Highly Disaggregated Land Unavailability

Chandler Lutz
University of North Carolina, Charlotte

Ben Sand
York University

Land Unavailability Data:
https://github.com/ChandlerLutz/LandUnavailabilityData

July 20, 2026

Abstract

Standard empirical proxies contradict the canonical prediction that supply constraints amplify house price cycles. We resolve this puzzle by constructing the Land Unavailability–Machine Learning (LU-ML) Indices from high-resolution satellite imagery. Unlike existing measures, our indices capture the nonlinear, time-varying, and heterogeneous impacts of physical geography. We provide two new results: (1) physical constraints are a key determinant of cross-sectional price dynamics, driven largely by the intensive margin (steep slopes); and (2) looser constraints significantly mitigate the price effects of demand growth by facilitating supply expansion, though this affordability channel has severely weakened with nationwide tightening.

JEL Classification: R30, R31, R20

Keywords: Land Unavailability, Supply Constraints, House Prices, Entrepreneurship

Amid a crisis of U.S. home affordability, households and legislators have targeted supply constraints as both a culprit for high prices and a lever for policy response (The White House, 2025).[1] This focus matches economic theory: building constraints steepen the housing supply curve, amplifying the price impact of demand shocks. Yet standard empirical proxies for these constraints have limited explanatory power for house prices (Davidoff, 2013, Guren et al., 2021, and Figure 1), while looser constraints, according to these measures, *do not* accommodate supply expansion or mitigate the house price impacts of demand growth (Davidoff, 2013, Louie et al., 2026, and Figure 2). These results, taken at face value, have wide-reaching ramifications for housing markets and upend urban theory, implying that (1) supply constraints explain little of the heterogeneous price growth across cities; (2) demand shocks have similar price and construction effects regardless of supply elasticities; and (3) the extensive literature that exploits supply constraints to sort housing markets or for exogenous variation is unfounded.[2]

This paper resolves this disconnect between theory and evidence by developing novel supply constraint measures that combine Land Unavailability (LU) estimates from highly detailed satellite imagery with machine learning (ML) techniques. The resulting LU-ML indices capture the severity of geographic building restrictions across U.S. housing markets at multiple levels of spatial disaggregation. Unlike previous proxies, the LU-ML indices exhibit strong explanatory power for house prices and reveal that looser constraints do facilitate supply expansion and thus dampen the house price effects of demand growth. These results provide new and robust empirical validation for canonical housing and urban theory.

More specifically, we construct two distinct indices: the LU-ML Price Pressure Index and the LU-ML Supply Index. While both use the same input data and methodology, they capture distinct economic mechanisms. The Price Pressure Index directly targets house price growth, while the Supply Index targets housing unit expansion. Akin to conventional constraint proxies (Saiz, 2010; Baum-Snow and Han, 2024), the Supply Index isolates the relationship between land unavailability

[1] See also The New York Times (2022); Office of the Governor of New York (2023); Office of the Governor of California (2025); Goldman Sachs (2025); United States Senate Committee on Banking, Housing, and Urban Affairs (2025). In the academic literature, see Molloy (2020) and Saiz (2023).

[2] See, for example, Mian and Sufi (2011); Chaney et al. (2012); Mian et al. (2013); Mian and Sufi (2014); Adelino et al. (2015); Aladangady (2017); Chetty et al. (2017); Ganong and Shoag (2017); Stroebel and Vavra (2019); Baldauf et al. (2020); Gao et al. (2020); Kaplan et al. (2020a); Griffin et al. (2021); Howard and Liebersohn (2021); Campello et al. (2022); Gupta et al. (2022, 2025); Fonseca et al. (2025).

and the physical housing stock. By using ML to measure the expectation of housing unit growth given land unavailability, $\mathbb{E}(\text{HU growth} \mid \text{LU})$, the index explicitly captures how tightly geographic constraints restrict housing expansion. In contrast to existing measures, interacting our Supply Index with income growth reveals that looser constraints actively dampen the price impacts of rising demand via increased construction.

Yet physical construction is not the only channel through which geography shapes housing markets. Demand shocks are first capitalized into prices as construction lags (Oh and Yoon, 2020), especially in constrained areas (Glaeser et al., 2008). Housing is also an asset class: prices reflect forward-looking expectations (Cheng et al., 2014; Kaplan et al., 2020b; Chopra et al., 2025), and agents price risk conditional on local supply constraints (Sinai and Souleles, 2005). In constrained markets, the inability of supply to absorb shocks increases asset volatility (Paciorek, 2013), demanding a higher risk premium (Han, 2013; Favilukis et al., 2017) and altering the option value of development (Titman, 1985; Nathanson and Zwick, 2018; Murphy, 2018).

The LU-ML Price Pressure Index explicitly captures this capitalization channel through $\mathbb{E}(\text{HP growth} \mid \text{LU})$. By accounting for expectation effects, it better explains heterogeneous price growth over both the short and long run. Moreover, because the index is largely unrelated to natural amenity demand—and uses an out-of-sample cross-fitting design to purge local, idiosyncratic demand and supply factors—it cleanly isolates the impact of geographic building restrictions on asset prices.

To build the LU-ML Price Pressure Index, we start with the foundational premise from Saiz (2010) that exogenously unavailable land (due to steep slopes, water, and wetlands) steepens the housing supply curve. Yet unlike Saiz and subsequent work (e.g., Baum-Snow and Han, 2024), which aggregate these three types of LU equally within a single, uniform circle for each city, we measure slope, water, and wetlands separately using multiple spatial definitions both within and around a city. This granularity is crucial: the price impact and risk profile of wetlands differ from those of steep slopes, just as the effects of a constraint at the urban core differ from those at the periphery. Our LU-ML approach captures both of these dynamics and allows for data-driven heterogeneity in how physical features impact local markets.

To compile the LU-ML Price Pressure Index for CBSAs, we use 90 features for each city (slope,

water, and wetlands LU measured using 30 different spatial definitions within and around the city). We then train an XGBoost model (Chen and Guestrin, 2016) that uncovers the relationship between physical constraints and prices, $\mathbb{E}(\text{HP growth} \mid \text{LU})$. The key advantages of this approach over standard linear proxies are twofold. First, by jointly evaluating 30 distinct spatial boundaries, the algorithm moves beyond measuring total unweighted land unavailability (Saiz, 2010; Baum-Snow and Han, 2024) and instead learns the spatial layout of physical constraints. It effectively reads a topographic map, distinguishing, for example, a highly restrictive peninsula (e.g., San Francisco) from water distributed harmlessly along a distant periphery (e.g., Houston). Second, it captures the rich nonlinearities and interactions inherent to spatial development, such as the compounding impact of steep slopes meeting coastal boundaries.

Moreover, to avoid other local idiosyncrasies and mitigate model overfitting, we generate the final LU-ML indices via cross-fitting (Chernozhukov et al., 2018). Because each city's index is an out-of-sample prediction generated by a model trained exclusively on other markets, the algorithm cannot conflate a city's physical geography with unique determinants of local supply or demand. Instead, the model learns generalizable rules regarding how physical layouts impact prices nationwide and applies those rules to the hold-out city's topographic blueprint. While this out-of-sample design purges city-specific factors, we also employ empirical specifications that further control for housing demand to isolate the effect of geographic constraints.

The LU-ML indices offer substantially greater flexibility than existing proxies. They scale to various time periods and multiple levels of disaggregation, from cities (CBSAs) down to zip codes and census tracts, for all housing markets in the contiguous U.S. The LU-ML indices also adapt to various empirical specifications: the model can be trained to target cumulative long-run appreciation for long-difference designs or period-specific price growth for panels where the indices reflect how geographic constraints bind over time. This breadth deepens our understanding of the relationship between geographic supply constraints and housing markets.

Figure 1 compares the explanatory power of the LU-ML Price Pressure Index to the standard Saiz (2010) and Baum-Snow and Han (2024, BSH) elasticity proxies via instrument relevance tests for house prices. Details on the data and specifications are in the figure notes. Panel A replicates leading studies across various empirical designs, including panels and long differences. In every

replication, the BSH proxy yields F-statistics for house prices below 10. The F-statistics for the Saiz proxy clear the traditional threshold of 10 but never reach 50, a modern weak instrument threshold suggested by Keane and Neal (2024). The Saiz and BSH proxies thus have limited explanatory power, implying that these standard measures of supply constraints are largely unrelated to cross-sectional house price growth. In sharp contrast, the F-statistics for the LU-ML Price Pressure Index exceed the Keane and Neal threshold across all specifications.

Panel B uses our preferred specification:

$$\Delta \log HP_{i,d,t} = \psi_i + \gamma_{d,t} + \beta \cdot SC_{i,d,t} + \varepsilon_{i,d,t} \tag{1}$$

$\Delta \log HP_{i,d,t}$ is the year-over-year (YoY) log first-difference in house prices for CBSA i in Census Division d at month t, and SC signifies a supply constraint proxy. Since the dependent variable is in log changes (price growth), time-invariant level factors are differenced away, and the CBSA fixed effects (ψ_i) capture heterogeneous city-level trends. The specification thus accounts for differences in both static city-level attributes, due to amenities or regulation, and secular trends, such as persistent price appreciation in "Superstar Cities" (Gyourko et al., 2013) or the secular boom in the "Sand States" (Davidoff, 2013). The Census division-by-time fixed effects non-parametrically control for shocks to each Census division over time.

