mobile homes.

A discontinuity in the mobile-home share at the boundary is consistent with municipal regulation that differentially burdens mobile homes, but it could also reflect a discontinuous change in the relative *taste* for manufactured housing—if, for example, households sort across the boundary on unobserved preferences. Even if the gap is regulatory in origin, its size does not reveal the implicit tax: the same share discontinuity is consistent with a small tax when the two housing types are close substitutes, or a large tax when they are not. Separating the two requires a model.

3 Model

I model the market for owner-occupied detached single-family housing to quantify the role of municipal zoning in the geography of manufactured homes. The model features CES demand over differentiated housing types and perfect competition in production. It links the RDD estimates to

[9] The welfare calculation is correspondingly partial-equilibrium: it measures the compensating-variation loss from the differential implicit tax on MH within the market for detached housing. To first order, that object depends on the tax wedge and the counterfactual expenditure share on MH rather than on whether regulation also changes homeownership rates, so the results should be understood as the welfare cost of distorting the composition of detached housing, not the full welfare effect of municipal housing policy across all types and tenures.

two economic fundamentals: the implicit differential tax on manufactured housing (MH) imposed by cities, t, and the elasticity of substitution between MH and site-built housing, σ.

To keep the identification argument in focus, I suppress boundary-area and product-segment subscripts throughout this section; all objects are understood to hold within a fixed product segment $\times$ boundary-area cell. Product segments and boundary areas are reintroduced as fixed effects in the empirical strategy (Section 4.1).

Let $d \geq 0$ denote the distance from a parcel to the nearest city boundary, and let $\ell \in \{I, U\}$ index the regulatory regime on each side, with $\ell = I$ inside the incorporated area and $\ell = U$ in the adjacent unincorporated area. Any structural object f that varies across locations is then a function $f_\ell(d)$ of side and distance, defined for every $d \geq 0$ on each side. Define the boundary difference operator

$$\Delta_\ell f(d) \equiv f_I(d) - f_U(d).$$

Δ_ℓ computes the inside-minus-outside difference between two parcels equidistant from the boundary on opposite sides; taking $d \to 0$ isolates discontinuities at the boundary and delivers the empirical object estimated by the RDD. Level differences between the inside and outside of the city are permitted throughout; the identification argument turns on which structural objects are assumed continuous as $d \to 0$.

3.1 Demand

Consumers choose between manufactured ($h = M$) and site-built ($h = S$) housing. At each location (ℓ, d), aggregate demand takes the CES form

$$Q_\ell(d) = \left[\alpha_{M,\ell}(d)^{1/\sigma}\, q_{M,\ell}(d)^{\frac{\sigma-1}{\sigma}} + \alpha_{S,\ell}(d)^{1/\sigma}\, q_{S,\ell}(d)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}},$$

where $\sigma > 0$ is the elasticity of substitution between housing types and $\alpha_{h,\ell}(d)$ summarizes the population-weighted preference for type h at (ℓ, d). Differences in average tastes and spatial sorting both manifest as variation in $\alpha_{h,\ell}(d)$ across locations, while heterogeneity in tastes across individuals within a location is governed by σ.[10] Standard CES demand gives relative quantities

$$\frac{q_{M,\ell}(d)}{q_{S,\ell}(d)} = \frac{\alpha_{M,\ell}(d)}{\alpha_{S,\ell}(d)} \left(\frac{p_{M,\ell}(d)}{p_{S,\ell}(d)}\right)^{-\sigma},$$

or in logs,

$$\Delta_h \ln q_{h,\ell}(d) = \ln\left(\frac{\alpha_{M,\ell}(d)}{\alpha_{S,\ell}(d)}\right) - \sigma \Delta_h \ln p_{h,\ell}(d), \tag{1}$$

where Δ_h denotes the difference across housing types.

3.2 Cross-boundary substitution

Equation (1) is best read as a *lower-tier* demand system, conditional on the side of the boundary a household locates on. Households may also substitute across the boundary itself. I nest the within-side choice inside a standard upper-tier problem of choosing which side to live on, with

[10] The representative-agent CES form arises as the aggregate of individual households each choosing one home type under idiosyncratic extreme-value taste draws (Anderson et al., 1992).

cross-side elasticity η, derived in full in Appendix Section A. For identification, one property of that structure matters: the upper tier shifts the *level* of housing demand on each side but leaves the within-side relative demand for MH versus site-built housing in Equation (1) unchanged. The design therefore allows cross-boundary substitution: households deterred from buying MH inside the city may reappear as MH buyers just outside it. It requires *weak separability*: municipal regulation must not shift the relative taste for MH versus site-built housing discontinuously at the boundary, beyond its effect through relative prices inside the municipality itself.[11]

A stronger alternative would be a single demand system over the four options (M, I), (S, I), (M, U), and (S, U) with unrestricted cross-effects. Under that structure, a city tax on MH could directly change the county-side MH–site-built ratio, and the boundary RDD would conflate within-location substitution with resorting across space. The nested formulation rules out that confounding channel while allowing regulation to reallocate households across locations; a nested logit with a lower-tier discrete choice between MH and site-built homes has the same implication.

3.3 Supply

Houses are produced by perfectly competitive developers and sold at marginal cost. The sale price of a detached single-family home (SFH) of type $h \in \{M, S\}$ at location (ℓ, d) decomposes additively into three components:

$$p_{h,\ell}(d) = \underbrace{c_h}_{\substack{\text{Construction}\\\text{cost}}} + \underbrace{r_\ell(d) \cdot \bar{T}}_{\substack{\text{Land}\\\text{cost}}} + \underbrace{t_{h,\ell}}_{\substack{\text{Regulatory}\\\text{cost}}} . \tag{2}$$

The construction cost c_h captures non-land production inputs—factory fabrication and transport for MH, on-site labor and materials for site-built homes—and is type-specific but location-invariant. The land cost is the product of the per-acre land rent $r_\ell(d)$ that a parcel at (ℓ, d) would command absent zoning and the lot size $\bar{T}$ required to site one detached SFH in the sample. The regulatory wedge $t_{h,\ell}$ is measured in dollars per home and captures permit fees, design-standard compliance costs, quota rents capitalized into the price of restrictively zoned land, and other type- and side-specific local barriers (Glaeser and Gyourko, 2002, 2018). I maintain that this undistorted per-acre land rent is common to both housing types at any given (ℓ, d), formalized as Assumption 2 in Section 3.5.[12]

Equation (2) can be microfounded by a Leontief production technology in which non-land capital and land are perfect complements: a fixed-lot, fixed-structure recipe produces one detached home, with no scope for substitution between structure and land on the intensive margin. The Leontief restriction is reasonable for the detached-SFH sample, where lot configuration and dwelling footprint are largely fixed by site-development practice and manufactured homes are available in a limited set of single-floor configurations. Because I condition on lot size and living area in the empirical analysis, the assumption of perfect complements does not imply that developers never

[11] Weak separability could fail if regulation diverts would-be MH buyers from inside the city to the immediately adjacent unincorporated area. The unincorporated MH share would then be inflated by displaced demand, and the estimated share discontinuity would overstate the within-location effect of regulation on housing composition. The sign of this bias is positive—it inflates the measured share effect and the recovered substitution elasticity—but it leaves the dollar wedge t unaffected, since t is identified from the price double-difference in Section 3.3 and does not depend on where marginal buyers locate.

[12] In jurisdictions that prohibit MH outright, only site-built homes are constructed and $r_\ell(d)$ reflects site-built valuations alone, so the framework does not recover t in those areas. I return to this limitation in the welfare analysis (Section 3.6).

substitute between land and structure; it implies that they are effectively unable to do so within a fine-grained product segment.[13]

The parameter of interest is the differential regulatory wedge on manufactured homes inside incorporated areas, in dollars:

$$t \equiv \Delta_\ell \Delta_h t_{h,\ell} = \underbrace{[t_{M,I} - t_{S,I}]}_{\text{MH–site-built wedge inside}} - \underbrace{[t_{M,U} - t_{S,U}]}_{\text{MH–site-built wedge outside}}. \tag{3}$$

Applying $\Delta_\ell \Delta_h$ to Equation (2) gives

$$\Delta_\ell \Delta_h \, p_{h,\ell}(d) = t \qquad \text{for all } d, \tag{4}$$

since construction costs c_h are location-invariant and difference out in Δ_ℓ, while land costs $r_\ell(d)\bar{T}$ are type-invariant at every (ℓ, d) and difference out in Δ_h. The double difference recovers t in dollars at every distance from the boundary, not only in the limit $d \to 0$. Under a multiplicative (Cobb-Douglas) alternative, land enters log prices with a type-specific share, the land term does not difference out in Δ_h, and t is identified only at the cutoff under an additional cross-side continuity assumption on per-acre rents.

The equilibrium condition (Section 3.4) involves the wedge in ad-valorem rather than dollar form,

$$\tau \equiv \lim_{d \to 0} \Delta_\ell \Delta_h \ln p_{h,\ell}(d).$$

The log double difference depends on both level discontinuities at the boundary: the differential wedge t and the common site-built jump $b \equiv \lim_{d \to 0} \Delta_\ell \, p_{S,\ell}(d)$, which absorbs any discontinuity in per-acre land rents together with any difference in the site-built regulatory wedge. Appendix B derives the mapping,

$$\tau = \ln\left(1 + \frac{b+t}{p_{M,U}}\right) - \ln\left(1 + \frac{b}{p_{S,U}}\right) \approx \frac{t}{p_{M,U}} + b\left(\frac{1}{p_{M,U}} - \frac{1}{p_{S,U}}\right), \tag{5}$$

where $p_{h,U}$ is the price of a type-h home just outside the boundary. The leading term is the dollar wedge expressed as a fraction of the price of the taxed good; the second term reflects that a dollar jump common to both types is a different proportional change for each, and vanishes when $b = 0$. Both t and b are estimated by the price RDD (Section 4.1), so the conversion introduces no objects beyond those already estimated.

The empirical specification (Section 4.1) takes the level price $p_{h,\ell}(d)$ as the dependent variable and includes a common distance polynomial for both housing types, operationalizing the supply-side restriction that land enters prices additively through a type-independent index $r_\ell(d)\bar{T}$. A distance gradient that differed across MH and site-built homes would imply differential capitalization of locational amenities into the two housing types and is ruled out by type-independent per-acre rents.

[13] A theoretical literature in urban economics derives the same additive cost structure from a binding minimum lot size, under which $\bar{T}$ is the regulatory floor rather than an unconstrained cost-minimizing choice, and this form has been adapted for empirical work on land use regulation (Saiz, 2010; Gyourko and Krimmel, 2021). Notably, Banzhaf and Mangum (2019) argue that zoning regulations will be capitalized in home prices as an additive per-unit "ticket", rather than as a proportional increase, and show that in this case a traditional hedonic model using log price as the outcome is mis-specified.

3.4 Equilibrium

Applying Δ_ℓ to Equation (1) yields the equilibrium condition:

$$\lim_{d\to 0} \Delta_\ell \Delta_h \ln q_{h,\ell}(d) = \lim_{d\to 0} \Delta_\ell \ln\left(\frac{\alpha_{M,\ell}(d)}{\alpha_{S,\ell}(d)}\right) - \sigma \lim_{d\to 0} \Delta_\ell \Delta_h \ln p_{h,\ell}(d). \tag{6}$$

The change in relative quantities at the boundary equals the change in relative demands minus the change in relative log prices, scaled by the substitution elasticity: when σ is large and the two types are close substitutes, a small differential tax causes a large shift away from MH. The supply side (Equation (4)) pins the level price discontinuity to t in dollars; the log-relative form needed here is the ad-valorem wedge τ of Equation (5).

3.5 Identification

Identification of σ and t from Equation (6) rests on two assumptions. Assumption 1 restricts the demand side and delivers the quantity discontinuity; Assumption 2 restricts the supply side and delivers the price wedge. I state and discuss each in turn, then combine them.

Assumption 1: relative demand continuity.

(i) *Within-side smoothness.* For each $\ell \in \{I, U\}$, the one-sided limit $\lim_{d\to 0^+} \alpha_{h,\ell}(d)$ exists for $h \in \{M, S\}$.

(ii) *Cross-side continuity of the relative weight.* $\lim_{d\to 0} \Delta_\ell \ln(\alpha_{M,\ell}(d)/\alpha_{S,\ell}(d)) = 0$.

Assumption 1 is a restriction on the unobserved distribution of preferences at the boundary. It permits level differences in α_M and α_S individually—households just inside the city may differ from those just outside in ways that shift their demand for housing of both types—but rules out a discrete jump in the *ratio* α_M/α_S at $d = 0$. Households on either side of the boundary need not be alike; they need only value MH relative to site-built housing in the same way. This condition makes the share discontinuity attributable to the regulatory wedge rather than to a compositional break in who lives where.

Two sufficient conditions help fix intuition. If utility is additively separable between housing type and local amenities, then α_M/α_S does not depend on amenities at all, and Assumption 1 holds regardless of sorting or amenity discontinuities. Alternatively, if amenities are continuous at the boundary, then no preference-based sorting on relative MH demand can occur, and Assumption 1 holds even under non-separable utility. My identifying assumption is strictly weaker than either: it permits non-separable utility, discontinuous amenities, and sorting on preferences, provided they do not shift the relative-preference mean at the cutoff.

Three features of the research design support Assumption 1. First, each boundary area is defined within a single school district, removing the main channel for preference-based sorting: a household crossing the boundary retains the same access to public schools. Second, the CDP placebo (Table 9) provides a direct check: Census Designated Places have amenity composition similar to incorporated places but lack zoning authority, and the absence of a share discontinuity at CDP boundaries indicates that preference-based sorting does not on its own generate the observed pattern. Third, the null estimate on γ_1 in the site-built price RDD (Table 3) rules out the bid-rent channel through which preference-based sorting on relative MH demand would operate: a population with differentially high relative demand for site-built homes would capitalize that preference into the price of land suitable for site-built homes, and the null γ_1 is inconsistent with such capitalization. None is a direct test of Assumption 1; each rules out a specific mechanism by which it could fail. A direct test of Assumption 1 is possible in principle but requires household-level demographic and choice data sufficient to estimate the distribution of α_M/α_S and check whether it shifts at the boundary, which is beyond the scope of the parcel-level records used here.

