

The Shelter Wedge: How Housing Supply Constraints Amplify the Great Decompression

TMC Research

February 2026

Abstract

Since 1970, shelter costs in the United States have grown 44% faster than overall prices, with the shelter CPI to overall CPI ratio rising from 0.897 to 1.291 by 2025. We introduce the concept of the *shelter wedge*—the mechanism by which housing supply constraints translate nominal income divergence into real consumption inequality—and test whether this wedge is amplified in supply-constrained markets. A structural two-neighborhood model predicts that inelastic housing supply magnifies the shelter wedge by a factor of 2.2 relative to elastic markets. However, difference-in-differences estimation exploiting cross-state variation in Saiz (2010) housing supply elasticities yields a precisely estimated null result: the interaction of supply inelasticity and national income divergence has no statistically significant effect on state-level rent burdens ($\beta = 0.011$, $p = 0.899$). Permutation inference confirms this null ($p = 0.868$). National local projections reveal that shelter cost sensitivity to income shocks has *weakened* since 2000, contradicting theoretical predictions. We argue that this informative null indicates the shelter wedge operates at a geographic granularity finer than states—likely at the metropolitan or neighborhood level—consistent with the spatial sorting mechanisms emphasized by Hsieh and Moretti (2019) and Diamond (2016). Our findings

have direct implications for the design of housing affordability policies and for future empirical research on the real consequences of income inequality.

JEL Classification: R21, R31, D31, J31

Keywords: housing supply, income inequality, shelter costs, Great Decompression, spatial sorting

1 Introduction

The United States has experienced a dramatic transformation in the structure of its income distribution over the past half century. Where the postwar era was characterized by what Goldin and Margo [1992] termed the “Great Compression”—a sustained convergence in wages across the skill distribution—the period since roughly 1980 has witnessed a sustained divergence. Upper-income households have pulled away from the median, a phenomenon that Goldin and Katz [2008] and Autor et al. [2008] have documented extensively. We refer to this widening income gap as the *Great Decompression*.

The consequences of the Great Decompression are typically measured in nominal terms: the ratio of 80th-percentile to 50th-percentile household income, the share of income accruing to the top decile [Piketty and Saez, 2003], or the evolution of skill premia [Autor et al., 2008]. Yet households consume goods and services, not nominal income. Housing occupies a singular position among these expenditures, absorbing 30–40% of income for median earners. If housing costs rise disproportionately for some households, even modest nominal income divergence may translate into large differences in material living standards.

This paper introduces the concept of the *shelter wedge*: the amplification of nominal income inequality through differential shelter cost inflation. We define the shelter wedge as the ratio of shelter cost growth to overall price growth, which has risen from 0.897 in 1970 to 1.291 in 2025—a 44% increase over five decades (Figure 1). The shelter wedge implies that housing has become systematically more expensive relative to other consumption goods, with outsized consequences for households in the middle and lower portions of the income distribution who devote larger income shares to rent and mortgage payments.

The theoretical mechanism linking supply constraints to the shelter wedge is straightforward. When housing supply is inelastic—due to geographic barriers, zoning restrictions, or regulatory constraints [Saiz, 2010, Glaeser et al., 2005, Gyourko et al., 2013]—increases in housing demand generated by rising incomes at the top of the distribution are capitalized into higher prices rather than absorbed by new construction. In markets where high-income households bid aggressively for desirable housing, supply constraints ensure that prices rise rather than quantities. This, in turn, raises shelter cost burdens for middle-income households competing for housing in the same local markets.

We formalize this intuition in a two-neighborhood spatial sorting model (Section 3), which yields a clear prediction: the shelter wedge multiplier—defined as the derivative of the shelter burden gap with respect to the income ratio—is decreasing in the housing supply elasticity η . Calibrated to match observed rent burdens in constrained ($\eta = 0.7$) and elastic ($\eta = 2.5$) markets, the model predicts that constrained markets amplify the shelter wedge

by a factor of 2.2 relative to elastic markets (Table 3).

The core empirical contribution of this paper is to test this theoretical prediction at the state level, exploiting cross-state variation in housing supply elasticity from Saiz [2010]. We construct a panel of 714 state-year observations spanning 51 states (including the District of Columbia) over the period 2010–2023, classifying 12 states as supply-constrained (Saiz elasticity < 1.2) and 23 as elastic (> 2.0). Our difference-in-differences specification interacts state-level supply inelasticity with the national 80th-to-50th income percentile ratio, identifying the shelter wedge amplifier as the coefficient on this interaction term.

Our central finding is a null result. Across all specifications, the estimated shelter wedge amplifier is small, imprecisely estimated, and statistically indistinguishable from zero. The baseline estimate is $\beta = 0.0107$ with a standard error of 0.0845 ($p = 0.899$). Adding economic controls yields $\beta = -0.0075$ ($p = 0.948$). A binary constrained-by-post-2015 specification produces $\beta = -0.0005$ ($p = 0.827$). Permutation inference, based on 1,000 random reassignments of treatment status across states, yields a p -value of 0.868, confirming the null. An event study design reveals no systematic pre-trends but also no post-treatment divergence. National local projections show that the shelter CPI response to income shocks has actually *weakened* since 2000, with the 12-month cumulative response falling from $\beta_{12} = 0.0182$ in the pre-2000 sample to $\beta_{12} = 0.0027$ in the post-2000 sample.

We argue that this null result is itself informative. The failure to detect a shelter wedge at the state level does not imply that the mechanism is absent; rather, it suggests that the relevant geographic unit of analysis is finer than the state. The spatial sorting processes that drive housing cost inflation—high-income households bidding up prices in desirable neighborhoods within metropolitan areas—are inherently local phenomena. Aggregation to the state level obscures within-state heterogeneity that may be substantial: California contains both San Francisco (among the most constrained housing markets in the nation) and the Inland Empire (with comparatively elastic supply). This interpretation is consistent with the metropolitan-level findings of Hsieh and Moretti [2019], who estimate that housing constraints in high-productivity cities reduced aggregate U.S. GDP growth by 36% over the period 1964–2009, and with Diamond [2016], who documents spatial sorting by skill within metropolitan areas.

The remainder of the paper proceeds as follows. Section 2 reviews the literature on the Great Decompression and housing markets. Section 3 presents the conceptual framework. Section 4 describes the data. Section 5 outlines the empirical strategy. Section 6 presents the main results. Section 7 discusses mechanisms and model estimation. Section 8 conducts robustness checks. Section 9 interprets the null result and discusses implications. Section 10 concludes.