$SC_{i,d,t}$ represents the effective, time-varying supply constraint. By using static Land Unavailability features to generate out-of-sample predictions of house price growth separately for each period, the Price Pressure Index dynamically captures how geographic constraints amplify broader housing cycles. In contrast, the BSH and Saiz proxies are static cross-sectional measures. To make them comparable in a panel setting, we interact these static proxies with the YoY log first-difference in national house prices (Guren et al., 2021).

Using the panel approach, BSH elasticity is again uncorrelated with house price growth across all sample periods, while Saiz elasticity yields an F-statistic barely above 10 only for the 1980–2024 and 2000–2011 time periods. Altogether, the results from Figure 1 imply that standard supply constraint proxies are largely uninformative for city-level house price growth, especially after accounting for heterogeneous trends and regional shocks. Conversely, the LU-ML Price Pressure Index has strong explanatory power, congruent with geographic supply constraints impacting house prices and validating our LU-ML approach that combines multiple LU estimates via ML techniques.

Finally, note that the results in Figure 1 extend to house prices measured at various levels of disaggregation (CBSAs, counties, zip codes, and census tracts) and all major publicly available housing datasets.

Davidoff (2016) highlights a key concern when using land unavailability as a predictor of house prices: households likely view steep-sloped areas and water as valuable natural amenities, and the demand for these amenities has increased over time. Prices in areas with high LU thus might rise not because of construction constraints, but rather due to household demand for these amenities.

The panel structure in Equation 1 directly addresses Davidoff's critique. Because the dependent variable is in log changes, time-invariant amenity heterogeneity is differenced away, and the CBSA fixed effects account for city-specific trends in amenity demand over time. Yet we also control for various city-level demand factors to capture heterogeneous housing demand across cities—such as initial college share to account for skill-biased migration and resulting endogenous amenity formation (Diamond, 2016)—and natural amenity demand using an index of Google Maps reviews for hiking, water, and natural recreation (Lutz, 2025b). Although we document a natural amenity housing premium, where markets with more attractive natural amenities experience higher house price growth, these amenities account for less than 4% of the variation in the LU-ML Price Pressure Index after accounting for standard factors. The relationship between the LU-ML Index and house prices thus remains large and highly significant after including local controls, and sensitivity analyses (Cinelli and Hazlett, 2020) further show that this relationship persists even under extreme scenarios of hypothetical omitted variables.

The strong predictive performance of the LU-ML Price Pressure Index emanates from its integration of diverse LU estimates within a flexible, data-driven framework. A closer examination of the underlying features reveals that steep slopes are the dominant driver of house price growth, far exceeding the impact of water. Economically, this suggests that constraints on the intensive margin—where slopes sharply increase the marginal cost of densification and infrastructure—are more restrictive for price appreciation than constraints on the extensive margin, such as water boundaries. Wetlands play a negligible role, as they are "soft" constraints often overcome by land-improvement technologies (Dahl, 2011).

The XGBoost algorithm also uncovers important interactions and nonlinearities that standard

proxies miss (Kaplan et al., 2020a). We find significant interaction effects between slope and water, as well as a concave relationship between slope ruggedness and prices. Increasing LU due to steep slopes initially raises house prices, but rugged terrain eventually dampens them. This nonlinearity mirrors the results in Burchfield et al. (2006). They find that moderate slopes induce expansive, scattered sprawl—an intensive margin effect as developers pursue easier-to-build and cheaper single-family homes rather than large apartment buildings—but high mountains act as hard barriers limiting the extensive margin. Our results suggest that while moderate topography restricts supply (raising prices), extreme slopes eventually limit agglomeration and market access, and lower the value of the housing stock.

The LU-ML Supply Index follows the same methodology as our Price Pressure Index but uses the log change in housing units as the target variable to measure $\mathbb{E}(\text{HU growth} \mid \text{LU})$. We use the LU-ML Supply Index to address a fundamental housing question: Do looser constraints dampen the house price effects of income growth? Examining existing supply constraint proxies, Louie et al. (2026) definitively conclude that the answer is "No," contrary to standard theory. They use the following specification across U.S. CBSAs:

$$\begin{aligned}\Delta \log(HP_i) = \alpha + \beta_1 \Delta \log(Income_i) + \beta_2 LessConstrained_i \\ + \beta_3[\Delta \log(Income_i) \times LessConstrained_i] + \varepsilon_i\end{aligned} \tag{2}$$

$LessConstrained_i$ is an indicator equal to 1 for CBSAs with looser-than-average supply constraints. For a valid supply constraint proxy, β_3 must be negative, meaning that looser supply constraints mitigate the house price effects of income growth.

Figure 2 plots the output for separate regressions using the BSH and Saiz proxies, as well as the LU-ML Supply Index. Data and implementation details are provided in the figure notes and generally follow Louie et al. (2026).

The first two columns of the figure replicate Louie et al. (2026): using the BSH and Saiz proxies, looser supply constraints fail to ameliorate the house price effects of income growth (Panels A and B). These proxies thus reject the standard narrative that supply elasticity dampens price appreciation with increasing demand.

The third column of Figure 2 shows the results using the LU-ML Supply Index. In marked contrast to other measures, the LU-ML Index reveals substantial heterogeneity in the income

elasticity of house prices. A 1% increase in income raises house prices by just 0.40% in less constrained cities versus 0.86% in more constrained cities. The difference in elasticities ($\beta_3 = -0.468$) is large in magnitude and statistically significant (Panel B). Indeed, looser supply constraints mitigate over 50% of the price impact of income growth found in more constrained cities (β_3/β_1), validating the LU-ML Indices as supply measures and highlighting the crucial role of supply restrictions in long-run affordability and housing market dynamics. Contributing to the LU-ML Index's improved identification of constraints is its low correlation with local housing demand: regressing our Supply Index on key demand factors yields an R^2 of just 0.079, compared with 0.20 and 0.28 for the BSH and Saiz proxies.[3]

Within a decadal panel framework, we also provide novel estimates of income elasticities spanning 1990 to 2020. These preferred specifications accommodate the cyclical housing fluctuations across decades that are often amplified by supply constraints (Case and Shiller, 2003; Glaeser et al., 2008; Mian and Sufi, 2009; Glaeser et al., 2014), while accounting for national macroeconomic conditions, unobserved heterogeneous city-specific trends, and other housing demand factors.

Less constrained cities, according to the LU-ML Supply Index, produce more housing units given average income growth and expand supply with rising demand. Increased supply then dampens prices, consistent with the preliminary evidence in Figure 2. These fundamental dynamics further validate the LU-ML Indices as supply constraints.

We also document that the effects of supply constraints are nonlinear and time-varying. We first find stark nonlinearities in mitigation effects. While looser constraints monotonically lead to increased housing production, only in the least constrained cities (the 5th quintile) does supply expansion statistically alleviate the price effects of demand growth. This result not only highlights the heterogeneous impacts of constraints across housing markets but also suggests that econometric specifications should extend beyond traditional binary classifications (Glaeser et al., 2006; Louie et al., 2026) or continuous linear specifications (Ganong and Shoag, 2017).

In terms of time-varying effects, we document a severe aggregate tightening of U.S. housing markets over the last three decades. While U.S. cities were relatively unconstrained in the 1990s,

[3]The local factors include the Lutz (2025b) natural amenity demand index and the 1990 values of the following variables: college share, median income, housing density, employment share in manufacturing, and employment share in computer services. The LU-ML Supply Index used in the regression targets the 1990–2020 long difference in CBSA housing unit growth.

geographic constraints became much more restrictive by the 2010s as the aggregate national supply curve shifted inward and steepened, with highly elastic local markets becoming exceedingly rare. Despite this aggregate tightening, a structural advantage persists: less constrained cities produced more housing units given average income growth in every decade from 1990 to 2020, congruent with a long-run secular growth premium. However, the differential ability of these markets to absorb increases in demand has fluctuated significantly. Less constrained markets exhibited only a modest additional supply response to income growth during the 1990s. This responsiveness surged during the 2000s boom before collapsing entirely in the 2010s, highlighting a severe weakening of the supply channel even in traditionally unconstrained areas (Glaeser and Gyourko, 2025). Consequently, it was only during the 2000s that the gap in income elasticities between more- and less-constrained cities was substantial enough to effectively ameliorate the price effects of rising demand.

Finally, we exploit the disaggregated nature of the LU-ML Supply Index to examine effects within cities from 1990–2020. In response to income growth, looser constraints facilitate housing supply expansion for both counties and zip codes. However, the localized price mitigation effects of this supply expansion only materialize for counties. This implies that within highly integrated, hyper-local housing markets, spatial arbitrage diffuses the localized price impacts of new supply.