Assumption 2: type-independent land cost.

(i) *Equality of land rents absent zoning.* The per-acre land rent that a parcel would command absent zoning, $r_\ell(d)$, is common to both housing types at every (ℓ, d): $r_{h,\ell}(d) = r_\ell(d)$, so that $\Delta_h\, r_{h,\ell}(d) = 0 \;\forall\, (\ell, d)$.

Assumption 2 concerns the land rent that would prevail absent zoning, not observed land prices. If zoning makes MH-suitable land artificially scarce, the resulting land-price gap is a quota rent capitalized into MH lots and is itself part of the differential wedge t the design recovers—not a violation of the assumption. The assumption fails only if *undistorted* land rents differ by type, which requires a complementarity between structure type and unpriced parcel amenities—for instance, if site-built homes are built on more desirable land, such as waterfront parcels. Even then, such a differential biases the double difference only if it shifts discontinuously at the boundary: a type-differential in land quality common to both sides differences out in Δ_ℓ. The evidence points against an upward bias from this channel, since MH inside cities sit on smaller lots in lower-income neighborhoods (Section 4.4), the opposite of the land-quality advantage that would inflate t.

What Assumption 2 buys is that the land term drops out of the price double difference *at every distance*, not merely in the limit at the boundary. Because land enters Equation (2) additively and identically for both types, $\Delta_h r_\ell(d) = 0$ at every (ℓ, d), so the level-price design does not require land rents themselves to be continuous across the boundary: any cross-boundary jump in r_ℓ—from city services, taxes, or amenities capitalized into land—differences out within Δ_h without further restriction.

This is the same logic that underlies the geographic difference-in-discontinuities design (Butts, 2023). In that setting, treatment switches on at a spatial boundary, but so do other determinants of the outcome, so a simple cross-boundary RDD is contaminated by a discontinuity that has nothing to do with treatment. The remedy is to difference along a second dimension—in the canonical case, time—on which the confounding discontinuity is assumed constant: the pre-period boundary jump estimates the confound and is netted out of the post-period jump. Here, housing type h plays the role that time plays there. The confound is the discontinuity in land rents at the city line; site-built homes play the role of the pre-period, absorbing that jump; and differencing MH against site-built homes at the same location nets it out. Assumption 2 is the analogue of the assumption that the confounding discontinuity is constant across periods, and it is what makes $\Delta_\ell\Delta_h\, p$ interpretable as the regulatory wedge alone.

Recovering σ and t. Combined with Equation (2), Assumption 2 delivers $\Delta_\ell\Delta_h\, p_{h,\ell}(d) = t$ at every d, since construction costs difference out in Δ_ℓ and land costs difference out in Δ_h at every location. Substituting the ad-valorem wedge $\Delta_\ell\Delta_h \ln p = \tau$ of Equation (5) into Equation (6) and imposing Assumption 1(ii) reduces the equilibrium condition to

$$\lim_{d\to 0} \Delta_\ell\Delta_h \ln q_{h,\ell}(d) = -\sigma\,\tau, \tag{7}$$

from which σ and t are recovered using the two boundary discontinuities estimated by the RDD. With two goods, relative quantities and shares carry the same information: $q_M/q_S = s_M/(1-s_M)$, so the left-hand side of Equation (7) is the boundary discontinuity in the log-odds of the MH share, denoted $\beta^{RDD,q}$, and is estimated directly by a logit RDD on the MH indicator (Equation (13)). The right-hand side is obtained from the level price discontinuities through Equation (5). Appendix B collects both mappings.

The model accommodates price spillovers through equilibrium land rents: restrictions on MH in cities may cause households to bid up the price of land in the neighboring county, but per-acre rents are common to both housing types at every (ℓ, d) and difference out within Δ_h identically. It also accommodates substitution across the city-county line through the upper-tier cross-boundary choice described above. What the RDD identifies is therefore the distortion to the composition of housing types within locations, not the full effect of zoning on the spatial distribution of households. I ignore any effects of zoning on aggregate housing supply in order to focus on how zoning distorts the composition of housing types.

3.6 Welfare

The CES structure delivers a closed-form expression for the welfare cost of the implicit tax. Evaluating at a representative inside location, the dual price index for the lower-tier CES aggregate over MH and site-built homes is:

$$P = \left[\alpha_M p_M^{1-\sigma} + \alpha_S p_S^{1-\sigma}\right]^{\frac{1}{1-\sigma}}$$

where α_M, α_S and p_M, p_S refer to inside-boundary values. Inside cities, the implicit tax raises the effective price of MH by the ad-valorem wedge τ, from p_M to $e^\tau p_M$. Let ω_M denote the expenditure share on MH absent the implicit tax. I measure the cost of the tax by its compensating variation: the additional income required at taxed prices to leave a household as well off as it was absent the tax.

Proposition 1. *Under the CES demand structure in Section 3.1, the nested cross-boundary choice in Section 3.2, and the competitive supply assumptions in Section 3.3, the compensating variation of the implicit tax on manufactured homes, expressed as a fraction of housing expenditure at a representative inside location, is*

$$\frac{CV}{E} = \left[\left((e^\tau)^{1-\sigma} - 1\right)\omega_M + 1\right]^{\frac{1}{1-\sigma}} - 1. \tag{8}$$

Three statistics are jointly sufficient for this cost: the elasticity of substitution σ, the ad-valorem tax τ, and the counterfactual MH expenditure share ω_M.

Proof. See Appendix C. □

Two features of this welfare measure matter. First, the expression captures the full private cost of the tax to households—the compensating variation—not merely the Harberger deadweight loss triangle. In a standard analysis of an explicit tax, the distinction matters: the tax revenue rectangle is a transfer, not a social loss, so only the triangle represents inefficiency. The implicit tax on MH, however, generates no revenue. The costs that constitute the implicit tax—compliance expenses, regulatory friction, restrictions to less desirable parcels—are real resource costs, not transfers to a tax-collecting authority. The full compensating variation is therefore a measure of the efficiency loss from anti-MH regulations.[14]

Second, the welfare cost is a function of the *counterfactual* expenditure share ω_M—the share of spending on MH that would prevail absent regulation—rather than the observed share. Because regulation suppresses MH consumption, the observed expenditure share understates the relevant base for welfare calculations. I calibrate ω_M using the estimated parameters in Section 7. This is a partial-equilibrium welfare calculation for the lower-tier distortion in housing composition within locations. A full welfare analysis that also valued resorting across the boundary would additionally require the upper-tier cross-side elasticity η and data on how overall demand shifts across space.

4 Estimation

I estimate the two key parameters of the model, t and σ, using a two-step procedure. First, I estimate a boundary RDD to recover the discontinuity in MH shares at city boundaries, $\hat{\beta}^{RDD,q}$. Second, I estimate a parallel RDD to recover the discontinuity in relative prices at city boundaries, $\hat{\beta}^{RDD,p}$. I combine these estimates with the average city MH share, $\bar{s}_M$, to solve for t and σ using the equilibrium condition in Equation (7).

4.1 Empirical Strategy

The model imposes that comparisons across the boundary be made within a product segment: an apples-to-apples contrast of MH and site-built homes of the same size, lot, and vintage. I define product segment j by the interaction of living area decile, parcel size decile, and built decade (1980s, 1990s, 2000s, 2010s), yielding 400 possible segments, of which only 8 contain no MH. This granular segmentation alleviates the concern that an observed decrease in MH shares at municipal boundaries arises from a general municipal bias against small homes or against larger lots, characteristics that are bundled with MH: Figure 1 plots MH shares across the living area × lot size grid and shows that MH are concentrated among small homes on the largest lots.

Motivated by the log-odds identity in Equation (13), I first estimate a boundary RDD on the MH indicator using a logit specification:

$$\Pr(\mathrm{MH}_{ib} = 1 \mid \cdot) = \Lambda\big(\beta^{RDD,q} \cdot \mathrm{Incorp}_i + f(\mathrm{dist}_i) \times \mathrm{Incorp}_i + \alpha_{j(i),b}\big)\,, \tag{9}$$

where MH_{ib} is an indicator for whether parcel i in boundary area b is a manufactured home, $\Lambda(\cdot)$ is the logistic CDF, Incorp_i indicates whether the parcel is inside the city, $f(\mathrm{dist}_i)$ is a third-degree polynomial in distance to the boundary, and $\alpha_{j(i),b}$ are fixed effects for each product segment by boundary area. The coefficient of interest, $\beta^{RDD,q}$, is the within-segment discontinuity in the log-odds of MH at the city boundary; by Equation (13) it is also the discontinuity in the relative log quantity of MH versus site-built homes. A linear-probability variant of Equation (9) is reported in Appendix E as Equation (14).

[14] To the extent that some components of the implicit tax are transfers (e.g., permit fees that fund municipal services, or zoning restrictions which are capitalized into lot prices), the efficiency loss would be smaller than the full compensating variation. This welfare measure is thus an upper bound on the deadweight loss. It is also a measure of private cost only: if manufactured homes impose negative externalities on neighbors, part of the implicit tax may be corrective and the efficiency loss smaller still. Section 7.3 returns to this possibility.

Next, motivated by the expression for the differential implicit tax in Equation (3), I extend the RDD framework to estimate a discontinuity in relative prices at city boundaries. The approach is similar to a "difference-in-discontinuities" design (as in Butts (2023)), except that I leverage cross-sectional variation in housing types rather than a change in treatment over time. Intuitively, site-built prices absorb any local changes in amenities or common land costs, so any change in the relative price must be due to differential regulation. I estimate the following regression:

$$p_{ib} = \gamma_1 \text{Incorp.}_i + \gamma_2 \text{MH}_i + \beta^{RDD,p} \cdot \text{Incorp}_i \times \text{MH}_i + f(\text{dist}_i) \times \text{Incorp}_i + \alpha_{j(i),b} + \epsilon_{ib}, \qquad (10)$$

where p_{ib} is the sale price of parcel i in boundary area b and all other terms are as defined previously. The parameters γ_1 and γ_2 capture the average additive price effect for homes inside incorporated areas and mobile homes, respectively, so that the coefficient on their interaction, $\beta^{RDD,p}$, captures the discontinuity in the MH–site-built price *gap* at the city boundary measured in dollars. Here, unlike in the share RDD, the inclusion of product segment $\times$ boundary area fixed effects is critical, since the equivalence of land costs across housing types holds only conditional on lot size.

The identifying assumption is that non-regulatory factors affecting *within-location relative* demand for MH are continuous at the boundary, as discussed in Section 3.5.

4.2 Boundary Effects on MH Shares

Results are provided in Table 2. Consistent with the descriptive evidence, the MH share falls sharply on the city side of the boundary, and the estimated discontinuity is large and significant at the 1% level in both specifications. Column (1) includes only boundary area fixed effects, so the estimates conflate the effect of any differential regulatory burden on MH with the effect of municipal regulations that target other housing characteristics typically bundled with MH, such as small homes or relatively larger lots. My preferred specification, Column 2, interacts boundary area FEs with product segment FEs to estimate within-segment differences in shares. The estimated discontinuity in the log-odds of MH at the city boundary is $\hat{\beta}^{RDD,q} \approx -1.04$.

To translate this estimate into a concrete change in housing composition, I invert the log-odds identity at the average inside-city MH share. The observed share inside cities is $\bar{s}_M = 2.1\%$, so absent the discontinuity the MH share inside cities would be

$$\tilde{s}_M = \Lambda\Big(\Lambda^{-1}(\bar{s}_M) - \hat{\beta}^{RDD,q}\Big) \approx 5.8\%.$$

Municipal regulations thus reduce the MH share inside cities from roughly 5.8% to 2.1%, a decline of about 63% relative to the counterfactual share.

4.3 Boundary Effects on Relative Prices

Results are shown in Table 3. Crossing from inside the city to the adjacent unincorporated area, quality-adjusted manufactured home prices fall by about \$7,700 relative to site-built homes, an effect significant at the 1% level. Mapping this estimate to the log double difference at the boundary mean price levels (Section 4.5) gives a relative effect of about 5.9%. Because the double difference nets out movements common to both housing types, this relative movement is a price premium on manufactured homes inside cities — the wedge the model interprets as the implicit regulatory tax.

The estimate depends on conditioning finely on characteristics. With boundary fixed effects alone (Column 1), the differential is \$14,900 (s.e. \$7,100); adding boundary × product-segment fixed effects (Column 2) roughly halves it to \$7,700 and cuts the standard error to \$2,500, while R^2 rises from 0.38 to 0.72. Both movements have the same source. The coarse specification compares manufactured to site-built homes of different sizes, lots, and vintages, so its interaction term picks up the compositional differences documented in Section 4.4 — inside-city MH are larger structures on smaller lots — along with the regulatory wedge. Restricting comparisons to homes in the same product cell removes that compositional component, and because most price variation in this market is characteristic variation, it also removes most of the residual variance, so the coefficient becomes markedly more precise even though the sample is nearly unchanged.

The estimated effect of incorporated status on site-built home prices is economically small and statistically insignificant: γ_1 implies a price effect of 3.7% of the boundary mean site-built price, with a 95% confidence interval running from -0.8% to 8.2%. The interval is wide enough that modest common shifts cannot be excluded, so I do not read this estimate as a tight bound on amenity differences. It does rule out the bid-rent channel through which sorting on *relative* MH demand would operate, the object the design requires (Section 3.5). Many municipal amenities, such as libraries and parks, are non-excludable, so any capitalization effects will be absorbed by controls for distance to the boundary. Moreover, municipalities typically represent a bundle of amenities and higher taxes, so that the marginal homebuyer may be indifferent between the two sides of the boundary if local public services are valued roughly at cost.