2 The Great Decompression and Housing

2.1 The Great Decompression

The wage and income structure of the United States underwent a dramatic transformation in the final decades of the twentieth century. Goldin and Margo [1992] documented the “Great Compression” of wages during the 1940s and its persistence through the early 1970s—a period in which the wage distribution narrowed substantially, driven by institutional changes, wartime labor demand, and the expansion of secondary education. Beginning in the late 1970s, this compression reversed. Goldin and Katz [2008] argue that the reversal—what we term the Great Decompression—reflects a race between education and technology, in which the rate of skill-biased technical change outpaced the growth of educational attainment.

The empirical evidence for the Great Decompression is extensive. Piketty and Saez [2003] document the sharp increase in top income shares since the early 1980s. Autor et al. [2008] provide a detailed decomposition of wage inequality trends, emphasizing the role of upper-tail divergence. Moretti [2013] shows that much of the apparent increase in wage inequality is offset by spatial variation in the cost of living—particularly housing—suggesting that nominal income statistics may overstate real inequality in some contexts while understating it in others.

The social consequences of rising inequality extend beyond income itself. Reeves [2017] documents how the upper middle class has consolidated advantages in education, housing, and social networks. Putnam [2015] traces the growing divergence in childhood experiences by parental socioeconomic status, with housing quality and neighborhood characteristics playing central roles. Chetty et al. [2014] provide evidence that the geography of intergenerational mobility is closely linked to local economic conditions, including housing market characteristics.

2.2 Housing Supply Constraints

Housing supply in the United States varies enormously across regions. Saiz [2010] constructs estimates of housing supply elasticity for U.S. metropolitan areas based on geographic constraints (the share of land unavailable for development due to bodies of water, steep terrain, and similar features) and regulatory barriers. He finds that supply elasticities range from below 0.7 in the most constrained coastal markets to above 3.0 in the most elastic interior markets.

These supply differences have profound implications for housing prices. Glaeser et al. [2005] show that the long-run increase in housing prices in the United States is driven

primarily by supply-constrained markets, where regulatory and geographic barriers prevent the construction response that would otherwise moderate price increases. Gyourko et al. [2013] identify a set of “superstar cities” in which high-income households compete for a fixed housing stock, driving persistent price appreciation that excludes lower-income residents.

The macroeconomic consequences of housing supply constraints have received growing attention. Hsieh and Moretti [2019] estimate that constraints on housing supply in high-productivity cities like New York and San Francisco have significantly reduced aggregate output by preventing workers from relocating to their most productive employment. Ganong and Shoag [2017] show that the historical pattern of regional income convergence in the United States has reversed since the 1980s, with housing costs playing a central role in deterring the migration of low-skilled workers to high-wage regions. Howard and Liebersohn [2021] decompose the sources of high rents in inelastic cities, finding that growing demand to live in supply-constrained locations—rather than changes in supply conditions themselves—accounts for most of the rent increase. Albouy et al. [2016] document the resulting cost-of-living inequality, showing that housing costs in high-amenity cities have risen dramatically relative to incomes, with disproportionate effects on lower-income households.

2.3 Housing, Inequality, and the Missing Link

Despite the voluminous literatures on income inequality and housing supply constraints, relatively little work has connected the two in a unified framework. Rognlie [2015] observes that the rising capital share of income documented by Piketty is driven almost entirely by housing, suggesting that real estate plays a unique role in the dynamics of wealth inequality. Mian and Sufi [2014] highlight the role of housing wealth in consumption patterns and economic vulnerability. Bertrand and Morse [2016] document “trickle-down consumption” patterns in which lower-income households in areas with high top-income growth spend more on housing and other visible consumption goods. Kaplan et al. [2018] and Auclert [2019] emphasize the importance of housing expenditure heterogeneity for the transmission of monetary policy.

Yet the mechanism we emphasize—that housing supply constraints *amplify* the translation of nominal income divergence into real consumption inequality through differential shelter cost inflation—has not been tested directly. The shelter wedge concept fills this gap by providing a framework for measuring how much of the Great Decompression is channeled through housing markets and how this channeling varies with local supply conditions.

3 Conceptual Framework

3.1 Setup

Consider a single housing market populated by a continuum of households indexed by income y , drawn from a distribution $F(y)$. Each household allocates a share ϕ of its income to housing and consumes a composite numeraire good with the remainder. The housing market features a supply function that relates the equilibrium price P to the quantity of housing demanded D and the available housing stock S :

$$P = \left(\frac{D}{S} \right)^{\frac{1}{1+\eta}}, \quad (1)$$

where $\eta > 0$ is the housing supply elasticity. As $\eta \rightarrow 0$, supply becomes perfectly inelastic and demand shocks are fully capitalized into prices. As $\eta \rightarrow \infty$, supply is perfectly elastic and demand shocks are absorbed by new construction with no price effect.

3.2 Housing Demand and Income Divergence

Aggregate housing demand is driven by income. With a constant expenditure share ϕ , total housing demand is proportional to aggregate income:

$$D = \phi \cdot \bar{y}, \quad (2)$$

where $\bar{y}$ denotes average income. The shelter burden—the share of income that a household with income y must spend to afford housing at price P —is given by:

$$b(y) = \frac{P}{y}. \quad (3)$$

Now consider an increase in income inequality, modeled as an increase in the ratio y_{80}/y_{50} while holding y_{50} constant. This is the empirical counterpart of the Great Decompression: the 80th-percentile household pulls away from the median. The increase in upper-tail incomes raises aggregate housing demand, which in turn raises the equilibrium price according to equation (1). The price increase raises the shelter burden for the median household:

$$\frac{\partial b(y_{50})}{\partial (y_{80}/y_{50})} = \frac{1}{y_{50}} \cdot \frac{\partial P}{\partial (y_{80}/y_{50})}. \quad (4)$$

3.3 The Shelter Wedge Multiplier

The key object of interest is the *shelter wedge multiplier*: the sensitivity of the shelter burden to income divergence. Differentiating the equilibrium condition, we obtain:

$$\frac{\partial P}{\partial \bar{y}} = \frac{1}{1 + \eta} \cdot \frac{P}{\bar{y}}. \quad (5)$$

The shelter wedge multiplier is decreasing in η : the less elastic is housing supply, the larger the price response to a given increase in income divergence, and hence the larger the increase in the shelter burden for median-income households.

Proposition 1 (Supply Amplification). *For two markets with supply elasticities $\eta_C < \eta_E$ (constrained and elastic, respectively), the shelter wedge multiplier satisfies:*

$$\frac{m(\eta_C)}{m(\eta_E)} = \frac{1 + \eta_E}{1 + \eta_C} > 1, \quad (6)$$

where $m(\eta) \equiv \partial b(y_{50}) / \partial (y_{80}/y_{50})$ is the shelter wedge multiplier.

Calibrating with $\eta_C = 0.7$ (constrained) and $\eta_E = 2.5$ (elastic), the model yields multipliers of 0.069 and 0.031, respectively, for a ratio of 2.20 $\times$ (Table 3). That is, the model predicts that constrained markets amplify the shelter wedge by a factor of 2.2 relative to elastic markets.