In a novel application, we use the LU-ML Price Pressure Index to revisit the housing collateral channel for entrepreneurship. Exploiting the index's high spatial resolution, we implement a zip-code-level research design that controls for heterogeneous local trends and time-varying regional shocks—a level of spatial precision previously unavailable with coarser approaches.[4] We find that rising house prices significantly increase small business formation over the long run (1994–2022). However, this effect has evolved over time: it was lower during the 1990s but strengthened markedly during the 2000s housing boom. While it fell during the subsequent bust, the effect rebounded and strengthened even more so in the wake of the COVID-19 pandemic, suggesting that housing collateral has become increasingly important for entrepreneurs in the modern economy.

[4]See Chaney et al. (2012), Adelino et al. (2015), Bahaj et al. (2020), and Campello et al. (2022).

1 Data

The U.S. Geological Survey (USGS) provides the two primary datasets necessary to measure land unavailability due to steep slopes, water, and wetlands. The first is the USGS National Elevation Dataset (NED) 3DEP 1 arc-second Digital Elevation Model (DEM). While the original Saiz dataset relied on 3 arc-second DEM data with a resolution of approximately 90 meters, our data provide continuous coverage of the United States at a significantly finer resolution of approximately 30 meters. From the DEM data, we calculate the slope of the terrain to determine the percentage of unavailable land. Following Saiz, we classify any land with a slope exceeding 15% as undevelopable. Our second main dataset is the National Land Cover Database (NLCD) 2011, which uses Landsat imagery to classify land use. From this, we identify categories that constitute unavailable land, specifically water (oceans, lakes, rivers) and wetlands.

Other geospatial data, shapefiles, and census data are from the NHGIS (Schroeder et al., 2025) and TIGER/Line shapefiles (U.S. Census Bureau, a), and satellite imagery is from Google Maps. House price datasets are from Freddie Mac, the Federal Housing Finance Agency, and Zillow. The Bureau of Economic Analysis tabulates income at the CBSA level. We obtain zip-code-level business patterns from the U.S. Census Bureau (b) and geographic crosswalks from Lutz (2025a). Finally, journal repositories provide data for replicating Chaney et al. (2012), Mian and Sufi (2014), Stroebel and Vavra (2019), and Guren et al. (2021).

2 The Construction of the LU-ML Indices

The groundbreaking work of Saiz (2010) provides the foundation for this paper, as it was the first to use satellite imagery and GIS methods to compute proxies of LU. Saiz computes the share of land unavailable for construction within a 50 km radius around the centroid of each MSA's first central city using 1999 MSA/NECMA definitions. He combines the physical area of land undevelopable due to steep slopes, water, and wetlands to create a single proxy for each city. By pooling these distinct physical features into a single unweighted measure, Saiz effectively assumes that different types of geographic constraints are perfectly substitutable: that water or wetlands impact housing supply identically to steep slopes. This approach is at odds with Burchfield et al. (2006), for example, who find that sloped terrain has a larger impact on development patterns than

water or wetlands.

Figure 3, Panel A illustrates the limitations of the Saiz approach using Google satellite imagery for the Los Angeles-Long Beach-Anaheim CBSA. The blue outlined area marks the Los Angeles-Long Beach CBSA polygon boundary (2023 delineations) and the red circle represents a 50 km radius around the first central city centroid (in this case, the Los Angeles central city). Saiz uses the area within the red circle to calculate land unavailability. Clearly, the location of the first central city centroid determines the geography used in the Saiz calculation: the red circle in Panel A captures Los Angeles proper but fails to cover the eastern reaches of the CBSA, northern cities such as Lancaster and Palmdale (combined 2000 population over 230,000), and major economic hubs such as Anaheim and Irvine. In the greater LA area, these are the exact regions where much of the new construction occurs.

More generally, the Saiz circle creates systematic, spatially correlated measurement error. It tends to under-cover the spatially expansive CBSAs in the Southwest, while over-covering the more compact jurisdictions in the Northeast. Cities are also heterogeneous along other dimensions, including their polygon shapes and the location of their growth frontiers. By applying a uniform approach to diverse markets, the Saiz measure often mischaracterizes the relevant supply constraints facing developers.

To address these limitations, we construct new measurements of LU that scale to various levels of spatial aggregation including CBSAs, counties, zip codes, and census tracts. As differences in geographic boundaries present notable challenges for researchers, an immediate benefit of our approach is that we match the delineations found in other datasets. This removes a source of measurement uncertainty for researchers studying housing markets.

Panels B–D in Figure 3 show an example of our approach for the Los Angeles-Long Beach-Anaheim CBSA. We consider various buffered polygons around the first principal city (Panel B), buffers around the CBSA polygon (Panel C), and circles around the principal city centroid (Panel D).

In Panel B, we calculate land unavailability within the first principal city polygon (the most populous principal city; orange polygon), corresponding to the Los Angeles principal city. Then,

we buffer this polygon by 10%[5] (inner yellow polygon) and calculate land unavailability. From there, we sequentially increase the buffer size by 10pp until the buffer distance reaches 150% (outer yellow polygon) of the original first principal city polygon. The result is 16 nested polygons (the principal-city polygon plus 15 buffers) around the first principal city. These polygons allow us to directly capture land unavailability at the urban core of a CBSA.

We apply the same approach to the overall CBSA polygon (Figure 3, Panel C), yielding five additional polygons corresponding to the CBSA polygon with buffers ranging from 0% to 20% (in 5pp increments). The advantage of buffering the CBSA polygon is that it covers the entire administrative definition of the market and accounts for idiosyncrasies in the polygon's shape. Finally, in Panel D, we expand Saiz's original approach by considering multiple circles around the first principal city centroid instead of just within a 50 km radius circle: we exploit nine distinct circles with radii around the first principal city centroid ranging from 20 to 100 km (in 10 km increments). In the case of Los Angeles in Panel D, larger circles encompass both northern and southern LA, including the cities of Anaheim, Irvine, Lancaster, and Palmdale missed by the Saiz circle. Overall, Figure 3 exemplifies how we calculate land unavailability for 30 different polygons for each CBSA. Since we calculate the share of land unavailable due to three constraints (steep slopes, water, and wetlands) across these 30 distinct spatial footprints, we have 90 total LU features per CBSA.

The second step of our approach synthesizes these disaggregated land unavailability estimates using machine learning techniques. Specifically, we construct two complementary metrics: the LU-ML Price Pressure Index and the LU-ML Supply Index. While both LU-ML indices exploit the same methodology and LU estimates (steep slopes, water, and wetlands, measured at various levels of disaggregation), we build them from different outcome variables. The LU-ML Price Pressure Index targets house price growth, isolating how physical supply constraints push up prices, $\mathbb{E}(\text{HP growth} \mid \text{LU})$. The LU-ML Supply Index instead targets the log change in housing units, capturing the direct relationship between geographic barriers and the physical housing stock, $\mathbb{E}(\text{HU growth} \mid \text{LU})$.

The LU-ML framework also accommodates various empirical settings. Depending on the

[5]All buffer calculations use the EPSG:5070 NAD83/Conus Albers projection. A percentage buffer of $x\%$ corresponds to an absolute buffer distance equal to $x\%$ of the square root of the polygon's total area.

econometric specification, the algorithm generates either a time-invariant cross-sectional proxy for each geographic unit or a time-varying panel of LU-based predictors.

Here, we describe our methodology within a panel data setup for the LU-ML Price Pressure Index targeting CBSA YoY house price growth, but this framework easily extends to long differences, other spatial definitions, and the LU-ML Supply Index.

Since the economic impact of physical constraints may vary over time (e.g., constraints may bind more tightly during booms), we estimate the model separately for each time period using a repeated cross-fitting approach (Chernozhukov et al., 2018). We also train the model only using data outside of the local labor market—CBSAs for city-level analyses and Commuting Zones (CZs) for disaggregated areas like counties, zip codes, and census tracts—to ensure that the algorithm captures the effects of physical geography rather than idiosyncratic local supply and demand factors.

For each time period in the panel, we implement five repeats of five-fold cross-fitting. Within each fold, we train the XGBoost algorithm on a training set of labor markets and predict house price growth out-of-sample for the holdout (test) set. The average prediction across all repeats for a given geographic unit and time period constitutes the final index value. The XGBoost algorithm handles many correlated or weak predictors, unlike standard regression techniques, allowing the data to non-parametrically measure the expectation of price growth given geographic supply constraints. XGBoost also allows the LU-ML indices to capture rich interactions and nonlinearities that standard techniques miss.

Specifically, for each time period t in the panel dataset:

1. Randomly assign local labor markets (CBSAs or Commuting Zones) into five folds.
2. Designate one fold of labor markets as the holdout (test) set, while the remaining folds constitute the training set.
3. Train the XGBoost algorithm on the training data.[6]
4. Predict house price growth out-of-sample for the holdout set.