Finally, I estimate a large and significant negative effect of MH status on home prices. The average price for MH is roughly 58% lower than site-built homes in the same boundary area. Adjusting for differences in living area, lot size, and age with product segment fixed-effects decreases the discount to 41%, well in line with industry estimates of the cost savings from factory-built housing: traditional single- and double-wide manufactured homes enjoy a cost advantage of between 40% and 65% over similar site-built housing (Herbert et al., 2023). To my knowledge, this is the first empirical estimate of the price differential using market prices rather than construction cost estimates or surveys.[15]

The fact that MH enjoy substantial price discounts relative to site-built homes while simultaneously representing a small share of the housing stock indicates that their unobserved quality must be substantially lower than that of site-built homes. These quality differences may arise from more rapid depreciation, greater exposure to weather risks (Sutter and Poitras, 2010), or social stigma. However, they are not an issue for my identification strategy so long as they do not change differentially at municipal boundaries.

4.4 Boundary Effects on Physical and Neighborhood Characteristics

The estimates so far compress municipal regulation into a single price wedge. But zoning ordinances do not operate on prices directly; they operate on characteristics and locations: design and installation standards dictate what a manufactured home must be to be permitted, and district-level placement rules dictate where it may be located. If these non-price margins bind, they should affect the characteristics of homes near the boundary. In this section I examine six physical and

[15] Mimura et al. (2013) use American Housing Survey data on self-reported home values for a sample of rural houses to estimate a hedonic price model. They find a MH discount of nearly 75% conditional on observable characteristics. Because their geographic detail is limited, they cannot adjust for local amenities or neighborhood characteristics, biasing their estimate upwards in magnitude if MH are differentially located in lower-quality areas.

neighborhood characteristics as outcomes, and estimate whether each changes at the boundary separately for site-built and manufactured homes. The exercise serves two purposes. For site-built homes, smoothness of characteristics across the boundary is a validation test of the design, analogous to a covariate balance test in a standard RDD. For manufactured homes, discontinuities are informative about mechanism: they reveal which margins of regulation lie behind the price wedge and discipline its interpretation.

Table 4 reports the results. Each column re-estimates the difference-in-discontinuities specification of Equation (10), replacing the sale price with the characteristic listed in the column header: the neighborhood's Black population share, per-capita income, and population density, and the parcel's lot size, living area, and year built. The *Incorp.* row is the boundary discontinuity for site-built homes, the omitted category; the *Incorp.* $\times$ *MH* row is the differential discontinuity for manufactured homes. Unlike the preferred price specification, these regressions include only boundary area fixed effects: living area, lot size, and vintage define the product segments, so segment fixed effects would absorb the outcomes under study.

Site-built homes do not change significantly at municipal boundaries on any of the six dimensions. This validates the RDD design: the boundary does not cut through pre-existing gradients in neighborhood quality or parcel characteristics for the dominant housing type.

Manufactured homes, by contrast, shift at the boundary in ways that reflect two distinct margins of regulation. First, design and quality mandates appear to raise the quality of MH permitted within cities: MH inside cities are roughly 62 square feet larger and are located in neighborhoods with per-capita incomes roughly $4,800 higher. To the extent consumers value these improvements, t captures not only artificial supply restriction but also mandated quality upgrading, and the welfare loss from pure exclusion is smaller than the price gap alone implies. Second, the lot size and neighborhood composition results suggest placement restrictions that confine MH to specific zones: relative to site-built homes, MH inside cities occupy lots that are 0.32 acres smaller than comparable MH outside cities, and are correspondingly located in denser neighborhoods and in neighborhoods with Black population shares 0.6 percentage points higher. These patterns suggest that within cities, MH are directed toward smaller and less desirable parcels, which suppresses their prices relative to the counterfactual and biases t downward. The two physical-characteristic effects—lot size and living area—are significant at the 1% level and the Black-share effect at the 5% level, while the income and density effects reach only the 10% level; the estimates are modest in magnitude but run consistently in the direction of both mechanisms. On net, the relative price effect is positive, so the scarcity and quality-mandate channels dominate any placement penalty.

Importantly, the demographic effects should be interpreted as lower bounds in magnitude. Neighborhood characteristics are assigned from the Environmental Impact Frame (EIF), a dataset reporting population counts by demographic category at the 0.01-degree grid cell level (roughly 1 km) using administrative tax and census records. Because each parcel inherits the average over its grid cell, the assignment spatially smooths demographics across the boundary: cells that straddle or sit near the city line blend incorporated and unincorporated populations, mechanically compressing the measured cross-boundary discontinuity toward zero. The true differential changes in neighborhood demographics are therefore likely larger than those reported.

These characteristic discontinuities also carry a methodological implication: because MH inside cities are systematically larger and on smaller lots than their unincorporated counterparts, a raw price comparison conflates regulatory costs with compositional differences in the housing stock. Conditioning on living area decile, lot size decile, and construction decade—the product segment fixed effects in the preferred specification—ensures the estimated price gap reflects the cost of regulation for comparable units while holding constant land acquisition and construction costs.

4.5 Recovering Model Parameters

Because the price RDD is estimated in levels, I convert its two dollar coefficients into the ad-valorem wedge using the sample counterpart of Equation (5) (Appendix B),

$$\widehat{\tau} = \ln\left(1 + \frac{\hat{\gamma}_1 + \hat{\beta}^{RDD,p}}{\bar{p}_{M,U}}\right) - \ln\left(1 + \frac{\hat{\gamma}_1}{\bar{p}_{S,U}}\right),$$

where $\bar{p}_{h,U}$ is the mean price of a type-h home just outside the boundary. At the coefficients in Table 3 and the boundary mean prices, $\widehat{\tau} \approx 0.059$: the dollar wedge of \$7,700 implies an implicit tax of roughly 6% of the price of a manufactured home.[16] The conversion involves two implementation choices, the treatment of the common jump $\hat{\gamma}_1$ and the price base; alternatives—setting the common jump to zero, or normalizing on a base built from the two MH cells—yield $\widehat{\tau}$ between 0.035 and 0.055, and the elasticities they imply lie well inside the confidence interval reported in Table 5.

Combining this mapped price effect with the equilibrium condition in Equation (7) and the log-odds identity in Equation (13) allows me to solve for σ, the elasticity of substitution between MH and site-built homes:

$$\sigma = -\frac{\hat{\beta}^{RDD,q}}{\widehat{\tau}} = -\frac{-1.045}{0.059} \approx 17.6.$$

An elasticity of substitution of roughly 18 suggests that MH and site-built homes are close substitutes within a product segment. In other words, among households shopping for homes of a given living area, vintage, and plot size on a given side of a specific boundary area, MH and site-built homes are highly interchangeable.

The high substitution elasticity partly reflects small baseline MH shares and the granular product segmentation. If a price change caused even one in one hundred homebuyers to switch from site-built to MH, the MH share would increase substantially.

Because σ, the counterfactual share, and the welfare objects in Section 7 are nonlinear combinations of the two RDD coefficients, I compute confidence intervals for all of the structural objects with a cluster bootstrap and collect them in Table 5. Each of the 999 draws resamples boundary areas with replacement, re-estimates the share and price RDDs jointly, and recomputes every object, so the intervals account for the covariance between the two discontinuities. I report percentile intervals, which are robust to the heavy tails that ratio objects such as σ inherit from draws where the price coefficient in the denominator is small. The elasticity is estimated imprecisely, with a 95% interval of $[3.6, 60.8]$. Much of that width comes from the conversion of the dollar wedge to ad-valorem form rather than from the share discontinuity: because the common jump $\hat{\gamma}_1$ enters $\widehat{\tau}$ and is itself imprecisely estimated, the interval for $\widehat{\tau}$ is roughly twice as wide as it would be were the common jump set to zero, and σ inherits that width through the denominator. The estimates

[16] Interpreted structurally, the level specification treats the municipal wedge as additive in dollars, and Equation (5) converts it to the proportional form the CES mapping requires. The backup log-price RDD estimates the log double difference directly and yields the larger value 0.076. The two need not coincide: the log specification recovers the proportional wedge only if the underlying regulatory cost is itself proportional to the price, whereas the additive structure here implies a fixed dollar cost, so the log specification is misspecified in exactly the way Banzhaf and Mangum (2019) describe.

imply substantial substitution: the lower end of the interval is well above unit elasticity, although the data cannot pin down the exact magnitude. The wedge $\widehat{\tau}$ itself remains bounded away from zero in every bootstrap draw. The welfare loss is estimated far more precisely than σ, a point I return to in Section 7.

5 Robustness

I assess the robustness of the baseline share and price RDDs along six dimensions. Two conclusions emerge. The share discontinuity holds its sign and significance throughout but attenuates in narrow bandwidths, where incorporated boundaries are sparse. The level price discontinuity — the input for the welfare calculation — is positive and significant in every bandwidth and polynomial specification, but it is sensitive to the treatment of extreme sales, which is why the baseline winsorizes the price tails. The connection to the welfare estimate is direct: the loss is to first order $\omega_M \hat{\beta}^{RDD,p}$, so movements in the price wedge pass through to the welfare cost roughly one-for-one, whereas imprecision in σ does not.

First, I vary the bandwidth and polynomial order of the global RDD. Table 6 reports the share results: column (1) reproduces the baseline 2 km cubic specification; columns (2)–(3) hold the cubic fixed and shrink the symmetric bandwidth to 1 km and 0.5 km; columns (4)–(5) hold the 2 km bandwidth fixed and reduce the polynomial order to linear and quadratic. The estimated log-odds discontinuity is essentially invariant to polynomial order but attenuates sharply at narrower bandwidths, becoming statistically indistinguishable from zero at 1 km and 0.5 km. This pattern reflects the fact that incorporated boundaries are sparse: a 0.5 km window contains less than half the baseline sample, and within that narrow window the cubic polynomial absorbs much of the level shift in MH shares. Table 7 reports the analogous exercise for the price diff-in-discontinuities. The relative price effect $\hat{t}$ is positive and statistically significant in all five specifications, with point estimates ranging from \$7,700 to \$11,600 per home (compared to a baseline of \$7,700). The average MH discount is stable across columns. The site-built discontinuity is statistically insignificant in the cubic baseline and at narrower bandwidths (columns 1–3); it turns positive and significant only when the distance polynomial is restricted to linear or quadratic (columns 4–5, roughly \$7,000 to \$8,400), indicating that the cubic polynomial is needed to absorb a smooth land-cost gradient near the boundary and motivating its use in the baseline.

Second, I re-estimate the share RDD using the local-polynomial methods of Calonico et al. (2014) with MSE-optimal bandwidth selection (Table 8), residualizing the outcome on boundary × product-segment fixed effects before estimation. The estimated discontinuity is negative and highly significant in every specification. With residualized fixed effects, the MSE-optimal bandwidth is 0.30 km and the estimated discontinuity is about 1.4 percentage points—roughly 55% of the baseline 2 km estimate of 2.4 percentage points, consistent with the bandwidth sensitivity in Table 6.

Third, the baseline share RDD restricts to parcels with an observed market transaction, since the price RDD requires sale prices. To check that this restriction does not drive the share results, I re-estimate the share RDD on the full sample of parcels rather than only those that transact (the *All parcels* row of Appendix Table 15). The estimated discontinuity in the preferred segment × boundary specification is -2.2 percentage points, against -2.4 in the transacting-parcel baseline, indicating that selection into transactions is not the source of the baseline share result.

Fourth, I re-estimate the price diff-in-discontinuities using assessor-reported values rather than transaction prices (Table 11). I deflate assessed values to 2023 dollars using the same housing price index based on the assessment year. Assessor values may address bias in the estimated MH price discount if single-family homes are disproportionately likely to have a recorded market transaction as a “new” home. Their disadvantage is that they may be measured with error due to a limited set of comparables, attenuating the estimates. With only boundary fixed effects (Column 1), the estimated differential effect is $10,100, somewhat larger than but of the same order as the baseline transactions estimate of $7,700. Adding product-segment fixed effects (Column 2) attenuates the estimate to near zero ($300) and renders it statistically insignificant, suggesting that assessors apply uniform valuation rules within fine product cells that smooth over much of the geographic variation captured in transactions.

Fifth, I test whether the results hinge on the treatment of extreme sale prices. Appendix Section D re-estimates both RDDs leaving sale prices raw rather than winsorizing the tails at the 1st and 99th percentiles within housing type, as in the baseline. The share and log-price discontinuities are stable across the two treatments. The level price discontinuity, on which the welfare cost rests, is sensitive: with prices left raw, its standard error nearly doubles and it is no longer distinguishable from zero, and the implied welfare loss falls from 0.14% to 0.06% of expenditure. This is the expected consequence of estimating a level equation on a heavy-tailed outcome, and it is why the baseline winsorizes. A second sample rule, the requirement that lot size and living area be reported, also functions as a data-cleaning step: the parcels it drops are disproportionately manufactured (39% versus 6% of the retained sample), consistent with the removal of trailer-park and other manufactured units that CoreLogic does not identify.

Sixth, I run the share and price RDDs on Census Designated Places (CDPs), which are statistical areas defined by the Census Bureau and lack zoning authority. The share placebo (Table 9) yields a small and statistically insignificant discontinuity of -1 percentage point. The price placebo (Table 10) yields a CDP $\times$ MH coefficient of $4,200 that is not statistically significant. Together these placebos indicate that the baseline results are not driven by a generic spatial pattern at the edge of populated areas: although the CDP price point estimate is about half the baseline incorporation effect, it is imprecise and statistically indistinguishable from zero, whereas the incorporation effect is precisely estimated.