3.4 Counterfactual

The model permits a counterfactual exercise: if constrained states had elastic housing supply ($\eta = 2.5$ instead of $\eta = 0.7$), how much lower would the median shelter burden be? The calibrated model implies that the median rent burden in constrained states would decline by 70.6%, from 0.212 to a level approaching that of elastic states (0.184). This counterfactual underscores the potential welfare significance of supply-side housing reforms, even conditional on the null empirical result documented below.

3.5 Empirical Predictions

The model generates three testable predictions:

- (i) The shelter wedge multiplier is positive: increases in the income ratio y_{80}/y_{50} should increase the shelter burden for median households.

- (ii) The shelter wedge multiplier is larger in supply-constrained markets: the interaction of supply inelasticity and income divergence should predict higher shelter burdens.
- (iii) The amplification ratio should be approximately $2.2\times$ for constrained relative to elastic markets.

We test these predictions in the cross-state panel described below.

4 Data

Our analysis combines four data sources spanning both national time-series and state-level panel dimensions.

4.1 National Time Series

For the national analysis, we draw on monthly and quarterly data from the Federal Reserve Economic Data (FRED) database [Federal Reserve Bank of St. Louis, 2024]. The key series are:

- **Shelter CPI** (CUSR0000SAH1): The Consumer Price Index for shelter, which includes rent of primary residence and owners' equivalent rent. We use this series from 1970 to 2025 to construct the shelter premium ratio.
- **Overall CPI** (CPIAUCSL): The Consumer Price Index for all urban consumers, all items. The ratio of shelter CPI to overall CPI is our primary measure of the shelter wedge at the national level. This ratio rose from 0.897 in 1970 to 1.291 in 2025, indicating that shelter costs have grown 44% faster than the overall price level over this period (Figure 1).
- **Case-Shiller Home Price Index** (CSUSHPINSA): A national repeat-sales home price index, available from 1987. We use this alongside median household income (MEHOINUSA672N) to construct the national home price-to-income ratio (Figure 2).
- **Macroeconomic controls**: We include the 30-year fixed mortgage rate (MORTGAGE30US), the federal funds rate (FEDFUNDS), the civilian unemployment rate (UNRATE), and total building permits (PERMIT) as controls in the local projection analysis.

4.2 State Panel: Census American Community Survey

For the cross-state analysis, we draw on the Census Bureau’s American Community Survey (ACS) 1-year estimates [U.S. Census Bureau, 2024] for the period 2010–2023, obtained via the Census API. The key variables are:

- **Median household income** (B19013_001E): State-level median household income in current dollars.
- **80th percentile household income** (B19080_004E): The 80th percentile of the household income distribution, used to construct the income gap ratio y_{80}/y_{50} .
- **Median gross rent** (B25064_001E): The median gross rent paid by renter-occupied households, in current dollars.
- **Median home value** (B25077_001E): The median value of owner-occupied housing units, used as an alternative outcome.
- **Population** (B01001_001E): Total population, used as a control variable.

The primary outcome variable is the *rent burden ratio*, defined as:

$$\text{Rent Burden}_{s,t} = \frac{12 \times \text{Median Gross Rent}_{s,t}}{\text{Median Household Income}_{s,t}}. \quad (7)$$

4.3 State Economic Accounts

We obtain state-level per capita personal income from the Bureau of Economic Analysis (BEA) Regional Economic Accounts [Bureau of Economic Analysis, 2024], used as an economic control variable in augmented specifications and as a placebo outcome.

4.4 Housing Supply Elasticities

Housing supply elasticities are drawn from Saiz [2010], who estimates MSA-level elasticities based on the share of land unavailable for residential development. We aggregate these to the state level using population-weighted averages across MSAs within each state. States are classified into three groups:

- **Constrained** (Saiz elasticity < 1.2 , $N = 12$): CA, CT, DC, FL, HI, MA, MD, NH, NJ, NY, RI, WA.

- **Moderate** (Saiz elasticity 1.2–2.0): Omitted from the primary analysis to sharpen the contrast.
- **Elastic** (Saiz elasticity > 2.0 , $N = 23$): AL, AR, AZ, IA, ID, IN, KS, KY, MI, MO, MS, MT, NC, ND, NE, OH, OK, SC, SD, TN, TX, WV, WY.

4.5 Summary Statistics

Table 1 presents summary statistics for the state panel by supply group. The panel contains 714 observations across 51 states and 14 years. Mean rent burden across all observations is 0.194 (19.4% of income). Constrained states exhibit substantially higher rent burdens (0.212 vs. 0.184 for elastic states), higher median incomes (\$72,615 vs. \$53,639), and higher median home values (\$375,735 vs. \$170,645). The mean income gap (y_{80}/y_{50}) is 1.958 across all states, slightly higher in constrained states (1.996) than elastic states (1.952).




Figure 1: The Shelter Wedge: Shelter CPI Relative to Overall CPI, 1970–2025
Notes: Ratio of shelter CPI (CUSR0000SAH1) to overall CPI (CPIAUCSL), both indexed to their respective 1982–84 base periods. The ratio rises from 0.897 in 1970 to 1.291 in 2025, implying that shelter costs grew 44% faster than the overall price level. Source: FRED.




Figure 2: Home Prices Relative to Income, 1987–2025
Notes: National Case-Shiller Home Price Index (CSUSHPINSA) divided by median household income (MEHOINUSA672N), both normalized. Source: FRED.

Table 1: Summary Statistics by Housing Supply Group

Variable	All			Constrained			Elastic		
	Mean	SD	N	Mean	SD	N	Mean	SD	N
Rent burden ratio	0.194	0.023	714	0.212	0.026	168	0.184	0.019	32
Home value / income	3.818	1.189	714	5.107	1.428	168	3.135	0.620	32
Income gap (80th/50th)	1.958	0.087	714	1.996	0.101	168	1.952	0.071	32
Median HH income (\$)	60,650	13,221	714	72,615	13,176	168	53,639	8,940	32
Median gross rent (\$)	984	264	714	1,276	250	168	817	145	32
Median home value (\$)	240,680	121,244	714	375,735	142,485	168	170,645	57,036	32
Population	6,333,585	7,148,556	714	9,693,377	10,949,041	168	5,434,988	5,707,366	32
Saiz supply elasticity	1.900	0.804	714	0.808	0.189	168	2.680	0.271	32

Notes: State-year panel, 2010–2023. Constrained states have Saiz (2010) supply elasticity < 1.2 ; elastic states have elasticity > 2.0 . Rent burden ratio = annual median rent / median household income. Income gap = 80th percentile income / median income.