[6]To prevent overfitting, we partition the training set within each cross-fitting iteration to create an internal validation set (drawn randomly at the labor-market level to comprise 25% of the training data). We use this validation set to implement early stopping (XGBoost Documentation, 2025), which halts the addition of trees if the out-of-sample prediction error fails to improve for 25 consecutive rounds. The optimally tuned number of boosting iterations is then used to retrain the model on the full training set. While we use default values for all other hyperparameters, this data-driven constraint on model complexity, combined with ensembling predictions across five repeats of five-fold spatial cross-fitting, heavily restricts model variance and prevents the algorithm from fitting to idiosyncratic noise.

5. Repeat steps (2) to (4), rotating which fold serves as the holdout set.
6. Repeat steps (1) to (5) five times (five repeats) using different random spatial partitions of the local labor markets to minimize variance induced by data splitting.
7. Average the out-of-sample house price growth predictions across all repeats by geographic unit to generate the LU-ML Price Pressure Index for time period t.

Repeating this process for every period t yields a panel LU-ML Index measuring expected house price growth given geographic constraints. If the target variable is long-differenced, we generate a cross-sectional LU-ML Index by following steps (1) to (7) once.

By grouping our cross-fitting folds by local labor markets, we ensure that the out-of-sample predictions are strictly out-of-market. We likewise cluster standard errors at the same spatial boundaries (CBSAs or Commuting Zones) in all subsequent empirical specifications to account for within-market spatial correlation.[7]

Overall, the combination of high-dimensional spatial definitions and cross-fitting ensures the algorithm does not spuriously learn unintended local factors. Instead, the model's spatial features allow it to identify distinct topographic signatures and price them using generalizable rules learned entirely from other markets. The resulting LU-ML indices thus provide a geographic constraint measure purged of idiosyncratic local contaminants. Below we further show that the LU-ML indices are largely uncorrelated with amenities and local housing demand.

3 The LU-ML Price Pressure Index and House Prices

Figure 4, Panel A summarizes the explanatory power of the LU-ML Price Pressure Index relative to the BSH and Saiz supply constraint proxies using instrument relevance tests for house price growth. Specifically, we expand on the panel CBSA-level results in Figure 1 that used Freddie Mac data by aggregating results across various housing datasets (FHFA and Freddie Mac) at several levels of disaggregation (CBSA, county, zip code, and census tract) using Equation 1.

Panel A shows that existing supply constraint proxies have limited explanatory power for house prices. The BSH measure has negligible F-statistics across all estimation periods. The Saiz proxy

[7] As the LU-ML index is an out-of-sample prediction, using it in subsequent regressions introduces a "generated-regressor." Tests of the null hypothesis $\beta = 0$ remain valid but confidence intervals may understate uncertainty (Pagan, 1984; Murphy and Topel, 1985). We explicitly address this estimation uncertainty via spatial permutation tests that fully re-train the first-stage ML algorithm in every iteration, ensuring our standard errors are not artificially compressed (Appendix Figure B2 and Appendix Table B5).

has marginally better performance but consistently falls short of the $F = 50$ weak instrument threshold (black-dashed line) suggested by Keane and Neal (2024) and only noticeably clears the traditional $F = 10$ threshold during the 2000–2007 and 2000–2011 periods.

Unlike the BSH and Saiz proxies, the LU-ML Price Pressure Index exhibits robust explanatory power, with large F-statistics across specifications. These results thus document that geographic supply constraints do explain the variation in house price growth across cities, aligning with the canonical view that more restrictive supply constraints push up house price growth.

Panel B highlights another advantage of the LU-ML Index: its geographic coverage. The BSH and Saiz proxies often exclude 20–40% of available cross-sectional units, with samples skewed toward larger, more established metros. Empirical estimates using these proxies may thus suffer from sample selection bias. The LU-ML Indices, conversely, have nearly 100% coverage, only discarding Alaska and Hawaii.

Appendix Table A1 confirms that the LU-ML Price Pressure Index has more explanatory power than existing proxies across every house price dataset and spatial aggregation tested.

3.1 The Robustness of the LU-ML Index to Housing Demand

Davidoff (2016) challenges the validity of geographic supply proxies, arguing that the positive correlation between land unavailability and prices stems largely from household demand for natural amenities (e.g., water access, nice views), contaminating the interpretation of geographic restrictions as supply constraints. Using a multifaceted approach, we demonstrate that the LU-ML Indices reflect supply constraints rather than housing demand.

First, our preferred specification (Equation 1) addresses the core of the Davidoff critique. As the dependent variable is price growth ($\Delta \log HP$), time-invariant amenity values are differenced out. The inclusion of CBSA fixed effects (ψ_i) then absorbs heterogeneous city-level trends, including secular shifts in the value of local amenities and other local demand factors.

More broadly, as a supply constraint measure, the LU-ML Price Pressure Index should capture the cross-sectional price variation driven by physical building limitations while remaining largely unrelated to idiosyncratic local demand shocks. By design, the use of cross-fitting insulates the LU-ML Index from such shocks. We therefore expect the Index to exhibit systematic prediction errors in the tails of the house price growth distribution, as large demand shocks—such as a tech

boom or a manufacturing collapse—often dominate local price dynamics.

Figure 5, Panel A illustrates this by plotting the LU-ML out-of-sample (OOS) prediction errors (actual CBSA house price growth minus the LU-ML prediction) across CBSA house price growth bins from 2002–2024. In the middle of the distribution, where demand shocks are often moderate, the LU-ML prediction errors are small. The tails, however, yield sizable residuals. In the highest growth bin (far right), the prediction errors are large and positive: actual appreciation far exceeded what geographic constraints alone would predict, consistent with an outsized positive demand shock. Conversely, in the lowest growth bin, errors are large and negative, reflecting a negative demand shock pulling prices below their geography-implied levels.

Panel B explicitly links these prediction residuals to local housing demand. Specifically, it shows a binscatter of the LU-ML Index's out-of-sample prediction errors against CBSA total income growth (a standard proxy for local housing demand) from 2002–2023. The strong positive relationship ($t = 2.75$, $R^2 = 0.17$) implies that income growth pushes actual prices significantly away from the LU-ML predictions, consistent with the index capturing supply-based constraints rather than local demand.

Table 1 formally assesses the robustness of the relationship between the LU-ML Price Pressure Index and house price growth by directly controlling for local housing proxies in Equation 1. First, Column (1) shows the coefficient on the LU-ML Price Pressure Index, controlling only for CBSA and Census division-by-time fixed effects, over 2000–2023. As noted above, the LU-ML Index has strong explanatory power for the cross-section of house price growth ($t = 14.48$), with land unavailability-induced price pressure corresponding to higher price growth.

Column (2) controls broadly for local housing demand via income growth. As expected, higher income growth translates into higher house price growth ($t = 5.86$). However, the inclusion of income growth leads to nearly no attenuation in the coefficient on the LU-ML Index, confirming that it is largely uncorrelated with local housing demand.

We also control directly for natural amenities—the crux of Davidoff's land unavailability critique. The natural amenity index from Lutz (2025b) captures users' revealed preference for natural amenities via Google Maps reviews. Figure 6 first documents a natural amenity house price growth premium that varies with the housing cycle: cities with more desirable natural

amenities experience higher house price growth during housing expansions but lower growth during contractions, confirming Davidoff's supposition that demand for natural amenities has evolved over time.

To incorporate the static natural amenities index into our panel framework and also account for nonlinearities in the amenity premium, we interact quintiles of the natural amenity index with YoY national house price growth in Column (3). There is a clear gradient in the impact of natural amenities, even after controlling for total income growth, as more attractive amenities capitalize into higher house prices.

Using average YoY national house price growth (4.7%), the marginal effect of moving from the 1st to the 5th amenity quintile increases house price growth by 0.84pp ($t = 2.60$). This amenity premium is sizable, equivalent to roughly 2.7pp higher income growth per year.[8] Also, because the natural amenity index was measured using Google Maps reviews as of 2025, it is an *ex post* measure that inevitably captures the endogenous growth in amenity values driven by high-income sorting (Diamond, 2016). The index therefore likely overstates natural amenity demand, making our estimates for the LU-ML Index conservative in nature.

Despite the significance of natural amenities, the coefficient on the LU-ML Price Pressure Index in Column (3) moves only slightly from 0.661 to 0.642 and remains highly significant ($t = 14.27$). The LU-ML Index thus provides substantial supply-side information for house prices that is largely unrelated to amenities and housing demand.

Column (4) instead interacts the natural amenity quintiles with date fixed effects, allowing the impact of amenities across quintiles to vary non-parametrically over time. The results are similar.