6 Heterogeneity

I next test whether the boundary-level price effects align with the institutional mechanisms discussed in Section 2.1. If the differential implicit tax reflects local political economy and the legal scope of municipal authority, then t should be larger in geographically smaller cities (where incumbents internalize a larger share of any nearby MH externality and organized opposition is more electorally effective, as in Hankinson, 2018 and Mast, 2024) and in states whose institutions grant cities broader and less-constrained authority over land use.

I focus the heterogeneity analysis on the boundary-level relative price effect, rather than the share effect, for one statistical reason and one conceptual reason. Statistically, shares are downstream of the implicit tax: the share discontinuity equals $-\sigma \cdot \tau$ in the model, so boundary-level share effects mix variation in t with variation in the substitution elasticity σ, and inherit the noise from both objects. Conceptually, the relative price effect is the more direct measure of the regulatory wedge: under the additive supply structure, $\beta_b^{RDD,p}$ recovers t in dollars at every d, with no first-stage elasticity entering between t and the outcome. Empirically, share effects are noisy across boundary areas and only weakly correlated with the moderators below; the price effects are more informative. I report share-effect regressions in Appendix Table 17 for completeness.

The boundary-specific price effects are estimated using a simplified version of Equation (10), fit separately for every boundary area with a linear control in distance and product-segment fixed effects retained within each boundary-specific regression:

$$p_{ib} = \gamma_1 \text{Incorp.}_i + \gamma_2 \text{MH}_i + \beta_b^{RDD,p} \cdot \text{Incorp.}_i \times \text{MH}_i + f(\text{dist}_i) \times \text{Incorp}_i + \alpha_{j(i),b} + \epsilon_{ib}. \quad (11)$$

Since each regression is split by boundary area, this is equivalent to boundary area × product segment fixed effects in the pooled specification and keeps the relative price effect from conflating regulation with boundary-specific changes in lot size, living area, or vintage composition. A key difference from the pooled specification is that the linear control for distance is estimated separately for each boundary area, allowing municipalities to have different spatial patterns of amenities and land costs.

Table 12 relates the resulting boundary-level price effects $\hat{\beta}_b^{RDD,p}$ to two sets of moderators motivated by the institutional discussion above. The first is the (log) land area of the incorporated place — a proxy for the collective-action mechanism. The second set captures the state-level legal environment for municipal regulation of MH: an indicator for states with strong statutory protections for MH (e.g., preemption of local MH bans or restrictions on local design authority) from Dawkins (2011); the parcel-weighted state-level mean of the Wharton Residential Land Use Regulatory Index as a survey measure of overall stringency (Gyourko et al., 2008); and an indicator for Dillon's Rule states. Column 1 includes only city extent; Column 2 adds the state institutional moderators; Column 3 replaces them with state fixed effects, isolating within-state variation in city size.

The estimates align with the institutional mechanism. The coefficient on log land area is negative and significant: doubling the area of the city reduces the differential implicit tax by roughly \$3,100, consistent with the collective-action mechanism that smaller cities can sustain stricter exclusion. Among state institutional moderators, Dillon's Rule states show economically and statistically smaller differential taxes — the direction predicted by the legal scope of municipal authority — and the Dawkins index of strong state MH protections enters negatively, indicating that statutory state preemption translates into smaller boundary-level wedges. The Wharton index loads positively, consistent with the same forces that produce broader regulatory stringency also producing larger MH-specific wedges. The land-area effect survives the inclusion of state fixed effects in Column 3, so it is not merely picking up state-level correlates of city size.

The map in Figure 4 reinforces the same point: state averages of $\hat{t}_b$ are largest in the Northeast and parts of the West, where state institutions grant cities broad regulatory authority, and smaller in the Sunbelt Dillon's Rule states. The pattern is suggestive rather than decisive: the moderators explain a modest share of total variation in boundary effects, leaving substantial residual heterogeneity attributable to specific ordinances, enforcement practices, and unobserved local demand. But the heterogeneity that the observables do explain points consistently back to the scope of municipal authority as the proximate driver of the differential implicit tax on MH.

7 Quantifying Welfare Costs

How costly are anti-MH zoning regulations for city residents? In this section, I calibrate the welfare expression from Proposition 1 using the RDD estimates. The key result is that the average cost to city residents is small—a fraction of one percent of housing costs—because MH account for a small share of housing expenditure even in the absence of regulation. However, the costs are concentrated among lower-income households who would otherwise choose MH.

7.1 Calibration

The welfare formula in Equation (8) requires three inputs: the elasticity of substitution σ, the ad-valorem tax τ, and the counterfactual MH expenditure share ω_M. The first two are estimated directly from the boundary RDDs. To obtain ω_M, I convert the counterfactual quantity share to an expenditure share using the no-tax MH price relative to site-built housing. Since observed inside-city MH prices include the implicit tax, I remove the estimated wedge before forming the counterfactual price ratio, giving $p_M/p_S \approx 0.6$:

$$\omega_M = \frac{s_M \cdot (p_M/p_S)}{s_M \cdot (p_M/p_S) + (1 - s_M)} \approx \frac{0.058 \times 0.6}{0.058 \times 0.6 + 0.942} \approx 0.034,$$

where s_M is the share of MH inside cities absent the implicit tax, or 5.8%. If not for municipal regulations, MH would account for about 3.4% of housing expenditure inside cities. Plugging in to Equation (8) yields

$$\frac{CV}{E} = \left[(1.06^{1-17.6} - 1) \times 0.034 + 1\right]^{\frac{1}{1-17.6}} - 1 \approx 0.13\%.$$

The welfare estimate is far better determined than σ, with a 95% interval of $[0.04, 0.24]\%$ against an interval for the elasticity that spans an order of magnitude (Table 5). The reason is that to first order the loss is $\omega_M \tau$, which does not involve the elasticity at all; σ enters the exact CES expression only through a higher-order curvature adjustment, so the imprecision in the elasticity does not propagate to the welfare cost. What uncertainty remains comes from the wedge and the counterfactual expenditure share, and it leaves the qualitative conclusion intact: the average loss is a fraction of one percent of housing expenditure across the interval.

7.2 Incidence

The average welfare cost of 0.13% of housing expenditure is small because the vast majority of households would not choose MH even absent regulation. But the costs are not evenly distributed. Households who choose MH despite the implicit tax bear the full 5.9% price increase on their housing expenditure. These are disproportionately lower-income households for whom housing is a larger share of total spending, implying that the welfare cost as a share of income is higher still. The average cost therefore understates the policy significance of anti-MH zoning for the households most affected by it.

7.3 Caveats

The welfare costs reflect only distortions to the composition of housing types in cities where manufactured homes are allowed but discouraged. To the extent that anti-MH regulations cause parcels to remain undeveloped and displace homebuyers to other locations, the welfare costs will be larger than the compositional effect alone. The costs are also limited by the framework, which requires price observations on both sides of the boundary and thus omits the effect of blanket bans (or prohibitively costly regulations) on MH which are imposed by some cities. In the opposite direction, the compensating variation measures private cost: if manufactured homes impose negative externalities on neighbors, part of the implicit tax is corrective and the efficiency loss is smaller than the estimate implies (Munneke and Slawson, 1999). Finally, any differential economies of scale in MH production (relative to site-built homes) would further increase welfare costs by limiting the size of the market for MH and driving up the marginal cost of production (Gann, 1996).

8 Conclusion

Municipal land use regulations that target MH have a significant effect on the housing stock. Cities discourage MH by imposing a differential implicit tax of about 5.9% on manufactured homes, reducing the MH share inside city boundaries from a counterfactual of roughly 5.8% to just 2.1% — a decline of about 63% relative to the counterfactual share. Such regulations may have undesirable distributional consequences by limiting the supply of affordable housing within cities. A natural conjecture, which the boundary design does not test, is that displacing manufactured housing to the urban fringe also pushes low-income households further from job centers and public services; tracing that displacement and its consequences for access is left to future work.

The prevalence of anti-MH zoning, even in relatively small agglomerations with little obvious market power, suggests that incentives to impose such regulations are widespread. Anti-MH regulations may be economically efficient if MH impose fiscal or aesthetic disamenities on city residents, justifying the implicit tax. If instead anti-MH regulations are used by incumbent homeowners to limit housing supply or by the site-built home industry to suppress competition, then the implicit tax represents an inefficient distortion that reduces aggregate welfare and generates rents for incumbent homeowners and builders.

MH also represent one possible response to the "strange and awful" path of construction productivity in the US (Goolsbee and Syverson, 2023). If higher quality MH can be produced at a 25% discount to site-built homes, they could provide affordable housing for millions of households. Municipalities should consider the trade-offs of anti-MH zoning carefully, especially as the industry develops high-quality models and federal policymakers expand support for factory-built housing.

9 Tables and Figures

Table 1: Summary Statistics

Variable	Incorporated	Unincorporated
N	4,472,601	1,786,644
Manufactured (%)	2.1 (14.4)	8.3 (27.6)
Square feet (1000s)	2.1 (0.9)	2.1 (0.9)
Acres	4.0 (5617.1)	1.2 (44.9)
Year built	2000 (10)	2000 (11)
Share Black	9.0 (12.6)	8.2 (12.1)
AGI ($1000s)	151 (81)	149 (85)
Pop. density (per sq. km)	1336 (1008)	905 (1033)
Home price ($1000s)	436.4 (290.6)	409.8 (287.2)

Notes: Summary statistics for single-family homes within 2 km of an incorporated place boundary built since 1980. Each cell reports the variable mean with its standard deviation in parentheses. Prices are deflated to real 2023 dollars using the Federal Housing Finance Agency Home Price Index (FHFA HPI). Household demographic characteristics are from the Census Bureau's Environmental Impact Frame (EIF) and correspond to the 0.01° grid cell nearest to each parcel. *AGI ($1000s)* is adjusted gross income aggregated to the household rather than the individual level, so it reports mean income per tax filing unit and is not comparable to a per capita measure.

Figure 1: Manufactured Home Share by Living Area and Parcel Size Decile

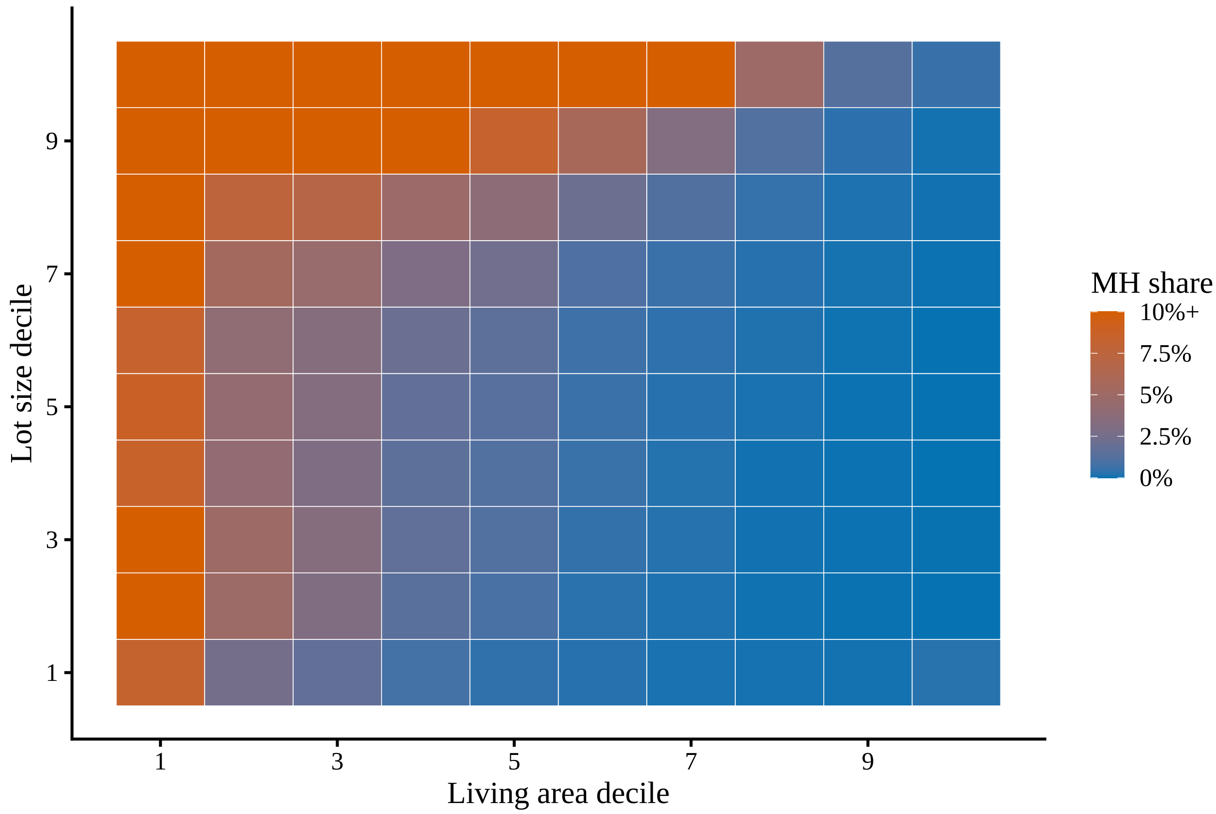

Notes: Figure reports the manufactured-home share among single-family homes by living area and parcel size decile. The sample includes single-family homes built after 1980 within 2km of a city line. Cells with a share of MH exceeding 10% are topcoded in the darkest shade of orange.