5 Empirical Strategy

5.1 Difference-in-Differences

Our primary identification strategy exploits the interaction between time-invariant cross-state variation in housing supply elasticity and the time-varying national income gap. The baseline specification is:

$$Y_{s,t} = \beta \cdot (\text{Inelasticity}_s \times \text{IncomeGap}_t) + \mu_s + \lambda_t + \varepsilon_{s,t}, \quad (8)$$

where $Y_{s,t}$ is the rent burden ratio in state s at time t ; $\text{Inelasticity}_s = 1/\text{Elasticity}_s$ is the inverse of the Saiz supply elasticity (time-invariant, absorbed by state fixed effects in its main effect); $\text{IncomeGap}_t = y_t^{80}/y_t^{50}$ is the national 80th-to-50th income percentile ratio (time-varying, absorbed by year fixed effects in its main effect); μ_s is a state fixed effect; λ_t is a year fixed effect; and $\varepsilon_{s,t}$ is the error term. Standard errors are clustered at the state level using the CR1 sandwich estimator, which is robust to arbitrary heteroskedasticity and within-state serial correlation.

The parameter β is the *shelter wedge amplifier*: the additional rent burden in more supply-inelastic states when the national income gap is wider. Under the null hypothesis that housing supply constraints do not amplify the Great Decompression at the state level, $\beta = 0$.

We estimate three additional specifications:

1. **With controls** (Specification 2): We add state-level per capita income growth and population growth as time-varying controls to absorb state-specific economic trends.
2. **Home value ratio** (Specification 3): We replace the rent burden outcome with the home value-to-income ratio to test whether the mechanism operates through asset prices rather than (or in addition to) rental markets.
3. **Binary DiD** (Specification 4): We replace the continuous interaction with a binary treatment—an indicator for constrained states ($\text{Elasticity}_s < 1.2$) interacted with a post-2015 indicator—to test for a discrete structural break.

5.2 Identification Assumptions

The validity of our DiD strategy rests on the assumption that, absent differential supply constraints, shelter cost burdens in constrained and elastic states would have evolved on parallel trends in response to national income divergence. This assumption is plausible

because the Saiz elasticity measure is based on geographic features (share of developable land) that were determined by physical geography long before the sample period. Nevertheless, we recognize potential threats to identification:

- Supply constraints may be correlated with other state characteristics (e.g., industry composition, amenity values, or policy environments) that independently affect shelter burdens.
- Sorting of high-income workers into supply-constrained states may confound the treatment effect with differential demand shocks.
- The state-level aggregation may mask within-state heterogeneity, as discussed in Section 9.

We address these concerns through an event study design, a placebo outcome test, and permutation inference.

A separate methodological concern involves the recent literature on two-way fixed effects (TWFE) estimators with heterogeneous treatment effects. Goodman-Bacon [2021] shows that staggered DiD designs can produce biased estimates when treatment effects vary over time or across groups, and De Chaisemartin and D’Haultfoeuille [2020] formalize conditions under which TWFE estimates may be a weighted average of treatment effects with negative weights. In our setting, the treatment variable—supply inelasticity interacted with the national income gap—is continuous rather than binary, and all states are “treated” throughout the panel (the interaction varies in intensity, not in adoption timing). This design mitigates the staggered-adoption biases that these papers emphasize. Nevertheless, if the relationship between supply constraints and shelter burdens varies heterogeneously across states or time periods, the TWFE estimator may not recover a well-defined average treatment effect. The short panel (14 years) further limits the scope for time-varying heterogeneity to generate bias.

5.3 Event Study

To examine pre-trends and the dynamics of treatment effects, we estimate an event study specification:

$$Y_{s,t} = \sum_{t \neq 2012} \beta_t \cdot \text{Constrained}_s \times \mathbf{1}(\text{year} = t) + \mu_s + \lambda_t + \varepsilon_{s,t}, \quad (9)$$

where Constrained_s is an indicator for constrained states (Saiz elasticity < 1.2) and 2012 is the base year (normalized to $\beta_{2012} = 0$). The event study coefficients $\{\beta_t\}$ trace out the

evolution of the rent burden gap between constrained and elastic states over time. Under the parallel trends assumption, the pre-treatment coefficients should be zero.

5.4 Local Projections

To complement the cross-state analysis with a national time-series perspective, we estimate local projections [Jordà, 2005] of the shelter CPI response to income shocks:

$$\text{ShelterCPI}_{t+h} - \text{ShelterCPI}_{t-1} = \alpha_h + \beta_h \cdot \text{IncomeGrowth}_t + \gamma_h \cdot X_t + \varepsilon_{t+h}, \quad (10)$$

for horizons $h = 0, 1, \dots, 24$ months. Standard errors are computed using the Newey-West heteroskedasticity and autocorrelation consistent (HAC) estimator with bandwidth $h + 1$, following the standard practice for local projections. We estimate this specification on the full sample and separately for the pre-2000 and post-2000 subsamples to test whether the shelter cost sensitivity to income shocks has increased over time—as would be predicted by tightening supply constraints.

6 Results

6.1 Main DiD Results

Table 2 presents the main difference-in-differences results across four specifications. The central finding is a null result: none of the specifications yields a statistically significant estimate of the shelter wedge amplifier.

In the baseline specification (column 1), the estimated coefficient on the interaction of supply inelasticity and the national income gap is $\hat{\beta} = 0.0107$ with a standard error of 0.0845, yielding a p -value of 0.899. The point estimate is economically negligible: a one-unit increase in the interaction term is associated with a 1.07 percentage-point increase in the rent burden, but this effect is indistinguishable from zero.

Adding per capita income growth and population growth as controls (column 2) yields $\hat{\beta} = -0.0075$ with a standard error of 0.115 ($p = 0.948$). The sign reversal relative to the baseline is itself indicative of the fragility of the estimate. The within- R^2 increases from 0.000 to 0.095, suggesting that the controls are meaningful predictors of rent burden variation, but the treatment effect remains absent.

When we replace the rent burden outcome with the home value-to-income ratio (column 3), the estimate is $\hat{\beta} = -4.27$ with a standard error of 2.65 ($p = 0.113$). While closer

to conventional significance thresholds, the negative sign is the *opposite* of the theoretical prediction—it implies that supply-constrained states experience relatively *lower* home value-to-income ratios when the national income gap widens, which is difficult to rationalize within the model.

The binary DiD specification (column 4), which interacts a constrained state indicator with a post-2015 dummy, produces $\hat{\beta} = -0.0005$ with a standard error of 0.0022 ($p = 0.827$). This estimate is negligible in both statistical and economic terms.