Next, Columns (5) to (8) progressively control for the time-varying impacts of local factors—initial college share, median income, housing density, employment shares—through interactions with date fixed effects. College share and median income capture skill-sorting and endogenous amenity demand (Diamond, 2016); accounting for housing density ensures that the Price Pressure Index is not just revealing already developed cities; and employment shares capture industry trends, such as manufacturing declines or tech booms. Even though these controls are important housing

[8]The income growth equivalent is calculated by dividing the marginal effect of the amenity interaction ($0.177 \times 0.0473 \approx 0.0084$) by the coefficient on total income growth (0.315), yielding $0.0084/0.315 \approx 0.0266$, or roughly 2.7 percentage points.

market determinants and, through their interaction with date fixed effects, soak up a substantial amount of variation, the coefficient on the LU-ML Index remains robust at 0.508 ($t = 11.81$). Thus, even after including an omnibus of foundational housing market controls, the coefficient on the Price Pressure Index remains large and highly significant. Finally, Column (9) employs state-by-date fixed effects. While non-parametrically absorbing all time-varying state-level macro factors and regulatory shifts attenuates the coefficient, the LU-ML Index remains highly significant ($t = 6.48$).

Using the regression in Column (8) and the omitted variable bias sensitivity framework of Cinelli and Hazlett (2020), Figure 7 assesses the robustness of the LU-ML Index to unobserved local housing demand by simulating the impact of a hypothetical omitted variable relative to the explanatory power of a set of observed covariates. For example, Panel 1A explores the stability of the LU-ML coefficient relative to a potential omitted variable with explanatory power that is equivalent to, or a multiple of, the included natural amenities. The idea is simple. Table 1 accounts for natural amenities. Although these amenities have an economically meaningful and statistically significant impact on house price growth, there may also be a latent amenity factor that our array of controls misses. We thus want to assess the potential impact of such an omitted variable on the LU-ML coefficient, using the explanatory power of the observed amenity controls—their partial R^2 with both the treatment and the outcome—as a guide.

In Panel 1A of Figure 7, the 0x on the horizontal axis represents the unadjusted baseline, the coefficient on the LU-ML Price Pressure Index from Table 1, Column (8). The 1x on the horizontal axis simulates the impact of an omitted variable with explanatory power equal to the included natural amenities. The coefficient only falls slightly from 0.508 to 0.461 and remains highly significant ($t = 10.66$). The stability of the coefficient under the simulated latent variables stems from the relatively low explanatory power of natural amenities: using the partial R^2, they explain just 3.8% of the LU-ML Price Pressure Index and 2.9% of house price growth (Appendix Table A2).

The 2x and 3x on the horizontal axis examine the extreme cases where the omitted confounder has two or three times the explanatory power of the included natural amenity controls. Even under the severe scenario (3x), the coefficient on the LU-ML Price Pressure Index remains large at 0.362

with a t-statistic of 8.28.

Simulations based on the other housing demand factors—college share, median income, and computer services—in Panels 1B–2B likewise show that the coefficient on the Price Pressure Index is stable even under severe scenarios of hypothetical omitted variables. These results thus buttress the LU-ML Index as a supply constraint measure for house prices.

Beyond estimating the direct relationship, an extensive literature has used supply constraints as instrumental variables (IVs) for house prices to examine broader economic outcomes.[9] In such analyses, the validity of the instrument requires that supply constraints affect outcomes only through house prices (exclusion restriction) and are otherwise unrelated to unobserved determinants of those outcomes (exogeneity). Either violation contaminates the reduced form, with the instrument's reduced-form coefficient no longer reflecting the effect operating solely through house prices. Sensitivity analyses on the reduced form therefore probe both assumptions.

For example, in their study of housing wealth elasticities, Guren et al. (2021) note that the main threat to identification stems from natural amenity demand. Using the Price Pressure Index as the IV, we thus include quintiles of the Lutz (2025b) natural amenities index interacted with national house price growth in both their first-stage and reduced-form regressions. We then examine the sensitivity of the coefficient on the LU-ML Index to hypothetical omitted variables. Amenities appear to be a relevant predictor in the first stage ($t = 3.00$ for the top quintile), but Cinelli and Hazlett (2020) tests show that the coefficient on the LU-ML Index is stable even in the extreme scenario where a latent confounder has three times the explanatory power of the observed amenities (Appendix Figure A3). More crucially, in the reduced form, these amenities have no direct residual correlation with the outcome (all t-statistics < 1.31), meaning that the reduced-form coefficient on the LU-ML Index remains highly stable under severe scenarios of hypothetical omitted confounders. Together, these results support the identifying assumptions associated with using the LU-ML Index as an instrument in the Guren et al. (2021) design and its broader use in IV applications.

3.2 Decomposing the Geographic Drivers of Price Pressure

Having established that geographic constraints are important determinants of house price growth, we next open the "black box" to uncover the economic mechanisms driving these results. Specifically,

[9] See, for example, Figure 1 and footnote 2.

we investigate which physical constraints amplify prices. For example, steep slopes bind on the intensive margin by raising the marginal cost of density, whereas water boundaries constrain growth along the extensive margin as an urban border. Distinguishing between these mechanisms can clarify whether high prices stem from physical limits on a city's footprint or from the escalating marginal costs of densification.

We decompose the city-level LU-ML Price Pressure Index into its constituent geographic drivers using SHAP (SHapley Additive exPlanations) values (Shapley, 1953; Lundberg and Lee, 2017; Lundberg et al., 2020). This approach isolates the marginal contribution of each LU feature—steep slopes, water, and wetlands across various spatial definitions—to the final LU-ML Index. The SHAP value thus represents the specific growth premium (or discount) attributable to a single geographic feature. As derived by Shapley (1953) in the context of cooperative game theory, SHAP values are additive: the LU-ML Index for each city is the national average plus the sum of the marginal contributions of the geographic components.

Formally, let $\varphi_i^{(j)}$ denote the SHAP value for LU feature j in CBSA i: the marginal impact of feature j averaged over all possible coalitions of other features for city i. Here, j may represent a main effect or an interaction term.[10] We can decompose the LU-ML Price Pressure Index, the model's total predicted price growth for city i, as:

$$\text{LU-ML Index}_i = \mu + \sum_{j=1}^{M} \varphi_i^{(j)} \tag{3}$$

where μ represents the unconditional mean prediction across all CBSAs. A positive SHAP value ($\varphi_i^{(j)} > 0$) implies that LU feature j acts as a binding supply constraint for CBSA i, pushing its Price Pressure Index above the national baseline. Conversely, a negative SHAP value signals a looser-than-average supply constraint for feature j. This decomposition thus ranks geographic constraints by their economic magnitude, city-by-city.

Figure 8 decomposes the LU-ML Price Pressure Index—based on 2002–2024 Zillow house price growth—for representative cities. For each metro, we report the LU-ML Index, the excess price growth relative to the national baseline ($\sum \varphi_i^{(j)}$), and the prediction residual.[11] The arrows in the plot display aggregated SHAP values: the total marginal contribution of each LU type—steep slopes

[10] We isolate interaction effects using the SHAP interaction index (Lundberg et al., 2020).

[11] Because predictions are generated out-of-sample, $\hat{\mu}$ varies slightly across estimation folds, ranging from 130–136% with a mean of 134%.

(red), water (blue), and wetlands (teal)—to the LU-ML Price Pressure Index.[12] The arrow labels report the share of undevelopable land (averaged across spatial definitions) and the corresponding national percentile by LU type.

The figure highlights the substantial heterogeneity in how physical constraints manifest across different markets. Among these cities, Houston has the lowest Price Pressure Index (predicting 99% cumulative price growth over the sample), consistent with a geographic landscape that offers little resistance to supply. Its lack of steep-sloped terrain pulls price growth 22pp below the national average. In contrast, Salt Lake City has the highest Index (183%), where both rugged terrain (+31pp) and water barriers (+18pp) bind severely. San Francisco (LU-ML Index = 165%) and Miami (LU-ML Index = 146%) also exhibit high geographic pressure, though the specific drivers differ.

The SHAP values further reveal that geographic constraints are not monolithic; their impact is nonlinear and context-dependent. We observe three stylized facts. First, steep slopes are the dominant constraint for many markets. In flat cities like Houston, Miami, and Oklahoma City, the lack of topographical barriers acts as an elastic release valve, pulling the Price Pressure Index 10–22pp below the national average. Conversely, in Salt Lake City, rugged terrain (slope: 36%, 88th percentile) boosts predicted growth by over 30pp.

Second, the impact of slope on prices appears non-monotonic. San Francisco possesses significantly more LU due to slopes (48%) than Salt Lake City (36%), but the marginal contribution of slope to price growth is smaller (+9.7pp vs +31.1pp). This suggests a concave, non-monotonic relationship: while moderate constraints raise the marginal cost of density (increasing prices), extreme constraints may limit development so severely that the marginal pressure on prices diminishes. However, this may also reflect that in coastal cities like San Francisco, extreme water boundaries already severely restrict supply, reducing the marginal additional impact of topography. We examine this nonlinearity further below in Figure 10.

Third, water constraints appear to bind only at extremes. Consistent with Burchfield et al. (2006), who find that water boundaries generally do not constrain sprawl, Figure 8 documents that water exerts downward price pressure in Houston (−8.7pp) because, despite its relatively high rank

[12] We calculate the aggregate impact for each broad category by summing the SHAP values by LU type (e.g., Slope Impact$_i = \sum_{j \in \text{Slope}} \varphi_i^{(j)}$).