Figure 2: Housing Shares by Distance to City Boundaries

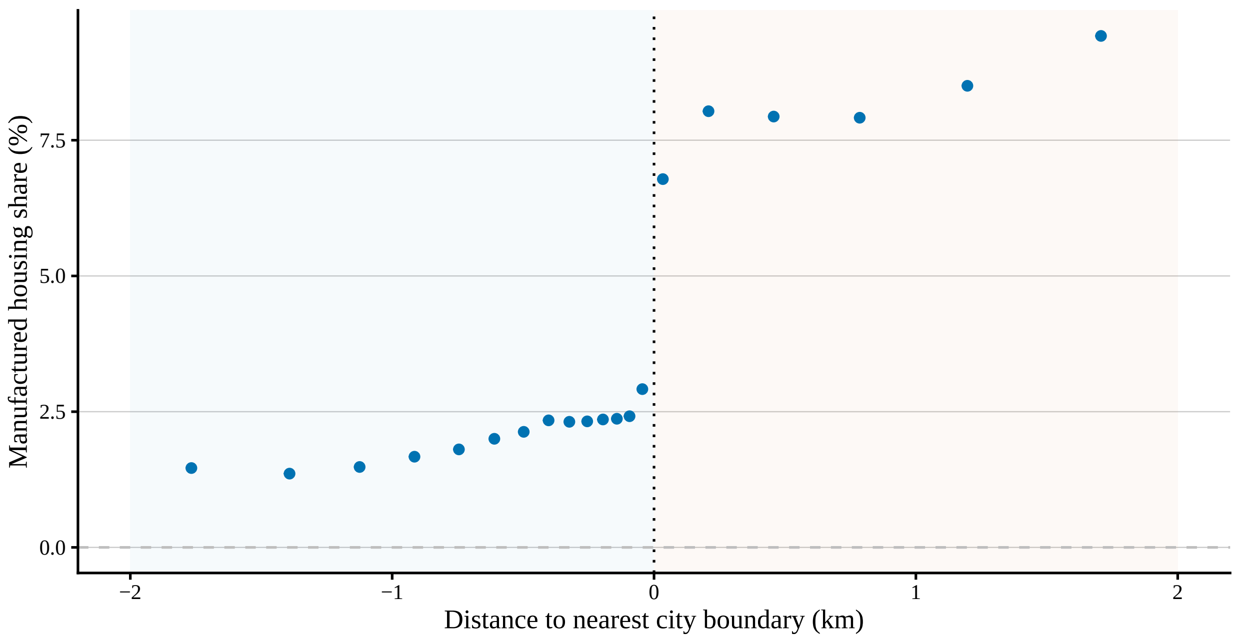

Notes: The figure shows a binned scatterplot of the probability of observing a mobile home. The plotted relationships between the outcome and the distance to the boundary adjust for county fixed effects following Cattaneo et al. (2024). Observations with negative distances are located inside an incorporated city. The figure is calculated using 6,259,245 parcels spread throughout 3,812 boundary areas.

Figure 3: Mobile and Site-Built Homes in Lumberton, NC

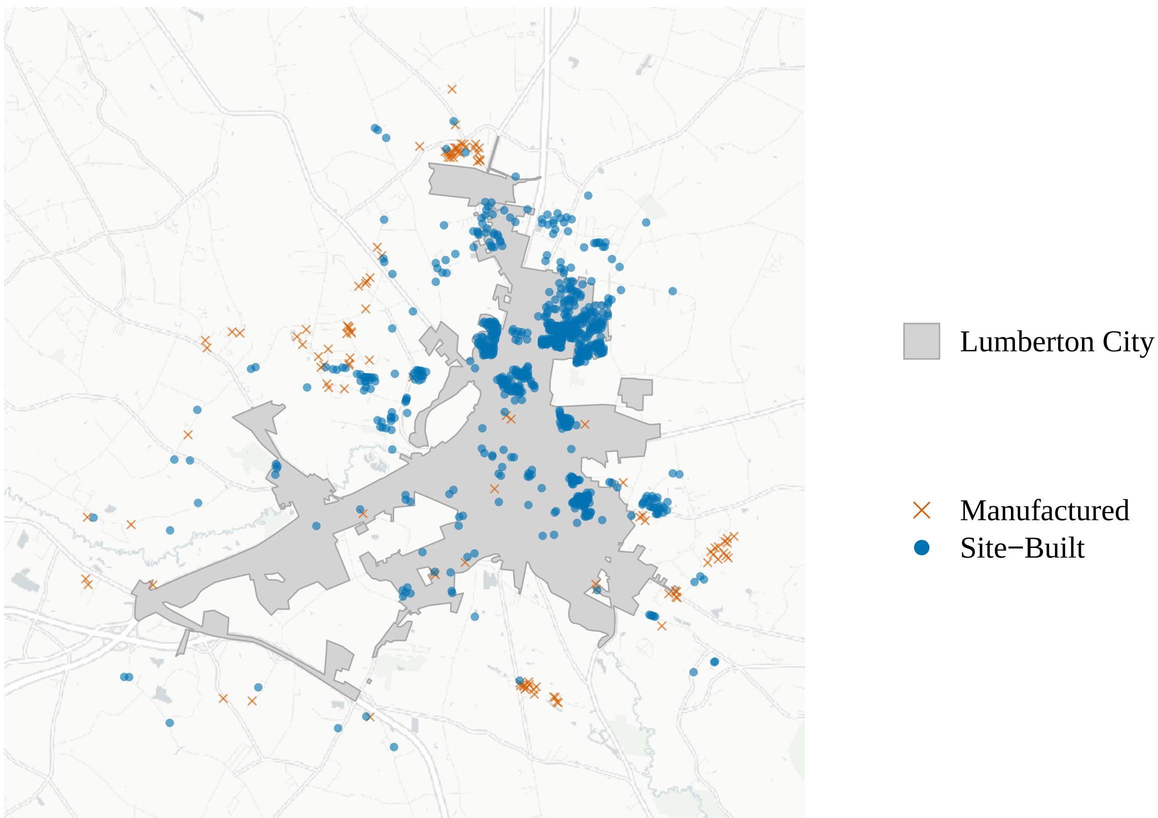

Notes: The map shows the location of manufactured and site-built single-family homes in Lumberton, North Carolina. The gray region indicates the area of the incorporated city. The figure depicts single-family homes built since 1980 within 2 km of the city boundary.

Table 2: RDD Estimates on Manufactured Home Shares (Logit)

Model:	(1)	(2)
Variables		
Incorp. ($\beta^{RDD,q}$)	-1.16*** (0.305)	-1.04*** (0.401)
Fixed-effects		
Place boundary	Yes	
Place boundary-Product segment		Yes
Fit statistics		
Observations	6,259,245	1,256,210
Squared Correlation	0.174	0.352
Pseudo R^2	0.292	0.357

Clustered (Place boundary) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Table reports logit regression results from Equation (9) using a sample of single-family homes built after 1980 within 2 km of a city line. Coefficients are reported on the log-odds scale, with cluster-robust standard errors in parentheses (clustered by boundary area). Each boundary area is defined as the nearest city interacted with the local school district so that all comparisons are made within these areas. Product segments are decile bins of living area and lot size interacted with construction decade (1980s–2010s). All specifications include a third-degree polynomial in distance to the boundary that is estimated separately for homes inside and outside the city.

Table 3: RDD Estimates on Sale Prices ($1000s)

Model:	(1)	(2)
Variables		
Incorp.	32.5* (19.3)	14.2 (8.8)
MH	-211.5*** (5.0)	-128.1*** (1.9)
Incorp. × MH	14.9** (7.1)	7.7*** (2.5)
Fixed-effects		
Place boundary	Yes	
Place boundary-Product segment		Yes
Fit statistics		
Observations	6,259,245	6,103,629
R^2	0.375	0.721
Within R^2	0.028	0.012
Dependent variable mean	428.8	431.7

Clustered (Place boundary) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Table reports regression results from Equation (10) using a sample of single-family homes built after 1980 within 2 km of a city line. Each boundary area is defined as the nearest city interacted with the local school district so that all comparisons are made within these areas. All specifications include a third-degree polynomial in distance to the boundary that is estimated separately for homes inside and outside the city. Coefficients are measured in $1000s of 2023 dollars. The smaller observation count in Column 2 is due to fixed-effect singleton observations that are dropped after adding boundary-area × product-segment fixed effects.

Table 4: Physical and Neighborhood Characteristics at City Boundaries

Dependent Variables:	Black	AGI (\$1000s)	Pop./km^2	Acres	Sq. ft.	Yr built
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Incorp.	-0.013	6.93	-189.9	0.119	104.4	0.668
	(0.011)	(7.79)	(130.7)	(0.076)	(68.6)	(1.05)
MH	-0.004	-41.3***	-111.7***	0.451***	-556.0***	-4.01***
	(0.002)	(1.92)	(28.2)	(0.039)	(13.5)	(0.306)
Incorp. × MH	0.006**	4.79*	64.1*	-0.324***	61.9***	-0.103
	(0.003)	(2.58)	(38.1)	(0.098)	(23.9)	(0.388)
Fixed-effects						
Place boundary	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	6,254,313	6,254,313	6,254,313	6,254,313	6,254,313	6,254,313
R^2	0.59969	0.39220	0.43836	0.17555	0.17160	0.18308
Dep. Var. Mean	0.088	150	1213	0.478	2081	2000

Clustered (Place boundary) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table reports results from RDD regressions of parcel characteristics at city-county boundaries using a sample of single-family homes built after 1980 within 2 km of a city line. All specifications include a third-degree polynomial in distance to the boundary that is estimated separately for homes inside and outside the city. Each *Place boundary* is defined as the interaction of the nearest city and the local school district, so all comparisons are made within these areas. *Black* is the fraction of the population that is Black in the parcel's grid cell; *AGI (\$1000s)* is mean adjusted gross income per tax filing unit, aggregated to the household rather than the individual level. Demographic variables are assigned from the Environmental Impact Frame (EIF), which reports population counts by demographic category at the 0.01-degree grid cell level using administrative tax and census records. Because each parcel inherits its grid-cell average, this spatial aggregation smooths demographics across the boundary and compresses the estimated demographic discontinuities toward zero. These models exclude a small number of parcels with invalid demographic data in the EIF due to statistical disclosure protections. Because *Acres* and *Sq. ft.* enter here as outcomes rather than as conditioning variables, both are winsorized at the 1st and 99th percentiles—the same thresholds applied to sale prices throughout—so that a few extreme parcels do not dominate these level regressions.

Table 5: Structural Objects: Bootstrap Standard Errors and Confidence Intervals

	Point estimate	Bootstrap SE	95% CI
Reduced-form discontinuities			
Share discontinuity, log-odds ($\hat{\beta}^{RDD}$)	-1.04	0.39	[-1.79, -0.24]
Price discontinuity, \$1000s ($\hat{\beta}^{RDD\text{-}\$}$)	7.69	2.51	[2.67, 12.37]
Common price jump, \$1000s ($\hat{\gamma}_1$)	14.18	8.79	[-3.46, 32.14]
Structural objects			
Ad-valorem implicit tax (τ)	0.059	0.020	[0.021, 0.097]
Elasticity of substitution (σ)	17.6	25.2	[3.6, 60.8]
Counterfactual MH share ($\tilde{s}_M$, %)	5.8	2.3	[2.6, 11.7]
Reduction in MH share (%)	63.5	15.6	[21.0, 81.3]
Counterfactual MH expenditure share (ω_M, %)	3.4	1.4	[1.5, 7.1]
Welfare loss (% of expenditure)	0.13	0.05	[0.04, 0.24]

Note: Table reports point estimates, bootstrap standard errors, and 95% percentile confidence intervals for the reduced-form boundary discontinuities and the structural objects derived from them. Sampling distributions come from a cluster bootstrap over boundary areas (999 draws): each draw resamples boundary areas with replacement, re-estimates the share RDD (Equation (9)) and the price RDD (Equation (10)) jointly, and recomputes every object, so the intervals account for the covariance among the discontinuities. The price discontinuity $\hat{\beta}^{RDD\text{-}\$}$ and the common jump $\hat{\gamma}_1$ are measured in \$1000s of 2023 dollars. The ad-valorem tax τ maps the wedge and the common site-built jump γ_1 into the log double difference via Equation (5), normalizing on the unincorporated cell prices; the counterfactual share $\tilde{s}_M$ inverts the log-odds identity at the inside-city MH share; and the welfare loss applies the exact CES expression in Equation (8). All draws produce a positive implicit tax τ.

Table 6: Robustness of Share RDD Estimates

Robustness dimension	Baseline	Bandwidth		Polynomial	
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Incorp.	-1.04***	0.329	-0.361	-1.23***	-1.27***
	(0.401)	(0.618)	(0.447)	(0.090)	(0.159)
Fixed-effects					
Place boundary-Product segment	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Bandwidth	2 km	0.5 km	1 km	2 km	2 km
Observations	1,256,210	410,883	760,172	1,256,210	1,256,210
Polynomial	Cubic	Cubic	Cubic	Linear	Quadratic

Clustered (Place boundary) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Boundary-clustered standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications use a logit model; coefficients are reported on the log-odds scale. All specifications include boundary $\times$ product-segment fixed effects. The row labeled “Bandwidth” reports the symmetric window around the boundary used in each column. The row labeled “Polynomial” reports the order of the global polynomial in distance interacted with incorporation status.

Table 7: Robustness of Price RDD Estimates

Robustness dimension	Baseline	Bandwidth		Polynomial	
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Incorp. × MH	7.7***	11.6***	9.6***	7.7***	7.7***
	(2.5)	(2.8)	(2.7)	(2.5)	(2.5)
Incorp.	14.2	9.1	4.6	7.0***	8.4***
	(8.8)	(11.5)	(9.0)	(2.0)	(2.5)
MH	-128.1***	-128.4***	-128.1***	-128.1***	-128.1***
	(1.9)	(2.5)	(2.2)	(1.9)	(1.9)
Fixed-effects					
Place boundary-Product segment	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Bandwidth	2 km	0.5 km	1 km	2 km	2 km
Observations	6,103,629	2,637,376	4,297,173	6,103,629	6,103,629
R^2	0.721	0.727	0.728	0.721	0.721
Within R^2	0.012	0.012	0.012	0.012	0.012
Polynomial	Cubic	Cubic	Cubic	Linear	Quadratic

Clustered (Place boundary) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Boundary-clustered standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include boundary × product-segment fixed effects. The row labeled "Bandwidth" reports the symmetric window around the boundary used in each column. The row labeled "Polynomial" reports the order of the global polynomial in distance interacted with incorporation status. Coefficients are measured in $1000s of 2023 dollars.