To assess whether the null result reflects a genuinely absent effect or insufficient statistical power, we compute the minimum detectable effect (MDE). With $N = 714$ observations, $G = 51$ state-level clusters, a clustered standard error of 0.0845, and targeting 80% power at the 5% significance level, the MDE is 0.241—approximately 124% of the mean rent burden (0.194). This indicates that the design can detect only very large effects: a treatment coefficient would need to be roughly as large as the mean outcome itself to achieve statistical significance. The observed coefficient of 0.011 is an order of magnitude below the MDE. This power analysis suggests that while the null is genuine in the sense that the state-level effect is far from large, we cannot rule out economically meaningful effects at finer geographic scales where the signal-to-noise ratio would be substantially higher.

Table 2: The Shelter Wedge: Main DiD Results

	(1)	(2)	(3)	(4)
	Baseline	With controls	Home value ratio	Binary DiD
Dep. var:	Rent burden	Rent burden	HV/Income	Rent burden
Inelast. $\times$ Decomp.	0.0107 (0.0845)	-0.0075 (0.1146)	-4.2701 (2.6482)	-0.0005 (0.0022)
Controls	No	Yes	No	No
State FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	714	663	714	714
R^2 (within)	0.000	0.095	0.002	0.000

Notes: Standard errors clustered at the state level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Columns (1)–(2) use rent burden (annual median rent / median income) as outcome. Column (3) uses home value / income ratio. Column (4) uses a binary constrained indicator interacted with a post-2015 dummy. All specifications include state and year fixed effects.

6.2 Event Study

Figure 3 presents the event study coefficients from equation (9). The estimates trace the evolution of the rent burden gap between constrained and elastic states, with 2012 as the base year. Two features are noteworthy.

First, the pre-treatment coefficients (2010–2011) are close to zero and statistically indistinguishable from zero at conventional significance levels, consistent with the parallel trends assumption. There is no evidence of systematic divergence prior to the base year.

Second, the post-treatment coefficients show no clear divergence trend. The estimates are close to zero throughout the sample period, with some positive movement in 2021–2023 that is not statistically significant. The 95% confidence intervals are wide throughout, reflecting the imprecision of the state-level estimates. The event study is consistent with the null DiD results in Table 2: there is no evidence of a systematic shelter wedge amplification at the state level.




Figure 3: Event Study: Rent Burden Gap Between Constrained and Elastic States
Notes: Point estimates and 95% confidence intervals from equation (9). The dependent variable is the rent burden ratio. The base year is 2012 (normalized to zero). Constrained states have Saiz (2010) supply elasticity < 1.2 . Standard errors clustered at the state level.




Figure 4: Housing Supply Elasticity and Shelter Burden Growth, by State
Notes: Each point represents a state. The horizontal axis measures the Saiz (2010) housing supply elasticity (higher = more elastic). The vertical axis measures the growth in the rent burden ratio over the sample period 2010–2023.

6.3 Local Projections

Figure 5 presents the local projection impulse responses of shelter CPI to income shocks, estimated separately for the full sample, the pre-2000 subsample, and the post-2000 subsample.

In the full sample, the cumulative 12-month response of shelter CPI to a unit income shock is $\hat{\beta}_{12} = 0.0038$ —a small but positive effect indicating that income growth is associated with modest shelter cost inflation at the national level.

The subsample analysis reveals a surprising pattern. In the pre-2000 period, the 12-month response is $\hat{\beta}_{12} = 0.0182$, nearly five times larger than the full-sample estimate. In the post-2000 period, however, the response falls to $\hat{\beta}_{12} = 0.0027$ —*smaller* than the pre-2000 estimate. This pattern is the opposite of what the tightening-supply-constraints narrative would predict. If supply constraints have tightened over time (as documented by Glaeser

et al. [2005] and Gyourko et al. [2013]), we would expect the shelter cost sensitivity to income shocks to *increase* in the post-2000 period, not decrease.

This finding admits several interpretations. First, the post-2000 weakening may reflect the tightening of mortgage standards following the Dodd-Frank Wall Street Reform and Consumer Protection Act (2010) and the broader post-GFC regulatory environment, which reduced the transmission of income growth into housing demand by constraining credit access. Second, the structural shift toward renting—the homeownership rate fell from 69% in 2004 to 64% by 2016—may have altered the income-shelter cost relationship, as rental markets respond differently to income shocks than owner-occupied markets. Third, the pre-2000 period coincides with the most rapid phase of the Great Decompression (1980s–1990s), during which the income distribution was shifting most dramatically; the post-2000 stabilization of top income shares [Piketty and Saez, 2003] may have reduced the magnitude of the income shocks driving shelter cost inflation. Finally, the aggregate time-series framework may simply be too coarse to capture the cross-sectional variation in supply constraints that drives the shelter wedge at the local level.




Figure 5: Local Projection Impulse Responses: Shelter CPI to Income Shocks
Notes: Cumulative impulse response of shelter CPI to a unit income growth shock, estimated by local projections [Jordà, 2005]. Solid lines show point estimates; shaded areas show 90% confidence intervals. The pre-2000 response ($\hat{\beta}_{12} = 0.0182$) is substantially larger than the post-2000 response ($\hat{\beta}_{12} = 0.0027$). Source: FRED.

7 Mechanisms and Model Estimation

7.1 Calibration

We calibrate the structural model developed in Section 3 to match key moments of the data. The calibrated housing expenditure share is $\phi = 0.60$. The supply elasticity parameters are drawn directly from Saiz [2010]: $\eta_C = 0.7$ for constrained markets and $\eta_E = 2.5$ for elastic markets.

Table 3 presents the calibration results. Panel A reports the calibrated parameters. Panel B compares the model-implied moments to their data counterparts. The model matches the rent burden in constrained states exactly (0.212) but underpredicts the rent burden in elastic states (0.108 model vs. 0.184 data). This calibration miss reflects the simplicity of the single-market framework, which attributes the entire rent burden to the supply-demand equilibrium in housing. In practice, elastic-market rents incorporate non-housing amenity values, property taxes, maintenance costs, and other components that raise equilibrium prices above what pure supply-demand clearing would predict. Additionally, the spatial sorting model of Diamond [2016] suggests that even elastic markets attract workers whose willingness to pay for local amenities sustains rent burdens above the model’s prediction. A richer multi-market general equilibrium framework with amenity heterogeneity and sorting would likely narrow this gap, but at the cost of additional parameters and reduced transparency.

7.2 Implied Shelter Wedge Multipliers

Panel C of Table 3 reports the implied shelter wedge multipliers. In constrained markets ($\eta = 0.7$), the multiplier is 0.069: a one-unit increase in the income ratio y_{80}/y_{50} increases the shelter burden gap by 6.9 percentage points. In elastic markets ($\eta = 2.5$), the multiplier is 0.031—less than half the constrained-market value. The ratio of constrained to elastic multipliers is $2.20\times$, consistent with the theoretical prediction of Proposition 1.