(6%, 82nd percentile), this coverage is far less constraining than the extreme coastal boundaries that elevate the national mean. This stark contrast with cities like San Francisco—where slightly higher water coverage (9%, 86th percentile) adds 20.9pp of price pressure—highlights the advantage of our LU-ML approach. Standard proxies that collapse geographic features into a single unweighted measure within a uniform radius (Saiz, 2010; Baum-Snow and Han, 2024) fail to capture such heterogeneity.

By exploiting 30 distinct spatial definitions, the XGBoost algorithm instead recognizes the spatial layout of the constraints. It captures that San Francisco's water forms a constrained peninsula choking off its urban core, whereas Houston's water is distributed primarily along its far southern periphery where it is far less restrictive to outward expansion. Importantly, because this index value is generated strictly out-of-sample via cross-fitting, the algorithm does not assign this premium by memorizing San Francisco's local market; rather, it identifies the severe price pressure of a peninsula by applying geographic rules learned from the spatial blueprints of other U.S. cities. More broadly, when water features reach extreme levels that often define hard urban boundaries—as in Miami, San Francisco, and Salt Lake City—they become binding supply constraints, adding 18–29pp to the LU-ML Index.

Finally, because the LU-ML Indices are generated out-of-sample, the residuals (Actual Growth − LU-ML Index) largely reflect demand-side factors rather than geographic supply constraints (Figure 5). Consequently, the large positive residuals in Miami (+70pp) and Houston (+37pp) likely stem from idiosyncratic demand shocks—such as COVID-era migration booms (Hudson, 2025) or recent affordability crises (Glaeser and Gyourko, 2025; Sherman et al., 2025)—beyond the price growth predicted by terrain alone.

To further illustrate why the LU-ML Index more accurately captures the impacts of constraints than existing proxies, consider the case of Chicago versus Las Vegas. Using the method in Saiz (2010) and Baum-Snow and Han (2024)—an unweighted sum across LU types (steep slopes, water, and wetlands) using a 50km radius circle around the principal city centroid—Chicago and Las Vegas have about an equal share of land unavailable for development at 39.5% and 38.2%, respectively.[13] Existing measures thus treat undevelopable land in Chicago and Las Vegas as the same.

[13]Using 2023 CBSA definitions for the Chicago-Naperville-Elgin, IL-IN and Las Vegas-Henderson-North Las Vegas, NV CBSAs.

But the nature of Land Unavailability differs markedly across the two cities. As shown via Google Maps imagery in Figure 9, Panel A, Chicago is constrained on one side by Lake Michigan but unconstrained elsewhere. Burchfield et al. (2006) find that water constraining a city on only one side does not impact development, as sprawl simply spreads in the other direction. In contrast, Las Vegas is completely surrounded by mountains (Panel B). Hence, even though Chicago and Las Vegas have the same unweighted total land unavailability, the source of the geographic constraints and how they bind differ markedly across the two cities. The Saiz and BSH proxies ignore these differences.

The LU-ML indices, however, use steep slopes, water, and wetlands across 30 spatial definitions for a total of 90 features for each CBSA. Combining this granularity with our ML approach resolves differences in constraints across cities. For example, Figure 9, Panel C plots the contribution of each land unavailability type to the LU-ML Index for Chicago and Las Vegas. For Chicago, its flat terrain pulls price growth nearly 24pp below the national mean. Chicago also has a high land unavailability share due to water (25%), placing it at the 96th percentile of all cities. Yet its undevelopable land due to water is not a binding constraint, as it contributes little to the LU-ML Index, consistent with a single-sided water boundary not impacting housing growth and thus prices. Likewise, Chicago's negative SHAP value for wetlands indicates that modest amounts of wetlands are not a binding constraint and are easily overcome with land improvement technologies (Dahl, 2011).

The bottom row of Panel C instead shows that the steep slopes surrounding Las Vegas constrain 28% of all land and push the LU-ML Index, $\mathbb{E}(\text{HP growth} \mid \text{LU})$, 22pp above the national average. In total, unlike existing proxies, the LU-ML Indices discern the differences in geographic constraints and the heterogeneity in how they affect housing markets.

3.3 The Structure of Price Pressure Across Cities

Figure 10 further decomposes the LU-ML Price Pressure Index to identify the specific features, functional forms, and spatial scales that drive geographic price pressures. Panel 1 plots the marginal contribution (aggregated SHAP values) against the LU share for each feature type. These plots confirm that the stylized facts discussed above extend to the full distribution of cities.

Panel 1A documents that the relationship between topography and price pressure is nonlinear.

In flat cities, the absence of slope exerts downward pressure on the LU-ML Index (20pp below the national mean), indicative of highly elastic housing supply. As terrain becomes rugged, the price pressure rises sharply. This is consistent with slopes acting as a binding constraint on the intensive margin: as Burchfield et al. (2006) note, moderate terrain ruggedness restricts dense construction and forces developers to leapfrog difficult topography, thereby increasing sprawl and infrastructure costs. However, the marginal contribution of steep slopes is non-monotonic: beyond an LU share of approximately 40%, the marginal impact diminishes and eventually reverses. This shape suggests that while moderate ruggedness constrains supply and raises prices, extreme topography may eventually sever market access or reduce agglomeration benefits, thereby dampening valuations.

Water, by contrast, exhibits a different pattern (Panel 1B). At low levels, the marginal contributions of water LU are widely dispersed, in line with Burchfield et al. (2006) who find that water often does not impact development patterns. However, high levels of water push prices upwards. This finding is consistent with water acting as a hard boundary to extensive margin growth. Panel 1C confirms that wetlands, conversely, act as a "soft constraint" with negligible price impacts. Moreover, congruent with land improvement technologies mitigating the constraints of wetlands (Dahl, 2011), moderate wetland coverage even exerts a slight downward pressure on prices relative to the national mean.

Panel 2 ranks the drivers of geographic price pressure by economic magnitude. In Panel 2A, we group features by LU type but also allow for interactions. This plot confirms that steep slopes are the dominant constraint, with an absolute impact of 20pp on the LU-ML Index for the average city relative to the national baseline. The second most important LU type is water (10pp). Interaction effects are also economically meaningful. The magnitudes of the slope $\times$ water interaction (6.8pp) and the slope $\times$ wetlands interaction (6.0pp) are each substantial, and combined they rival the main effect of water (10.0pp). These results empirically validate the use of XGBoost: a significant portion of geographic price pressure stems from the confluence of constraints, rather than just their additive sum.

Finally, Panel 2B identifies the relevant spatial scale of these constraints. We find that the 100km radius around the principal city centroid is the single most important spatial predictor (16.9pp). The predictive dominance of this broad measure—relative to the standard 50km radius used by

Saiz (2010) and Baum-Snow and Han (2024)—implies that modern housing supply constraints are fundamentally regional. This suggests previous literature may have systematically mischaracterized supply elasticity by restricting focus too strictly to the localized urban core. However, the non-trivial contributions of smaller buffers highlight that constraints bind simultaneously at the urban core and the periphery.

In sum, the LU-ML Price Pressure Index demonstrates that geographic supply constraints are neither linear nor strictly local. They operate through concave nonlinearities and threshold effects, compound via interactions, and bind broadly.

4 The LU-ML Supply Index and Income Elasticities

Existing supply proxies contradict the canonical prediction that looser constraints mitigate the price impacts of demand growth (Davidoff, 2013; Louie et al., 2026). To resolve this puzzle and further assess the relationship between housing markets and supply constraints, we employ the LU-ML Supply Index. As described in Section 2, our Supply Index shares the same high-resolution LU inputs and ML methodology as the Price Pressure Index but targets the log change in housing units, measuring $\mathbb{E}(\text{HU growth} \mid \text{LU})$. It therefore sorts housing markets based on the responsiveness of their physical housing stock to geographic constraints conditional on demand.

Using the LU-ML Supply Index in a long-differenced design from 2000–2020, Figure 2 provides preliminary evidence that the income elasticity of house prices is substantially lower in less constrained cities. These estimates use major CBSAs and Freddie Mac house prices but also extend to other house price datasets (e.g., FHFA and Zillow), wider spatial coverage using the comprehensive FHFA house price dataset, and county delineations (Appendix Figure B1 and Appendix Table B2).

This section expands on the results in Figure 2 using all CBSAs in a decadal panel from 1990–2020. As construction and house prices fluctuate within cities across decades, often amplified by supply constraints (Case and Shiller, 2003; Glaeser et al., 2008; Mian and Sufi, 2009; Glaeser et al., 2014), panels provide a tighter temporal alignment between income growth and housing market responses. Panels also enable more nuanced analyses of within-city, nonlinear, or time-varying dynamics, while increased degrees of freedom yield more precise estimates.