Table 8: Local Polynomial RDD Estimates on MH Shares

	(1)	(2)	(3)
Incorp.	-0.0475^{***}	-0.0137^{***}	-0.0143^{***}
	(0.0048)	(0.0016)	(0.0015)
Bandwidth h (km)	0.33	0.24	2.00
Bias bandwidth b (km)	0.74	0.51	2.00
Eff. observations	2,402,077	1,864,019	6,259,245
Boundary × Segment FE		Yes	Yes
Optimal bandwidth	Yes	Yes	

Note: Table presents local polynomial RDD estimates of the incorporation effect on manufactured home shares using `rdrobust` (Calonico et al., 2015). All specifications use a local linear polynomial with a triangular kernel and report conventional point estimates with robust bias-corrected standard errors in parentheses. Columns (1)–(3) use MSE-optimal bandwidth selection; Column (4) uses a fixed 2 km bandwidth. Fixed effects and demographic controls are implemented via residualization. Standard errors are clustered at the boundary area level.

Table 9: Placebo Effects of Census Designated Places on MH Shares

Model:	(1)
Variables	
CDP	-0.012 (0.027)
Fixed-effects	
CDP boundary-Product segment	Yes
Fit statistics	
Observations	2,140,859
R^2	0.601
Within R^2	0.0004
Dependent variable mean	0.072

Clustered (CDP boundary) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Table presents regression results from Equation (9) using a sample of single-family homes built after 1980 within 2km of a Census Designated Place (CDP) boundary. CDPs are unincorporated statistical areas defined by the Census Bureau and lack formal zoning authority. Each *CDP boundary* is defined as the interaction of the nearest CDP and the local school district, so all comparisons are made within these areas. All specifications include a third-degree polynomial in distance to the boundary that is estimated separately for homes inside and outside the CDP.

Table 10: Placebo Effects of Census Designated Places on Sale Prices ($1000s)

Model:	(1)
Variables	
CDP	31.8** (15.9)
MH	-123.9*** (2.4)
CDP × MH	4.2 (2.9)
Fixed-effects	
CDP boundary-Product segment	Yes
Fit statistics	
Observations	2,140,859
R^2	0.735
Within R^2	0.017

Clustered (CDP boundary) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Table reports regression results from Equation (10) using a sample of single-family homes built after 1980 within 2 km of a CDP boundary. CDPs are unincorporated statistical areas defined by the Census Bureau and lack formal zoning authority. All specifications include a third-degree polynomial in distance to the boundary that is estimated separately for homes inside and outside the CDP. Coefficients are measured in $1000s of 2023 dollars. Standard errors are clustered at the boundary area level.

Table 11: RDD Effects on Relative Assessed Values ($1000s)

Model:	(1)	(2)
Variables		
Incorp.	47.7** (18.9)	13.6 (9.5)
MH	-181.9*** (4.7)	-108.4*** (2.0)
Incorp. × MH	10.1 (6.3)	0.3 (2.7)
Fixed-effects		
Place boundary	Yes	
Place boundary-Product segment		Yes
Fit statistics		
Observations	6,253,478	6,097,851
R^2	0.355	0.696
Within R^2	0.015	0.007

Clustered (Place boundary) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Table reports estimates of Equation (10) using assessor-reported values rather than transaction sale prices. The sample includes all single-family parcels within 2 km of a city boundary and is not restricted to transacting parcels. All specifications include a third-degree polynomial in distance to the boundary estimated separately for homes inside and outside the city. Coefficients are measured in $1000s of 2023 dollars. Standard errors are clustered at the boundary area level. The smaller observation count in Column 2 is due to fixed-effect singleton observations that are dropped after adding boundary-area × product-segment fixed effects.

Table 12: Heterogeneity in Boundary Price Effects

Dependent Variable:	Boundary price effect ($1000s)		
Model:	(1)	(2)	(3)
Variables			
Log city land area	-6.07*	-6.24**	-4.18**
	(3.26)	(2.45)	(2.01)
State MH protections		-19.6***	
		(5.94)	
State WRLURI		9.92*	
		(5.59)	
Dillon's Rule state		-21.8***	
		(7.41)	
Fixed-effects			
State FE			Yes
Fit statistics			
Observations	2,627	2,619	2,626
R^2	0.01224	0.04549	0.08824
Within R^2			0.00457
Dep. Var. mean (weighted)	10.173	10.226	10.214

Clustered (State FE) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Table shows regressions of the boundary-specific RDD estimate of the incorporation effect on the MH relative price gap on characteristics of the boundary area. The dependent variable is measured in $1000s of 2023 dollars. Boundary-specific estimates use a linear distance control and product-segment fixed effects within boundary area. Observations are weighted by the number of parcels in the boundary area and standard errors are clustered by state. "State MH protections" is an indicator for states with a score above 2 on the Dawkins (2011) index of MH protections in state statute, which measures whether states preempt local bans, mandate installation standards, or otherwise limit local authority to restrict MH. "State WRLURI" is the parcel-weighted mean of the Wharton Residential Land Use Regulatory Index (Gyourko et al., 2008), a survey-based measure of the overall stringency of local land use regulation. "Dillon's Rule state" is an indicator for states that apply Dillon's Rule to municipalities, limiting the scope of their legal authority, based on the classification in Richardson et al. (2003).

Figure 4: Boundary Price Effects by State

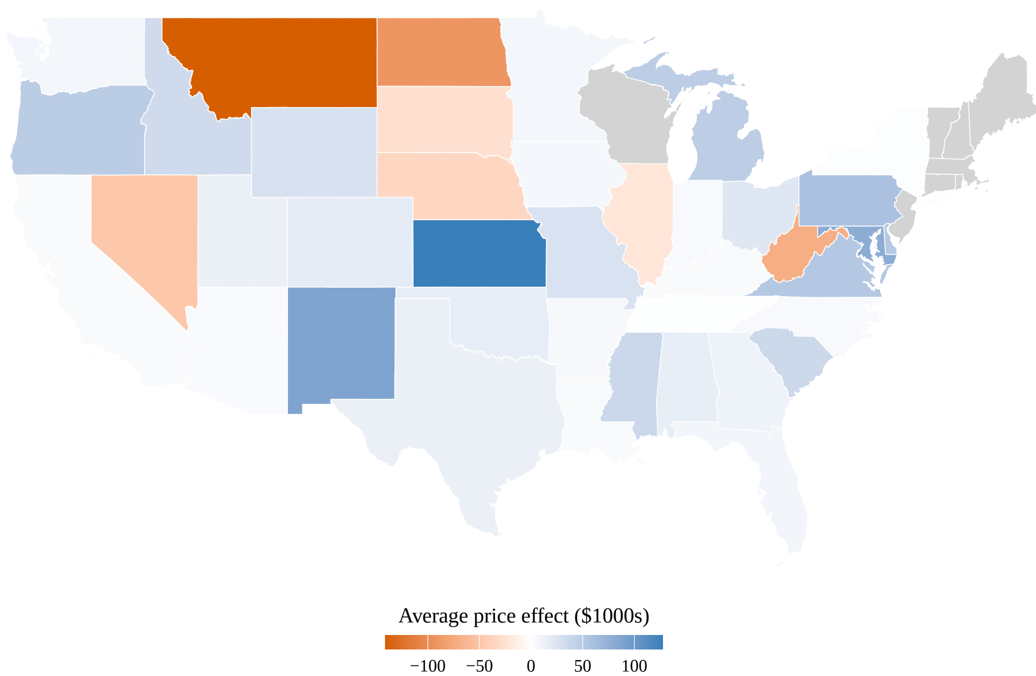

Notes: The map shows state-level weighted averages of boundary-specific estimates of the city effect on the MH relative price gap. Boundary-specific estimates use a linear distance control and product-segment fixed effects within boundary area. State averages use the same outlier-trimmed sample as Table 12 and weight boundary areas by parcel counts.

A Nested Demand System

The demand system in Section 3.1 is the lower tier of a two-level problem in which households first choose a side of the boundary and then a housing type. Let total housing consumption at distance d from the boundary be the CES aggregate across the two sides,

$$Q(d) = \left[\beta_I(d)^{1/\eta} Q_I(d)^{\frac{\eta-1}{\eta}} + \beta_U(d)^{1/\eta} Q_U(d)^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}},$$

where $\eta > 0$ is the elasticity of substitution across the two sides of the boundary, $\beta_\ell(d)$ is a preference weight on side ℓ, and $Q_\ell(d)$ is the lower-tier CES aggregate over MH and site-built housing defined in Section 3.1.

The problem separates into two stages. Conditional on the expenditure allocated to side ℓ, the household chooses its MH–site-built mix by minimizing lower-tier cost, which yields the conditional relative demand in Equation (1). The upper tier then allocates expenditure across the two sides according to the CES side-level price indices $P_\ell(d)$ implied by $Q_\ell(d)$. Because the mix chosen within a side depends only on the within-side relative price $p_{M,\ell}(d)/p_{S,\ell}(d)$ and the within-side taste ratio $\alpha_{M,\ell}(d)/\alpha_{S,\ell}(d)$, and not on the quantities consumed on the opposite side, the within-side relative demand for MH is invariant to the upper-tier allocation. This weak-separability property is the only feature of the nested structure used in identification (Section 3.2): regulation may reallocate households across the boundary through $\beta_\ell(d)$ and the side-level price indices, but it cannot alter the within-side MH–site-built ratio except through relative prices.

B From RDD Estimates to Model Objects

The equilibrium condition (Equation (7)) relates a relative log quantity to a log relative price. The RDD instead estimates a discontinuity in MH shares and two price discontinuities in dollars. This appendix records the two mappings between the estimated and model objects: the conversion of the level price discontinuities into the ad-valorem wedge τ, and the equivalence between the share discontinuity and the relative log quantity.

B.1 The ad-valorem wedge

The price RDD (Equation (10)) returns two dollar coefficients: $\beta^{RDD,p}$, which estimates the differential wedge t, and the coefficient on incorporated status γ_1, which estimates the common site-built jump at the boundary,

$$b \equiv \lim_{d\to 0} \Delta_\ell\, p_{S,\ell}(d).$$

In terms of the supply decomposition (Equation (2)), b collects the discontinuity in per-acre land rents at the city line and any difference in the site-built regulatory wedge across sides. Evaluating Equation (4) at the boundary gives $\lim_{d\to 0} \Delta_\ell\, p_{M,\ell}(d) - \lim_{d\to 0} \Delta_\ell\, p_{S,\ell}(d) = t$, so the manufactured-home level jump is

$$\lim_{d\to 0} \Delta_\ell\, p_{M,\ell}(d) = b + t.$$

Let $p_{h,U}$ denote the limiting price of a type-h home just outside the boundary, so that the corresponding inside prices are $p_{S,U} + b$ and $p_{M,U} + b + t$. The log boundary difference for each housing type is

$$\lim_{d\to 0} \Delta_\ell \ln p_{S,\ell}(d) = \ln\left(\frac{p_{S,U}+b}{p_{S,U}}\right) = \ln\left(1+\frac{b}{p_{S,U}}\right),$$
$$\lim_{d\to 0} \Delta_\ell \ln p_{M,\ell}(d) = \ln\left(\frac{p_{M,U}+b+t}{p_{M,U}}\right) = \ln\left(1+\frac{b+t}{p_{M,U}}\right).$$

Differencing across types yields the log double difference

$$\tau = \lim_{d\to 0} \Delta_\ell \Delta_h \ln p_{h,\ell}(d) = \ln\left(1+\frac{b+t}{p_{M,U}}\right) - \ln\left(1+\frac{b}{p_{S,U}}\right),$$

which is Equation (5). Applying $\ln(1+x) \approx x$ to each term gives the first-order form

$$\tau \approx \frac{b+t}{p_{M,U}} - \frac{b}{p_{S,U}} = \frac{t}{p_{M,U}} + b\left(\frac{1}{p_{M,U}} - \frac{1}{p_{S,U}}\right).$$

Two features of the mapping follow directly. First, the wedge is normalized by the manufactured-home price: t is a dollar cost borne by manufactured homes, and its ad-valorem equivalent is that cost as a fraction of the taxed good's price. A base that averaged over housing types would divide a cost borne by manufactured homes partly by the price of site-built homes and would understate the tax rate any household faces. Second, the common jump b enters τ even though it cancels in the level double difference: in logs, each dollar change is divided by the price of the good it applies to, and the two b terms offset only when $p_{M,U} = p_{S,U}$. Because manufactured homes are the cheaper type, a positive b raises τ. Neither feature affects the dollar wedge itself, which Equation (4) recovers at every distance regardless of b.

The sample counterpart replaces t and b with $\hat{\beta}^{RDD,p}$ and $\hat{\gamma}_1$ and the prices $p_{h,U}$ with the mean price of each type just outside the boundary. Section 4.5 reports the resulting $\widehat{\tau}$ and its sensitivity to the treatment of b (including the $b = 0$ special case $\tau = \ln(1 + t/p_{M,U})$) and to the choice of base. The cluster bootstrap resamples the two coefficients jointly, so sampling uncertainty in $\hat{\gamma}_1$ propagates to $\widehat{\tau}$ and to the objects downstream of it.