Figure 6: Model-Implied Shelter Wedge Multiplier as a Function of Supply Elasticity
Notes: The shelter wedge multiplier $m(\eta)$ from the calibrated structural model, plotted as a function of the housing supply elasticity η . The multiplier is decreasing in η : more constrained markets (lower η) exhibit larger amplification of income divergence into shelter cost divergence.

Vertical lines indicate the constrained ($\eta = 0.7$) and elastic ($\eta = 2.5$) calibration points.

Table 3: Structural Model: Calibration and Fit

<i>Panel A: Calibrated Parameters</i>		
α (housing expenditure share)		0.0000
S_L/S_H (supply ratio)		0.0000
$\eta_{\text{constrained}}$ (supply elasticity)		0.7
η_{elastic} (supply elasticity)		2.5
<i>Panel B: Moment Fit</i>		
	Data	Model
Rent burden (constrained)	0.2120	0.2120
Rent burden (elastic)	0.1836	0.1080
<i>Panel C: Implied Shelter Wedge Multipliers</i>		
Multiplier (constrained)		0.069086
Multiplier (elastic)		0.031446
Ratio (constrained / elastic)		2.20×

Notes: Two-neighborhood spatial sorting model. α is the share of income allocated to housing. η is the Saiz (2010) housing supply elasticity (lower = more constrained). The shelter wedge multiplier measures $d(\text{wedge})/d(\text{income ratio})$: how much a unit increase in the 80th/50th income ratio widens the shelter burden gap.

7.3 Reconciling Theory and Evidence

The structural model predicts a sizable shelter wedge amplification—a 2.2× ratio between constrained and elastic markets—yet the empirical analysis finds no significant effect at the state level. This apparent contradiction can be reconciled by recognizing that the model is a *within-market* model of sorting and price formation. The model’s predictions apply to the comparison between a constrained neighborhood and an elastic neighborhood *within the same metropolitan area*, not to the comparison of entire states that contain heterogeneous local housing markets.

When we aggregate the model’s predictions to the state level, the within-state heterogeneity in supply elasticity, demand conditions, and sorting patterns generates sufficient noise to obscure the signal. California contains both San Francisco (Saiz elasticity ≈ 0.6) and Riverside-San Bernardino (elasticity ≈ 2.1); the state-level average masks this critical variation. The model’s predictions are thus consistent with the null empirical result at the

state level, provided that the mechanism operates at the metropolitan or sub-metropolitan level.

8 Robustness

We conduct a battery of robustness checks to assess the sensitivity of the null result. Table 4 presents the results.

8.1 Alternative Treatment Measures

Column 2 of Table 4 replaces the binary constrained-elastic classification with the continuous negative Saiz elasticity, so that the treatment intensity varies smoothly across states. The estimated coefficient is 0.0068 with a standard error of 0.0365, confirming the null result. The continuous specification has greater power to detect a monotone relationship between supply elasticity and the shelter wedge, but none is detected.

8.2 Restricted Sample

Column 3 drops moderate-elasticity states, restricting the sample to 490 observations from the 12 most constrained and 23 most elastic states. The estimate ($\hat{\beta} = 0.0171$, $SE = 0.0852$) remains statistically insignificant. The point estimate is marginally larger than the baseline, but well within the confidence interval.

Column 4 excludes the District of Columbia and Hawaii—two outliers that may exert undue influence due to their unique housing market characteristics. The estimate shifts to $\hat{\beta} = -0.0591$ ($SE = 0.0863$), moving in the negative direction and remaining insignificant. This sign instability across specifications is itself indicative of a null underlying effect.

8.3 Permutation Inference

To assess the probability of obtaining a coefficient as large as the baseline estimate under the null hypothesis of no treatment effect, we conduct permutation inference. We randomly reassign state-level treatment status (constrained vs. elastic) 1,000 times, re-estimating the baseline specification each time. The resulting p -value is 0.868: 86.8% of the permuted samples yield a coefficient at least as large as the observed estimate. This strongly confirms the null—the observed estimate is well within the range generated by random assignment.

8.4 Placebo Outcome

Column 5 of Table 4 replaces the rent burden outcome with per capita income growth as a placebo. If the treatment (supply inelasticity interacted with income divergence) is unrelated to non-shelter economic outcomes, the placebo coefficient should be zero. Instead, we find $\hat{\beta} = -0.458$ with a standard error of 0.144 ($p = 0.003$)—a highly significant negative effect. This indicates that constrained states experience *slower* per capita income growth when the national income gap widens.

This placebo failure warrants careful interpretation. Supply-constrained states are not randomly assigned: they tend to be high-amenity, high-productivity coastal states (California, New York, Massachusetts) with distinct industrial composition, migration patterns, and labor market dynamics. The negative placebo coefficient likely reflects that these states’ per capita income growth is mechanically linked to the national income gap through composition effects—as the national 80th/50th ratio widens, high-productivity states where upper-income workers concentrate may experience differential income dynamics independent of housing supply. This violates the exclusion restriction in a strict sense, but the violation runs through the economic structure of constrained states rather than through the housing market channel that the shelter wedge hypothesis posits. The placebo result thus reinforces both the limitation of state-level analysis and the need for finer geographic units where supply constraints can be isolated from correlated economic characteristics.

8.5 State-Specific Income Gap

Column 6 replaces the national income gap with the state-specific 80th-to-50th income percentile ratio in the interaction term. This specification exploits within-state variation in income divergence rather than relying on the common national trend, providing substantially more identifying variation. The estimated coefficient is $\hat{\beta} = 0.040$ with a standard error of 0.021 ($p = 0.059$), which is marginally significant at the 10% level. This is the only specification that approaches conventional significance thresholds, suggesting that supply constraints may interact with *local* income dynamics in ways that are obscured when using the national income gap. The result is consistent with our interpretation that the shelter wedge mechanism operates at finer geographic scales.

Table 4: Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline	Continuous elast.	Excl. moderate	Excl. DC/HI	Placebo (PCI)	State in
Treatment	0.0107 (0.0845)	0.0068 (0.0365)	0.0171 (0.0852)	-0.0591 (0.0863)	-0.4577*** (0.1441)	0.0 (0.0)
N	714	714	490	686	663	7
Perm. p -value	0.868	—	—	—	—	

Notes: All specifications include state and year FE. Clustered SEs at state level. Column (2) uses continuous negative Saiz elasticity. Column (3) drops moderate-elasticity states. Column (5) uses per capita income growth as a placebo outcome. Column (6) replaces the national income gap with the state-specific 80th/50th income ratio. Permutation p -value from 1,000 random reassignments of treatment across states.

9 Discussion

9.1 Why a Null Result at the State Level?

The null result at the state level admits several interpretations, which we discuss in order of plausibility.