We employ the following baseline specification:

$$\begin{aligned}\Delta \log(Y_{it}) = &\alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 Constraint_{it} \\ &+ \beta_3[\Delta \log(Income_{it}) \times Constraint_{it}] + \varepsilon_{it}\end{aligned} \tag{4}$$

$\Delta \log(Y_{it})$ represents decadal growth in housing units or house prices for CBSA i, $\Delta \log(Income_{it})$ is decadal BEA total income growth, and α_i are CBSA fixed effects. Since the dependent variable is log-differenced, time-invariant factors (such as natural amenities) are differenced away and the CBSA fixed effects control for heterogeneous city-level trends, such as the persistent demand in "Superstar Cities" (Gyourko et al., 2013), the "Sand States" (Davidoff, 2013), or the Sunbelt (Glaeser and Gyourko, 2025). δ_t signifies decade (time) fixed effects and accounts for time-varying national shocks.

$Constraint_{it}$ is either the BSH, Saiz, or LU-ML supply constraint measure. For the static BSH and Saiz proxies, $Constraint_{it} = Constraint_i$ and its main effect (β_2) is absorbed by CBSA fixed effects (α_i) when included. To match the static nature of these proxies, we first train a time-invariant version of the LU-ML Supply Index on the long-difference in housing unit growth from 1990–2020. Below in Section 4.1, we target the decadal log change in housing units, allowing the LU-ML Supply Index to vary over time.

$Constraint_{it}$ can be an indicator equal to 1 for cities with looser-than-average supply constraints or the continuous version of each proxy, standardized to have zero mean and unit variance and signed such that larger values signal looser constraints. We prefer the continuous specifications as they leverage the full spatial variation in supply restrictions across cities.

We interpret β_3 as the "mitigation effect" for increased housing demand: do less constrained cities expand supply ($\beta_3 > 0$) and thus experience dampened prices ($\beta_3 < 0$) in response to elevated income growth.

Figure 11 plots estimates of β_3 from separate regressions of Equation 4 when the outcome is housing units (Panel 1) or FHFA house prices (Panel 2). Within each panel, the top estimates employ only time fixed effects, while the bottom estimates include both time and CBSA fixed effects.

The results for the BSH and Saiz indices echo Figure 2, implying that looser constraints *do not* mitigate the house price effects of income growth through expanded supply. Only the BSH index is

statistically significant with the correct sign when collapsed to a binary indicator for housing units (Panel 1A). For all other estimates, the existing proxies are statistically insignificant or produce the wrong sign.

In marked contrast to existing proxies, the LU-ML Supply Index consistently uncovers the canonical supply-side mitigation effect. For example, in the continuous specification with only time fixed effects, a one standard deviation loosening of constraints increases the income elasticity of housing units by 0.044 (Panel 1B). Compared to a baseline elasticity of 0.360 for a market with average constraints (Table 2), this mitigation effect implies that a one standard deviation loosening of constraints makes housing construction over 12% more responsive to demand growth ($(\beta_1 + \beta_3)/\beta_1$). The expanded supply response substantially offsets price appreciation, with a one standard deviation loosening of constraints lowering the corresponding elasticity for prices by 0.076 (Panel 2B).

Table 2 shows the full LU-ML regression results for the continuous specification both across CBSAs (using time fixed effects, Column 1) and within CBSAs (using time and CBSA fixed effects, Column 2). We demean income growth over the full BEA sample so that β_2 represents the secular growth premium: the excess long-run growth in housing units or house prices associated with a one standard deviation loosening of constraints, given average income growth. In Column (1), β_2 is positive for housing units (0.018) and slightly positive for house prices (0.005), reflecting the large and persistent demand for traditionally elastic markets (Glaeser and Gyourko, 2025). Observing a price growth premium for the Sand States (Arizona, California, Florida, and Nevada) led Davidoff (2013) to conclude that supply constraints were uncorrelated with prices during the 2000s, conditional on demand. Yet the estimates of β_3 in Table 2 instead show that more lax constraints accommodate new supply *and* alleviate price pressures, highlighting the large mitigation effects of looser constraints when local demand accelerates.

We control for persistent trends via CBSA fixed effects in Column (2). Intuitively, the CBSA fixed effects shift the analysis from a long-run comparison across cities to how cities differentially react to a deviation from their own local trend.

First, CBSA fixed effects reduce the baseline elasticity for housing units (β_1) from 0.360 to 0.084. So while the average city does expand its housing stock by 0.360% for a 1% increase in income

when comparing across cities, it produces far fewer new housing units when income accelerates over a single decade relative to the local trend. Since supply struggles to accommodate these short-run bouts of demand, the excess demand capitalizes directly into prices, pushing the baseline income elasticity from 0.631 to 0.736 after accounting for CBSA fixed effects. These results thus highlight the steep challenge of expanding supply to allay price growth in response to excess local demand.

In contrast, the total income elasticity of house prices for a city one standard deviation less constrained ($\beta_1 + \beta_3$) moves less with the inclusion of CBSA fixed effects, shifting from 0.555 to 0.610, even as its total income elasticity of housing units drops from 0.404 to 0.145. This suggests that accommodative supply responses help anchor the income elasticity of house prices for less constrained cities near their long-run mean, even when short-term demand spikes from its secular trend. Moving from Column (1) to (2) increases the magnitude of the mitigation effect for price growth (β_3) from -0.076 to -0.126. A purely cross-sectional comparison across cities (Davidoff, 2013; Louie et al., 2026) may hence underestimate how looser supply tempers price growth by conflating short-run price volatility with longer-run secular booms.

Returning to Figure 11, estimates of the LU-ML mitigation effect are robust across both the binary and continuous specifications. Yet while the precision for housing units remains similar across both models, the continuous specification strengthens the statistical precision of the mitigation effect (β_3) substantially for house prices (with the t-statistic increasing from -3.46 to -5.04 using time and CBSA fixed effects). This divergence coincides with a convex capitalization of costs and risk premiums into house prices (Paciorek, 2013; Han, 2013), suggesting that the price impacts of demand growth in constrained markets are highly nonlinear. We investigate these nonlinearities below by decomposing the mitigation effect across LU-ML Supply Index quintiles.

First, to further assess the robustness of these results, Appendix Table B4 employs alternative targets and specifications using the LU-ML Supply Index. Our findings are robust to using only 1990–2000 housing unit growth to train the LU-ML Supply Index, to winsorizing the LU-ML Supply Index at the 5th and 95th percentiles, and to using the median rather than the mean to define more versus less constrained cities in the binary specification. Targeting only 1990–2000 housing unit growth tests whether the previous results are driven by the housing boom of the 2000s or the lackluster growth in the 2010s, while winsorizing the continuous index or using a median

split for the binary indicator ensures that our results are not driven by outliers. In most cases, robustness results are stronger than those in Figure 11. Below in section 4.1, we further evaluate the robustness of the LU-ML Supply Index to various housing demand factors and the sensitivity of the mitigation effect to hypothetical omitted variables.

We also conduct a spatial permutation test for the mitigation effect (β_3 in Equation 4) to assess if our LU-ML machine learning approach captures the impacts of geographic constraints rather than complex spurious correlations learned by the algorithm. For each permutation of the data, we randomly shuffle the CBSA geographic identifiers. This permutation approach randomly reassigns all land unavailability features from one city to another. We then re-estimate the LU-ML Supply Index and Equation 4. As shown in Appendix Figure B2 and Appendix Table B5, our observed t-statistics fall entirely outside the empirical null distribution of the permuted samples ($p < 0.002$ based on 500 permutations). This collapse of predictive power under randomization confirms that the algorithm's performance relies on the relationship between local physical constraints and housing supply. Furthermore, the permutation test helps address the "generated regressor" problem (Pagan, 1984; Murphy and Topel, 1985). Because we fully re-train the first-stage ML algorithm in every iteration, the resulting empirical null distribution inherently incorporates algorithmic estimation uncertainty. Our observed t-statistics easily clear these permuted bounds, indicating that our inference on β_3 is robust to first-stage estimation uncertainty.

Next, Figure 12 estimates the secular growth premium and income elasticities by quintiles of the LU-ML Supply Index with the following specification:

$$\begin{aligned}\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \sum_{q=2}^{5} \beta_{2,q} LU\text{-}ML_{iq} \\ + \sum_{q=2}^{5} \beta_{3,q} [\Delta \log(Income_{it}) \times LU\text{-}ML_{iq}] + \varepsilon_{it}\end{aligned} \tag{5}$$

where $LU\text{-}ML_{iq}$ is an indicator equal to 1 if CBSA i falls into the q-th quintile of the LU-ML Supply Index. The most constrained cities in quintile 1 represent the reference category (omitted dummy).

Panel 1 plots the marginal secular growth premium relative to the most constrained cities in quintile 1 ($\beta_{2,q}$) when using only time fixed effects. For housing units (Panel 1A), the secular growth premium monotonically increases as geographic constraints soften, with the housing stock

in the least constrained cities (quintile 5) growing 5.2pp ($t = 13.0$) faster per decade than in the most constrained cities, conditional on average income growth. More broadly, the upward trend in the estimates for the excess housing unit growth in Panel 1A across LU-ML quintiles supports our approach, showing that our out-of-sample estimation of $\mathbb{E}(\text{HU growth} \mid \text{LU})$ captures the distribution of housing unit growth across cities. Our subsequent price results, which the algorithm never targets, then further validate the Supply Index as a constraint measure.