B.2 Shares and relative quantities

The share RDD estimates a discontinuity in the composition of the housing stock, while the left-hand side of Equation (7) is a relative log quantity. With two goods the two objects are linked by an identity. Define the MH share

$$s_{M,\ell}(d) \equiv \frac{q_{M,\ell}(d)}{q_{M,\ell}(d) + q_{S,\ell}(d)}.$$

Solving for the relative quantity,

$$\frac{q_{M,\ell}(d)}{q_{S,\ell}(d)} = \frac{s_{M,\ell}(d)}{1 - s_{M,\ell}(d)}, \tag{12}$$

so the relative log quantity equals the log-odds of the MH share. Taking the boundary double difference,

$$\lim_{d\to 0} \Delta_\ell \Delta_h \ln q_{h,\ell}(d) = \lim_{d\to 0} \Delta_\ell \ln\left(\frac{s_{M,\ell}(d)}{1 - s_{M,\ell}(d)}\right) \equiv \beta^{RDD,q}, \tag{13}$$

where $\beta^{RDD,q}$ is the discontinuity in the log-odds of the MH share at the boundary. A logit RDD on the MH indicator (Equation (9)) estimates $\beta^{RDD,q}$ directly, with no approximation. A linear-probability share RDD recovers the same object up to two first-order approximations; that mapping is reported in Appendix E.

C Proof of Proposition 1

Consider a household with expenditure E choosing between manufactured and site-built housing within product segment j. The household solves

$$\max_{q_M, q_S} \left[\alpha_M^{1/\sigma} q_M^{\frac{\sigma-1}{\sigma}} + \alpha_S^{1/\sigma} q_S^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad \text{s.t.} \quad p_M q_M + p_S q_S = E.$$

The indirect utility function is $V = E/P$, where the CES price index is

$$P = \left[\alpha_M p_M^{1-\sigma} + \alpha_S p_S^{1-\sigma} \right]^{\frac{1}{1-\sigma}}.$$

The expenditure function is $E(P, u) = P \cdot u$. The compensating variation of a price change from $\mathbf{p}^0$ (no tax) to $\mathbf{p}^1$ (with the implicit tax on MH) is

$$CV = E(\mathbf{p}^1, u^0) - E(\mathbf{p}^0, u^0) = \left(P^1 - P^0 \right) u^0,$$

so as a fraction of baseline expenditure $E = P^0 u^0$,

$$\frac{CV}{E} = \frac{P^1}{P^0} - 1.$$

Under the implicit tax, $p_M^1 = e^{\tau} p_M^0$ while $p_S^1 = p_S^0$, where τ is the ad-valorem wedge of Section 3.3. The ratio of price indices is therefore

$$\frac{P^1}{P^0} = \left[\frac{\alpha_M (e^{\tau} p_M)^{1-\sigma} + \alpha_S p_S^{1-\sigma}}{\alpha_M p_M^{1-\sigma} + \alpha_S p_S^{1-\sigma}} \right]^{\frac{1}{1-\sigma}}.$$

Dividing numerator and denominator inside the brackets by $\alpha_M p_M^{1-\sigma} + \alpha_S p_S^{1-\sigma}$ and using the definition of the counterfactual MH expenditure share,

$$\omega_M \equiv \frac{\alpha_M p_M^{1-\sigma}}{\alpha_M p_M^{1-\sigma} + \alpha_S p_S^{1-\sigma}},$$

the ratio simplifies to

$$\frac{P^1}{P^0} = \left[\omega_M (e^{\tau})^{1-\sigma} + (1 - \omega_M) \right]^{\frac{1}{1-\sigma}} = \left[\left((e^{\tau})^{1-\sigma} - 1 \right) \omega_M + 1 \right]^{\frac{1}{1-\sigma}}.$$

Subtracting one gives the result in Equation (8). $\square$

D Sample Construction and Restriction Sensitivity

This appendix documents how the sample is built from the CoreLogic parcel universe and how the reduced-form and welfare results respond to the sample restrictions described in Section 2.2.

D.1 Attrition and composition

Table 13 traces the parcel count through each restriction in the order I apply it, reporting the share of parcels retained at each step and the manufactured-home share among the survivors. Table 14 reports the composition of the final sample by state.

Two restrictions do most of the work. Dropping parcels with no reported lot size or living area removes the bulk of unidentified trailer-park units, which rarely report a lot size: 39% of these parcels are manufactured homes, against 6% in the retained sample. The restriction thus functions primarily as a data-cleaning step that removes trailer park units rather than as a neutral trim. Requiring at least one manufactured home on both sides of each boundary removes boundary areas with no manufactured housing to compare; a subset of these — 1,209 boundary areas — have manufactured homes just outside the city but none inside, and so are candidates for de facto exclusion that the intensive-margin design cannot speak to.

Table 13: Sample Construction Waterfall

Restriction	Parcels	Retained (%)	MH share (%)
Exported (1970+ vintages)	48,000,694	100.0	6.5
Year built $\geq$ 1980	38,367,447	79.9	6.7
Lot size and living area reported	37,684,213	78.5	6.1
Exclude New England	36,635,525	76.3	6.2
Valid geographic identifiers	34,452,511	71.8	6.4
Within 2 km of a city boundary	21,253,773	44.3	4.0
Arm's-length sale price observed	13,660,886	28.5	2.3
$\geq$ 50 parcels, both sides populated	11,466,914	23.9	2.4
$\geq$ 1 MH on both sides (final)	**6,259,245**	**13.0**	**3.9**

Note: Each row applies the restrictions in all rows above it. "Retained" is the share of the exported parcels (single-family homes built since 1970) that survive through the row; "MH share" is the manufactured-home share among survivors. The final row is the analysis sample.

D.2 Sensitivity to the outlier restrictions

The baseline winsorizes sale price at the 1st/99th percentiles within housing type and drops parcels with no reported lot size or living area, but does not trim the lot size or living area tails. Table 15 re-estimates the share and price discontinuities under an alternative treatment of this outlier rule, alongside variations in the vintage window and three restrictions that change which parcels, boundary areas, and product segments enter the sample rather than how outliers are handled.

Two patterns stand out. The share discontinuity ($\beta^{RDD,q}$) and the log-price discontinuity ($\beta^{RDD,\ln p}$) are stable whether I winsorize the price tails (the baseline) or leave them raw ("no price winsorize''). The level price discontinuity ($\beta^{RDD,p}$) is where the outlier rule bites: without

Table 14: Final Sample Composition by State

State	Parcels	Share of sample (%)	Retained (%)	MH share (%)
TX	959,213	15.3	20.3	2.6
FL	789,305	12.6	19.2	7.2
NC	723,393	11.6	35.3	3.7
CA	511,025	8.2	18.4	1.9
AZ	427,266	6.8	28.1	5.2
WA	340,363	5.4	31.7	4.8
CO	332,783	5.3	33.4	2.5
NV	295,470	4.7	41.9	1.4
TN	214,422	3.4	17.7	5.9
GA	200,899	3.2	9.7	3.3
SC	185,036	3.0	19.1	2.5
OR	179,614	2.9	34.0	6.1
OK	133,234	2.1	23.6	4.1
AR	106,447	1.7	23.2	3.5
IN	95,644	1.5	11.6	4.3

Note: States with the largest number of parcels in the final analysis sample. "Retained" is the share of the state's exported parcels (single-family homes built since 1980) that appear in the final sample.

winsorizing the extreme sales, its standard error nearly doubles and the discontinuity is no longer statistically distinguishable from zero. Because the welfare calculation rests on the level price wedge, winsorizing the price tails is what delivers a precise estimate of the regulatory cost.

The *All parcels* row relaxes a different restriction: it drops the requirement of an observed market transaction, so the share RDD runs on the full parcel universe within the baseline boundary areas rather than on transacting parcels alone. The price discontinuities are undefined without a sale price and are reported as "—' '. This row is the direct check on whether selection into transacting drives the share result.

The last two rows change the sample rather than the outlier rule. Like every other row, both start from the baseline outlier and vintage treatment; each then relaxes a single boundary-area restriction. The baseline keeps only boundary areas with at least one manufactured-home sale on *both* the incorporated and unincorporated side (the "both-sides MH" requirement), and, within those boundary areas, all product segments (decile bins of lot size and living area, interacted with decade built). *Common support* narrows the baseline further, dropping any boundary-area-by-product-segment cell that does not contain both manufactured and site-built homes; this is a materially smaller sample built around a different estimand (the discontinuity comparing housing types only where both are actually sold side by side), and its larger share discontinuity raises the welfare estimate accordingly. *Keep cities w/ no MH* goes the other way: it drops the both-sides MH requirement entirely, so boundary areas that never transact a manufactured home on one or both sides re-enter the sample. Both rows are reported for transparency about how the boundary-selection rule shapes the estimates, not as stability checks on the baseline.

Table 15: Reduced-Form Estimates across Sample Restrictions

Sample	N	$\beta^{RDD,q}$	$\beta^{RDD,p}$	$\beta^{RDD,\ln p}$
Baseline	**6,259,245**	**-0.024** (0.011)**	**7.69*** (2.48)**	**0.076*** (0.010)**
No price winsorize	6,259,245	-0.024** (0.011)	4.26 (4.55)	0.071*** (0.010)
Year built $\geq$ 1970	7,731,447	-0.018* (0.011)	6.05** (2.45)	0.066*** (0.010)
Year built $\geq$ 1990	4,763,772	-0.027*** (0.009)	9.88*** (2.85)	0.077*** (0.011)
All parcels	9,673,347	-0.022** (0.010)	—	—
Common support	1,256,210	-0.094*** (0.036)	7.89*** (2.22)	0.073*** (0.011)
Keep cities w/ no MH	11,466,914	-0.017*** (0.006)	4.96** (2.27)	0.075*** (0.010)

Note: Each row re-estimates the baseline share and price RDDs (Equations (9) and (10)) on a different sample. $\beta^{RDD,q}$ is the share discontinuity; $\beta^{RDD,p}$ is the boundary price discontinuity in levels ($1000s) and $\beta^{RDD,\ln p}$ in logs. *Baseline* is the sample used throughout the paper. All specifications include boundary area $\times$ product-segment fixed effects and a third-degree polynomial in distance to the boundary estimated separately for homes inside and outside the city. Product segments are decile bins of lot size and living area interacted with construction decade; the decile cuts are recomputed on each row's own sample, so a row reports the specification the paper would have used had that restriction been adopted at the outset. *Common support* is the exception: because the restriction is defined over product segments, that row inherits the baseline cuts. *All parcels* drops the requirement of an observed market transaction, so the price discontinuities are undefined and reported as "—"; it is restricted to the boundary areas of the baseline sample. Standard errors clustered on boundary area in parentheses; stars from $|t|$ (* 0.10, ** 0.05, *** 0.01).

E Linear Probability Specification

The preferred share estimate in Section 4.2 uses a logit RDD, which by Equation (13) maps to the relative log quantity object in the model without approximation. As a check, I also estimate the share discontinuity by OLS (a linear probability model) and map the resulting estimate to the model object via two first-order approximations.

The linear-probability analog of Equation (9) is

$$\mathrm{MH}_{ib} = \beta^{RDD,\mathrm{LPM}} \cdot \mathrm{Incorp}_i + f(\mathrm{dist}_i) \times \mathrm{Incorp}_i + \alpha_{j(i),b} + \epsilon_{ib}, \tag{14}$$

where all variables are as in Equation (9) and $\beta^{RDD,\mathrm{LPM}}$ is the discontinuity in the MH share at the boundary, measured in share units. Results are reported in Table 16. In the preferred segment $\times$ boundary specification, the estimated discontinuity is $\hat{\beta}^{RDD,\mathrm{LPM}} \approx -0.024$, i.e. a 2.4 percentage point decrease in MH shares at the boundary. Relative to the average inside-city share of 2.1%, this is a 113% reduction in MH.

To convert $\hat{\beta}^{RDD,\mathrm{LPM}}$ to the relative log quantity object on the left-hand side of Equation (7), two first-order approximations are required in addition to the share-to-quantity identity in Equation (12). For small shares of MH, the log-odds ratio approximates the log share,

$$\ln\left(\frac{s_{M,\ell}(d)}{1 - s_{M,\ell}(d)}\right) \approx \ln s_{M,\ell}(d), \qquad s_{M,\ell}(d) \ll 1, \tag{15}$$

and for small changes in shares, the change in log shares approximates the percentage change,

$$\Delta_\ell \ln s_{M,\ell}(d) \approx \frac{\Delta_\ell s_{M,\ell}(d)}{\bar{s}_M}, \tag{16}$$

where $\bar{s}_M$ is the average MH share on the incorporated side. Putting these together,

$$\begin{aligned} \lim_{d\to 0} \Delta_\ell \Delta_h \ln q_{h,\ell}(d) &= \lim_{d\to 0} \Delta_\ell \underbrace{\ln\left(\frac{s_{M,\ell}(d)}{1 - s_{M,\ell}(d)}\right)}_{(12)} \approx \lim_{d\to 0} \Delta_\ell \underbrace{\ln s_{M,\ell}(d)}_{(15)} \\ &\approx \underbrace{\frac{\lim_{d\to 0} \Delta_\ell s_{M,\ell}(d)}{\bar{s}_M}}_{(16)} = \frac{\hat{\beta}^{RDD,\mathrm{LPM}}}{\bar{s}_M}. \end{aligned} \tag{17}$$

Combining with the equilibrium condition and the level-mapped price effect yields the LPM-based elasticity

$$\sigma^{\mathrm{LPM}} = -\frac{\hat{\beta}^{RDD,\mathrm{LPM}}/\bar{s}_M}{\hat{\tau}} = -\frac{-0.024/0.021}{0.059} \approx 19,$$

close to the preferred logit-based estimate of 17.6.