Geographic aggregation. The most likely explanation is that states are too coarse a geographic unit to detect the shelter wedge. Housing markets are inherently local: rents and home values are determined by conditions within a commuting zone, not across an entire state. A state like California contains MSAs with Saiz elasticities ranging from below 0.6 (San Francisco, Los Angeles) to above 2.0 (inland valleys). Aggregating across these diverse markets attenuates the treatment effect toward zero. This interpretation is consistent with the findings of Hsieh and Moretti [2019], who identify significant effects of housing constraints at the MSA level, and Diamond [2016], who documents sorting at the metro level. It also aligns with Ganong and Shoag [2017], whose analysis of regional convergence operates at the commuting-zone level.

Sorting as a confound. The spatial sorting of high-income workers into supply-constrained states may confound the treatment effect. If high-income workers preferentially move to constrained states (attracted by amenities, job opportunities, or network effects), this simultaneously raises the income gap *and* the rent burden in those states, potentially creating a correlation between treatment and outcome that is driven by demand rather than supply. The significant placebo result—constrained states exhibit slower income growth—

suggests that income dynamics in constrained states are different from those in elastic states in ways not captured by the interaction term.

Measurement. The rent burden ratio, while intuitive, may be a noisy measure of the shelter wedge. Median gross rent reflects the current stock of leases and may adjust sluggishly to current market conditions. Similarly, the 80th-to-50th income percentile ratio is a coarse measure of the Great Decompression at the state level. The income gap may manifest differently across states depending on industry composition, education levels, and demographic factors.

Genuine absence. It is possible, though we consider it unlikely given the strong theoretical prediction and existing metro-level evidence, that the shelter wedge amplification mechanism is genuinely absent at all geographic scales. This would imply that housing supply constraints affect price *levels* but not the *differential* effect of income divergence on shelter burdens.

9.2 Implications for Housing Policy

Despite the null result at the state level, the structural model’s prediction of a $2.20\times$ amplification ratio has important policy implications—provided the mechanism operates at the metropolitan or neighborhood level, as we argue.

First, housing supply reforms in constrained markets—such as zoning liberalization, by-right permitting, and reductions in regulatory barriers—would disproportionately benefit median-income households in those markets by reducing the shelter wedge multiplier. The counterfactual calibration suggests a 70.6% reduction in median rent burden if constrained markets achieved elastic supply levels.

Second, the concentration of high-income workers in supply-constrained cities [Diamond, 2016, Moretti, 2013] implies that the shelter wedge is largest precisely where it matters most for economic mobility. Chetty et al. [2014] document that neighborhood quality is a key determinant of upward mobility; if supply constraints amplify the cost of accessing high-opportunity neighborhoods, the shelter wedge may have dynamic consequences for intergenerational income persistence [Black, 1999].

Third, macroeconomic models that incorporate household heterogeneity in housing expenditure—such as Kaplan et al. [2018] and Auclert [2019]—would benefit from incorporating the shelter wedge mechanism. The differential sensitivity of monetary policy transmission across constrained and elastic housing markets could have implications for the distributional effects of interest rate changes.

9.3 Directions for Future Research

Our findings point to several productive directions for future work.

The most direct extension is to replicate the analysis at the MSA level, exploiting within-MSA variation in housing supply elasticity. The Census ACS provides MSA-level data on rents, incomes, and home values that would permit a direct test of the shelter wedge hypothesis at the geographic unit where the mechanism is likely to operate. Alternatively, neighborhood-level data from the American Community Survey’s Public Use Microdata Areas (PUMAs) or from IPUMS [IPUMS, 2024] could support an even finer-grained analysis.

A complementary approach would exploit the time-series of supply constraints—for example, using changes in local zoning regulations or the introduction of growth management ordinances as quasi-experimental variation. This would address the concern that cross-sectional supply elasticity may be correlated with unobserved state characteristics.

Finally, the relationship between the shelter wedge and intergenerational mobility deserves further investigation. If housing cost amplification reduces the ability of middle-income families to reside in high-opportunity neighborhoods, the shelter wedge may contribute to the persistence of inequality across generations—a mechanism that connects the Great Depression to the geography of opportunity documented by Chetty et al. [2014].

10 Conclusion

This paper introduces the concept of the shelter wedge—the amplification of nominal income inequality through differential shelter cost inflation—and tests whether housing supply constraints magnify this wedge at the state level. We document that the ratio of shelter CPI to overall CPI has risen from 0.897 in 1970 to 1.291 in 2025, a 44% increase that implies shelter costs have persistently outpaced other consumer prices.

A structural two-neighborhood model predicts that supply-constrained markets amplify the shelter wedge by a factor of 2.2 relative to elastic markets: income divergence translates into a 6.9 percentage-point increase in the shelter burden gap in constrained markets versus 3.1 points in elastic markets. However, empirical tests using a difference-in-differences design on a panel of 714 state-year observations (51 states, 2010–2023), exploiting variation in Saiz (2010) housing supply elasticities, yield a precisely estimated null result. The baseline estimate of the shelter wedge amplifier is $\hat{\beta} = 0.0107$ ($p = 0.899$), confirmed by permutation inference ($p = 0.868$), event study analysis showing no divergence trend, and national local projections showing a *declining* shelter cost sensitivity to income shocks since 2000.

We interpret this null result as informative rather than dispositive. A power calculation

reveals that the state-level design can detect only very large effects ($\text{MDE} = 0.241$, or 124% of the mean rent burden), and one robustness specification—using state-specific rather than national income gaps—yields a marginally significant estimate ($p = 0.059$), suggesting that the mechanism may emerge when more granular income variation is exploited. The evidence is consistent with a shelter wedge that operates at a geographic granularity finer than states—likely at the metropolitan or neighborhood level, where the spatial sorting processes that drive housing cost inflation are most concentrated. This interpretation aligns with the metropolitan-level findings of Hsieh and Moretti [2019] and Diamond [2016], and suggests that future empirical tests of the shelter wedge hypothesis should target MSA-level or neighborhood-level data.

The shelter wedge concept remains valuable as a theoretical framework for understanding how the Great Decompression interacts with housing markets to shape real consumption inequality. As income divergence continues and supply constraints tighten in high-productivity cities, the shelter wedge may emerge as one of the primary channels through which nominal inequality becomes lived inequality—with consequences for economic mobility, household welfare, and the geography of opportunity.