F Heterogeneity in Share Effects

For completeness, Table 17 reports the analogue of Table 12 using the boundary-level normalized share effect $\hat{\beta}_b^{RDD,q}/\bar{s}_{M,b}$ as the outcome instead of the dollar price effect. As discussed in Section 6, the share effect is downstream of the implicit tax — it equals $-\sigma \cdot \tau$ in the model, so it mixes

Table 16: Linear Probability RDD Estimates on Manufactured Home Shares

Model:	(1)	(2)
Variables		
Incorp. ($\beta^{RDD,q}$)	-0.045*** (0.015)	-0.024** (0.011)
Fixed-effects		
Place boundary	Yes	
Place boundary-Product segment		Yes
Fit statistics		
Observations	6,259,245	6,103,629
R^2	0.163	0.487
Within R^2	0.012	0.004
Dependent variable mean	0.039	0.036

Clustered (Place boundary) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Table reports linear-probability regression results from Equation (14) using a sample of single-family homes built after 1980 within 2 km of a city line. Coefficients are in percentage points of the MH share. Each boundary area is defined as the nearest city interacted with the local school district. All specifications include a third-degree polynomial in distance to the boundary estimated separately for homes inside and outside the city. Standard errors are clustered at the boundary area level. The smaller observation count in Column 2 is due to fixed-effect singleton observations that are dropped after adding boundary-area × product-segment fixed effects.

variation in the regulatory wedge with variation in the substitution elasticity, requires normalization by the boundary's MH intensity, and is empirically much noisier. The resulting patterns do not mirror the price-effect regressions. Since the average normalized share effect is negative, negative moderator coefficients imply larger share reductions in magnitude. The point estimates therefore indicate larger share effects in geographically larger cities, in states with stronger MH protections, and in Dillon's Rule states, while the positive WRLURI coefficient indicates smaller share effects in states with greater overall land-use regulation. None of these moderator coefficients is statistically distinguishable from zero at conventional levels.

G Lumberton, NC Zoning Regulations

Section 35-15. Definitions of Basic Terms

Manufactured home, Class A. A manufactured home that meets or exceeds the construction standards promulgated by the U.S. Department of Housing and Urban Development that were in effect at the time of construction and that satisfies the following additional criteria:

1. The pitch of the manufactured home's roof shall have a minimum vertical rise of one foot for each five feet of horizontal run;
2. The exterior materials shall be of wood, hardboard, or aluminum comparable in composition, appearance, and durability to site-built houses in the vicinity;
3. A continuous, permanent masonry foundation, unpierced except for required ventilation and access, shall be installed under the manufactured home (this criterion does not apply to manufactured homes located on land leased to the homeowner);
4. The tongue, axles, transporting lights, and removable towing apparatus shall be removed subsequent to final placement; and
5. The manufactured home shall have a length not exceeding four times its width.

Manufactured home, Class B. A manufactured home that meets or exceeds the construction standards promulgated by the U.S. Department of Housing and Urban Development that were in effect at the time of construction but that does not satisfy the criteria necessary to qualify the house as a Class A manufactured home.

Source: Municipal Code, Article II, Section 35-15. Ordinance history: Ord. No. 813 (June 3, 1985) through Ord. No. 2021.06.02 (June 9, 2021).

Section 35-135. Residential Districts Established

1. The following residential districts are hereby established: A, R-20, R-15, R-11, PR-11, R-7, R-6, and R-3. Each of these districts is designed and intended to secure for the persons who reside there a comfortable, healthy, safe, and pleasant environment in which to live, sheltered from incompatible and disruptive activities that properly belong in nonresidential districts. Other objectives of some of these districts are explained in the remainder of this section.

Table 17: Heterogeneity in Boundary Share Effects

Dependent Variable:	Normalized share effect		
Model:	(1)	(2)	(3)
Variables			
Log city land area	-0.157	-0.152	-0.038
	(0.125)	(0.123)	(0.117)
State MH protections		-0.203	
		(0.265)	
State WRLURI		0.091	
		(0.268)	
Dillon's Rule state		-0.405	
		(0.248)	
Fixed-effects			
State FE			Yes
Fit statistics			
Observations	2,627	2,619	2,626
R^2	0.00828	0.01528	0.08047
Within R^2			0.00038
Dep. Var. mean (weighted)	-1.7718	-1.7697	-1.7736

Clustered (State FE) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: Table shows regressions of the boundary-specific RDD estimate of the incorporation effect on the normalized MH share, $\hat{\beta}_b^{RDD,q}/\bar{s}_{M,b}$, on characteristics of the boundary area. Boundary-specific estimates use a linear distance control and product-segment fixed effects within boundary area. Observations are weighted by the number of parcels in the boundary area and standard errors are clustered by state. "State MH protections" is an indicator for states with a score above 2 on the Dawkins (2011) index of MH protections in state statute, which measures whether states preempt local bans, mandate installation standards, or otherwise limit local authority to restrict MH. "State WRLURI" is the parcel-weighted mean of the Wharton Residential Land Use Regulatory Index (Gyourko et al., 2008), a survey-based measure of the overall stringency of local land use regulation. "Dillon's Rule state" is an indicator for states that apply Dillon's Rule to municipalities, limiting the scope of their legal authority, based on the classification in Richardson et al. (2003).

2. The A (agricultural) district is designed to protect agricultural lands and woodlands within the city's planning jurisdiction; for this reason, larger minimum lot sizes are required. This district is also intended to accommodate some types of uses that would be appropriate in more sparsely populated areas but would not be appropriate in the more intensely developed residential zones.

3. The R-20, R-15, and R-11 districts differ primarily in the density allowed as determined by the minimum lot size requirements set forth in Section 35-181. In addition, some types of manufactured homes are allowed to be used for single-family residential purposes in the R-20 zone.

4. The PR-11 district is intended to encourage the development of residential neighborhoods composed of a well-planned mixture of single-family, two-family, and multi-family dwellings. Except for the fact that this district allows planned residential developments (as described in Section 35-157), it is identical to the R-11 district.

5. The R-7 district is designed to accommodate single-family and two-family dwelling units.

6. The R-6 district is designed to accommodate single-family and two-family dwelling units, as well as some types of manufactured homes used as single-family residences.

7. The R-3 zone is designed primarily to accommodate higher density multifamily developments.

Source: Municipal Code, Section 35-135. Ordinance history: Ord. No. 813 (June 3, 1985); Ord. No. 2021.06.02 (June 9, 2021).

References

Anagol, S., F. V. Ferreira, and J. M. Rexer (2021). Estimating the economic value of zoning reform. Technical report, National Bureau of Economic Research.

Anderson, S. P., A. D. Palma, and J.-F. Thisse (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press. Google-Books-ID: xiXGtz47p5oC.

Arnott, R. and G. Young (1979). The Property Tax As a Tax on Durability. *The Canadian Journal of Economics / Revue canadienne d'Economique 12*(3), 485–494.

Banzhaf, H. S. and K. Mangum (2019, March). Capitalization as a Two-Part Tariff: The Role of Zoning.

Baum-Snow, N. and G. Duranton (2025, April). Housing Supply and Housing Affordability.

Bayer, P., F. Ferreira, and R. McMillan (2007, August). A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy 115*(4), 588–638.

Black, S. E. (1999, May). Do Better Schools Matter? Parental Valuation of Elementary Education. *The Quarterly Journal of Economics 114*(2), 577–599.

Butts, K. (2023, March). Geographic difference-in-discontinuities. *Applied Economics Letters 30*(5), 615–619.

Calder-Wang, S. (2021). The distributional impact of the sharing economy on the housing market. *Available at SSRN 3908062*.

Calonico, S., M. Cattaneo, and R. Titiunik (2015, June). rdrobust: An R Package for Robust Nonparametric Inference in Regression-Discontinuity Designs. *The R Journal*.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica 82*(6), 2295–2326.

Cattaneo, M. D., R. K. Crump, M. H. Farrell, and Y. Feng (2024, May). On Binscatter. *American Economic Review 114*(5), 1488–1514.

Chiumenti, N., A. Kulka, and A. Sood (2022, June). How to Increase Housing Affordability: Understanding Local Deterrents to Building Multifamily Housing.

Davis, M. A., W. D. Larson, S. D. Oliner, and J. Shui (2021, March). The price of residential land for counties, ZIP codes, and census tracts in the United States. *Journal of Monetary Economics 118*, 413–431.

Dawkins, C. J. (2011, August). *Regulatory Barriers to Manufactured Housing Placement in Urban Communities*. DIANE Publishing. Google-Books-ID: 5Rb2MZYsbgQC.

Dawkins, C. J. (2025, June). Local land use regulations and new mobile home concentration. *Urban Studies*, 00420980251335564.

Favilukis, J. Y. and J. Song (2023). Why zoning is too restrictive. *Available at SSRN 4535417*.

Gann, D. M. (1996). Construction as a manufacturing process? Similarities and differences between industrialized housing and car production in Japan. *Construction Management & Economics 14*(5), 437–450.

Genz, R. (2001, January). Why advocates need to rethink manufactured housing. *Housing Policy Debate 12*(2), 393–414. _eprint: https://doi.org/10.1080/10511482.2001.9521411.

Glaeser, E. and J. Gyourko (2018, February). The Economic Implications of Housing Supply. *Journal of Economic Perspectives 32*(1), 3–30.

Glaeser, E. L. and J. Gyourko (2002). The impact of building restrictions on housing affordability. *Federal Reserve Bank of New York, Economic Policy Review 2002*, 1–19.

Glaeser, E. L. and J. Gyourko (2025, May). America's Housing Supply Problem: The Closing of the Suburban Frontier?

Goolsbee, A. and C. Syverson (2023, January). The Strange and Awful Path of Productivity in the U.S. Construction Sector.

Grembi, V., T. Nannicini, and U. Troiano (2016). Do Fiscal Rules Matter? *American Economic Journal: Applied Economics 8*(3), 1–30.

Gyourko, J. and J. Krimmel (2021, November). The impact of local residential land use restrictions on land values across and within single family housing markets. *Journal of Urban Economics 126*, 103374.

Gyourko, J. and S. E. McCulloch (2024, October). The Distaste for Housing Density.

Gyourko, J. and R. Molloy (2015, January). Chapter 19 - Regulation and Housing Supply. In G. Duranton, J. V. Henderson, and W. C. Strange (Eds.), *Handbook of Regional and Urban Economics*, Volume 5 of *Handbook of Regional and Urban Economics*, pp. 1289–1337. Elsevier.

Gyourko, J., A. Saiz, and A. Summers (2008). A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index. *Urban Studies 45*(3), 693–729. _eprint: https://doi.org/10.1177/0042098007087341.

Hamilton, B. W. (1975). Zoning and property taxation in a system of local governments. *Urban studies 12*(2), 205–211. Publisher: Sage Publications Sage UK: London, England.

Hankinson, M. (2018, August). When Do Renters Behave Like Homeowners? High Rent, Price Anxiety, and NIMBYism. *American Political Science Review 112*(3), 473–493.

Hankinson, M. and A. Magazinnik (2023, July). The Supply-Equity Trade-Off: The Effect of Spatial Representation on the Local Housing Supply. *The Journal of Politics 85*(3), 1033–1047.

Herbert, C., C. Reed, and J. Shen (2023). Comparison of the Costs of Manufactured and Site-Built Housing. Technical report, Joint Center for Housing Studies of Harvard University.

Kulka, A., A. Sood, and N. Chiumenti (2026, March). Under the (Neighbor)Hood: Understanding Interactions Among Zoning Regulations. *The Review of Economics and Statistics*, 1–45.

Ma, L. (2024). Build What and for Whom? The Distributional Effects of Housing Supply. *The Distributional Effects of Housing Supply (October 01, 2024).*

Maantay, J. (2001, July). Zoning, equity, and public health. *American Journal of Public Health 91*(7), 1033–1041.

Mast, E. (2024, May). Warding off Development: Local Control, Housing Supply, and NIMBYs. *The Review of Economics and Statistics 106*(3), 671–680.

Mei, M. (2022, October). House Size and Household Size: The Distributional Effects of the Minimum Lot Size Regulation.

Mimura, Y., K. Love-Myers, A. L. Sweaney, and M. Leigh (2013, January). Comparing Values of Manufactured and Equivalent Site-Built Homes in the Rural United States. *Housing and Society 40*(2), 150–169. _eprint: https://doi.org/10.1080/08882746.2013.11430614.

Munneke, H. J. and V. C. Slawson (1999). A housing price model with endogenous externality location: a study of mobile home parks. *The Journal of Real Estate Finance and Economics 19*(2), 113–131.

Reid Wilson (2025, March). What You Need to Know About Land-Use Restrictions in Texas.

Richardson, J. J., M. Z. Gough, and R. Puentes (2003). *Is home rule the answer?: clarifying the influence of Dillon's rule on growth management.* Center on Urban and Metropolitan Policy, the Brookings Institution ....

Rollet, V. and L. Weiwu (2025). Measuring Winners and Losers from Increasing Housing Supply.

Saiz, A. (2010). The Geographic Determinants of Housing Supply. *The Quarterly Journal of Economics 125*(3), 1253–1296.

Sutter, D. and M. Poitras (2010, April). Do people respond to low probability risks? Evidence from tornado risk and manufactured homes. *Journal of Risk and Uncertainty 40*(2), 181–196.

Voorheis, J. L., J. Colmer, K. Houghton, E. Lyubich, M. Munro, C. Scalera, and J. Withrow (2023). Building the Prototype Census Environmental Impacts Frame. Technical report, National Bureau of Economic Research.

Zhang, H. (2022). Consumer Cities: The Role of Housing Variety.

Zhang, P. and P. Nguyen-Hoang (2023). Home rule and municipal revenue stability: New evidence from Texas. *Public Budgeting & Finance 43*(1), 38–60. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/pbaf.12332.

Zhang, P., P. Nguyen-Hoang, and N. Chen (2022). The impact of home rule on municipal boundary and fiscal expansion: Evidence from Texas. *Journal of Regional Science 62*(5), 1442–1466. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/jors.12610.

Zhou, Y. (2013, September). The decision to purchase a manufactured home: a nested logit model of determinants. *International Journal of Housing Policy 13*(3), 268–287. _eprint: https://doi.org/10.1080/14616718.2013.818784.