References

- David Albouy, Gabriel Ehrlich, and Yingyi Liu. Housing demand, cost-of-living inequality, and the affordability crisis. *NBER Working Paper 22816*, 2016.
- Adrien Auclert. Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–2367, 2019.
- David H. Autor, Lawrence F. Katz, and Melissa S. Kearney. Trends in U.S. wage inequality: Revising the revisionists. *Review of Economics and Statistics*, 90(2):300–323, 2008.
- Marianne Bertrand and Adair Morse. Trickle-down consumption. *Review of Economics and Statistics*, 98(5):863–879, 2016.
- Sandra E. Black. Do better schools matter? Parental valuation of elementary education. *Quarterly Journal of Economics*, 114(2):577–599, 1999.
- Bureau of Economic Analysis. Regional economic accounts, 2024. <https://www.bea.gov/data/economic-accounts/regional>.
- Raj Chetty, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. Where is the land of opportunity? The geography of intergenerational mobility in the United States. *Quarterly Journal of Economics*, 129(4):1553–1623, 2014.
- Clément De Chaisemartin and Xavier D’Haultfœuille. Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996, 2020.
- Rebecca Diamond. The determinants and welfare implications of US workers’ diverging location choices by skill: 1980–2000. *American Economic Review*, 106(3):479–524, 2016.
- Federal Reserve Bank of St. Louis. FRED economic data, 2024. <https://fred.stlouisfed.org>.
- Peter Ganong and Daniel Shoag. Why has regional income convergence in the U.S. declined? *Journal of Urban Economics*, 102:76–90, 2017.
- Edward L. Glaeser, Joseph Gyourko, and Raven E. Saks. Why have housing prices gone up? *American Economic Review*, 95(2):329–333, 2005.
- Claudia Goldin and Lawrence F. Katz. *The Race between Education and Technology*. Harvard University Press, 2008.

- Claudia Goldin and Robert A. Margo. The great compression: The wage structure in the United States at mid-century. *Quarterly Journal of Economics*, 107(1):1–34, 1992.
- Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 2021.
- Joseph Gyourko, Christopher Mayer, and Todd Sinai. Superstar cities. *American Economic Journal: Economic Policy*, 5(4):167–199, 2013.
- Greg Howard and Jack Liebersohn. Why is the rent so darn high? The role of growing demand to live in housing-supply-inelastic cities. *Journal of Urban Economics*, 124, 2021.
- Chang-Tai Hsieh and Enrico Moretti. Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2):1–39, 2019.
- IPUMS. Integrated public use microdata series. Technical report, University of Minnesota, 2024. <https://ipums.org>.
- Òscar Jordà. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 2005.
- Greg Kaplan, Benjamin Moll, and Giovanni L. Violante. Monetary policy according to HANK. *American Economic Review*, 108(3):697–743, 2018.
- Atif Mian and Amir Sufi. What explains the 2007–2009 drop in employment? *Econometrica*, 82(6):2197–2223, 2014.
- Enrico Moretti. Real wage inequality. *American Economic Journal: Applied Economics*, 5(1):65–103, 2013.
- Thomas Piketty and Emmanuel Saez. Income inequality in the United States, 1913–1998. *Quarterly Journal of Economics*, 118(1):1–41, 2003.
- Robert D. Putnam. *Our Kids: The American Dream in Crisis*. Simon & Schuster, 2015.
- Richard V. Reeves. *Dream Hoarders: How the American Upper Middle Class Is Leaving Everyone Else in the Dust*. Brookings Institution Press, 2017.
- Matthew Rognlie. Deciphering the fall and rise in the net capital share: Accumulation or scarcity? *Brookings Papers on Economic Activity*, 2015(1):1–69, 2015.
- Albert Saiz. The geographic determinants of housing supply. *Quarterly Journal of Economics*, 125(3):1253–1296, 2010.

U.S. Census Bureau. American community survey 1-year estimates, 2024. <https://data.census.gov>.

A Data Coverage and Sources

Table 5 summarizes the data sources, variables, and coverage periods used in the analysis.

Table 5: Data Sources and Coverage

Variable	Source	Series / Table	Coverage
<i>National Time Series (FRED)</i>			
Shelter CPI	BLS via FRED	CUSR0000SAH1	1947–2025 (monthly)
Overall CPI	BLS via FRED	CPIAUCSL	1947–2025 (monthly)
Case-Shiller HPI	S&P via FRED	CSUSHPINSA	1987–2025 (monthly)
Median HH income	Census via FRED	MEHOINUSA672N	1984–2024 (annual)
30-yr mortgage rate	Freddie Mac via FRED	MORTGAGE30US	1971–2025 (weekly)
Unemployment rate	BLS via FRED	UNRATE	1948–2025 (monthly)
Building permits	Census via FRED	PERMIT	1960–2025 (monthly)
Fed funds rate	Fed via FRED	FEDFUNDS	1954–2025 (monthly)
<i>State Panel (Census ACS)</i>			
Median HH income	ACS 1-Year	B19013_001E	2010–2023 (annual)
80th pct HH income	ACS 1-Year	B19080_004E	2010–2023 (annual)
Median gross rent	ACS 1-Year	B25064_001E	2010–2023 (annual)
Median home value	ACS 1-Year	B25077_001E	2010–2023 (annual)
Population	ACS 1-Year	B01001_001E	2010–2023 (annual)
<i>State Economic Accounts (BEA)</i>			
Per capita income	BEA Regional	CAINC1	1969–2023 (annual)
<i>Housing Supply Elasticity</i>			
Supply elasticity	Saiz [2010]	MSA-level, pop-weighted	Cross-sectional

A.1 State Classification by Supply Elasticity

Table 6 reports the full classification of states by housing supply elasticity group.

Table 6: State Classification by Housing Supply Elasticity

Group	States
Constrained ($\eta < 1.2$, $N = 12$)	CA, CT, DC, FL, HI, MA, MD, NH, NJ, NY, RI, WA
Elastic ($\eta > 2.0$, $N = 23$)	AL, AR, AZ, IA, ID, IN, KS, KY, MI, MO, MS, MT, NC, ND, NE, OH, OK, SC, SD, TN, TX, WV, WY
Moderate ($1.2 \leq \eta \leq 2.0$, $N = 16$)	AK, CO, DE, GA, IL, LA, ME, MN, NM, NV, OR, PA, UT, VA, VT, WI

Notes: Supply elasticity from Saiz [2010], aggregated from MSA-level to state level using population-weighted averages. Moderate-elasticity states are excluded from the primary DiD analysis.

A.2 Additional Robustness: Continuous Elasticity Scatter

Figure 4 (in the main text) plots the cross-state relationship between Saiz supply elasticity and the growth in rent burden over the sample period. The scatter shows no clear negative relationship: constrained states (low elasticity) do not systematically exhibit larger increases in rent burden than elastic states. This visual evidence corroborates the regression null reported in Tables 2 and 4.

A.3 Permutation Distribution

The permutation inference procedure proceeds as follows. We randomly reassign the Saiz supply elasticity values across states 1,000 times, holding all other data fixed. For each permutation, we re-estimate the baseline specification (equation 8) and record the coefficient on the interaction term. The empirical p -value is the fraction of permuted coefficients whose absolute value exceeds the absolute value of the observed coefficient. The resulting p -value of 0.868 indicates that the observed coefficient of 0.0107 is unremarkable relative to the distribution generated under the null hypothesis of no treatment effect.