Panel 1B displays the secular growth premium for house prices. Unlike physical housing units, baseline price growth does not exhibit a monotonic relationship across constraints.

Panel 2 plots the total income elasticity by quintile ($\beta_1 + \beta_{3,q}$). The income elasticity of housing units (Panel 2A) increases with LU-ML quintile and hence highlights how looser geographic constraints facilitate new supply across the distribution of cities. The estimates are large: the income elasticity increases by 0.185 ($t = 4.51$) moving from quintile 1 to 5 (Time FE only), meaning that supply expands 73% more in response to income growth in quintile 5 than in quintile 1. The difference across the most and least constrained quintiles is even larger when we also control for CBSA fixed effects, with $\beta_{3,q=5} = 0.209$ ($t = 6.33$).

Panel 2B shows the income elasticity of house prices by LU-ML quintile. The results are highly nonlinear. The difference in income elasticities between quintiles 1 and 4 inclusive of CBSA effects, the mitigation effect $\beta_{3,q=4}$, is just -0.161 ($t = -1.81$). In comparison, the mitigation for cities in quintile 5 (least constrained) relative to quintile 1 is -0.344 ($t = -4.00$). Hence, loose geographic constraints only ameliorate the price effects of demand growth at the extreme.

Appendix Table B7 shows a robustness check where we first residualize the LU-ML Supply Index to various local housing demand factors to ensure that the binning into quintiles is not confounded by demand.[14] The R^2 from the residualization regression is just 0.079, documenting the low correlation between the LU-ML Supply Index and housing demand. In contrast, the comparable R^2 statistics for the BSH and Saiz proxies are 0.20 and 0.28, indicating that these existing measures are confounded by housing demand. Ultimately, the nonlinear mitigation effects we document above remain largely unchanged with the demand-residualized LU-ML Supply Index.

[14]The local factors measured at the CBSA level are the Lutz (2025b) natural amenity demand index and the 1990 values of the following variables: college share, median income, housing density, employment share in manufacturing, and employment share in computer services.

Finally, from an econometric perspective, the nonlinear reduction in income elasticities with LU-ML quintiles suggests that regressions examining supply constraints should extend beyond simple binary classifications (Glaeser et al., 2006; Louie et al., 2026) or continuous linear specifications (Ganong and Shoag, 2017).

4.1 The Time-Varying LU-ML Supply Index

Thus far, we have used the growth in housing units from 1990–2020 (long difference) as the target variable for the LU-ML Supply Index. But our LU-ML approach is highly flexible. We can thus construct a time-varying LU-ML Supply Index by targeting the decadal log change in housing units, measuring $\mathbb{E}(\text{HU growth}_{i,t} \mid \text{LU})$. Unlike existing, static supply proxies, the time-varying LU-ML Supply Index reflects how constraints bind over time. This better captures the evolution of U.S. housing markets.

Table 3 first shows estimates of average income elasticities from 1990–2020 (Equation 4) using the time-varying LU-ML Supply Index. We standardize the index over the pooled sample to have zero mean and unit variance, such that higher values indicate looser effective constraints. Column (1) uses only time fixed effects. The secular growth premium in less constrained cities is positive for both housing units (Panel A) and house prices (Panel B), indicating that both housing units and prices grow faster over the long run in less constrained cities conditional on income.

A key advantage of the time-varying LU-ML Supply Index is that we can identify β_2 entirely from within-city variation and thus control for heterogeneous city-level trends, such as persistent housing demand. Column (2) does so by including CBSA fixed effects. β_2 thus no longer measures the cross-sectional secular growth premium. Instead, it measures the transitory growth premium—a city's deviation from its secular growth trend when constraints bind less severely than the long-run norm. β_2 for housing units persists and is nearly identical to Column (1). This suggests that when geographic constraints bind less *within* a city relative to its long-run trend, construction expands at the same rate as when comparing *across* cities. In other words, loosening geographic constraints yields the same proportional growth in housing units regardless of whether that loosening is measured across space or over time, highlighting the robust role of geographic constraints as a driver of housing unit growth.

Conversely, the β_2 for house prices falls to a statistical zero ($\beta_2 = 0.006$; $t = 0.75$) when

we include CBSA fixed effects in Panel B, Column (2). Given the same income growth, the higher long-run price appreciation in less constrained markets (Davidoff, 2013) stems entirely from persistent, city-level trends.

As in Table 2, controlling for CBSA fixed effects does not attenuate the mitigation effect (Column 2) for either quantities or prices, but rather strengthens it. Hence, within-city comparisons increase the response differential to income growth across constrained and unconstrained markets.

Column (3) further controls for various CBSA-level factors interacted with decade fixed effects. These include natural amenity quintiles and the 1990 levels of college share, median income, housing unit density, and local industry employment shares.[15] These time-varying controls account for notable dynamics that have impacted housing markets in recent decades, including changing amenity demand (Davidoff, 2016), endogenous skill-based sorting (Diamond, 2016), and manufacturing declines (Pierce and Schott, 2016). Note that Column (3) retains the CBSA fixed effects. The resulting elasticities thus measure the impact within cities in excess of time-varying, city-specific factors and local trends, and how this impact varies by geographic constraints.

Adding CBSA controls increases the baseline elasticity for housing units by 50% (to 0.105) and for house prices by 7% (to 0.811). Our previous results were thus conservative and biased downwards, especially for housing units. In other words, baseline city-specific factors mask a substantial portion of the response of housing supply, even when holding geographic constraints constant.

Nonetheless, β_3 for housing units remains remarkably stable, documenting that the additional controls do not alter the fundamental mitigation mechanism. Looser supply constraints facilitate substantial supply expansion when income rises, with an elasticity differential of 0.075 ($t = 8.33$) that makes cities one standard deviation less constrained 71% more responsive to income growth than the average CBSA.

Conversely, the mitigation effect for prices attenuates from -0.177 to -0.113. This attenuation directly highlights the role of city-specific factors in the capitalization of income growth across constrained and unconstrained cities. While housing units reflect physical construction, house prices behave as forward-looking assets that reflect demand composition, amenities, and skill-

[15]The included industries span computer services; construction; finance, insurance, and real estate; manufacturing; retail; and services (excluding computer services).

sorting (Diamond, 2016; Davidoff, 2016). Yet even after purging these variables, looser constraints significantly mitigate the price effects of income growth ($t = -4.37$), confirming that the LU-ML Supply Index captures a physical constraint mechanism distinct from shifts in unobserved local housing demand.

Figure 13 applies the Cinelli and Hazlett (2020) omitted variable bias sensitivity framework to test whether the mitigation effect in Column (3) is robust to unobserved local demand. For each of our key observed demand factors (natural amenities, college share, median income, and computer services employment share), we simulate how much the mitigation effect would attenuate if there existed a hypothetical omitted confounder with up to 3 times the explanatory power of that covariate.

For example, Panel 1A considers the stability of the mitigation effect for housing units relative to natural amenities. The baseline coefficient, from Column (3) in Panel A of Table 3, is 0.075 (0x on the horizontal axis). Intuitively, the Cinelli and Hazlett (2020) test assesses a scenario where an omitted confounder has equal explanatory power to natural amenities, meaning it has the same partial R^2 with both the treatment and the outcome as the observed natural amenities (1x on the horizontal axis). Under this scenario, the mitigation effect moves little, only dropping slightly from 0.075 to 0.068, demonstrating that unobserved local demand shocks (Davidoff, 2016) of this magnitude have little impact on our estimates. Even in extreme cases where the omitted confounder has double or triple the explanatory power (2x or 3x on the horizontal axis), the mitigation effect only falls modestly to 0.054. Mechanically, this stability stems from the low explanatory power of natural amenities in the original regression, as they explain 2.2% (partial R^2) of both the income growth–LU-ML Supply Index interaction and housing unit growth.

Panels 1B to 1D display these coefficient stability tests for the other main local demand proxies. The results likewise show that the mitigation effect for housing units is highly robust to even extreme hypothetical unobserved confounders. A potential omitted variable with up to 3 times the impact of initial college share—the observed covariate with the highest explanatory power of the interaction term—only reduces the mitigation effect to 0.057, an estimate that is highly significant and large in magnitude.

Panel 2 shows that the mitigation effect for house prices is broadly robust to unobserved local

housing demand. For natural amenities, median income, and computer services employment, the estimates remain statistically significant, even after accounting for hypothetical omitted variables with triple the explanatory power of these observed covariates.

Initial college share is the strongest observed confounder in our model and a foundational driver of skill-based sorting and local housing demand (Diamond, 2016). Even accounting for an additional omitted confounder with explanatory power equal to initial college share (1x on the horizontal axis) only attenuates the mitigation effect from -0.113 to -0.088, leaving the estimate highly significant ($t = -3.35$). Doubling the explanatory power of this latent confounder to twice that of initial college share still leaves a significant coefficient ($t = -2.30$). Only in the extreme case where an omitted variable has three times the explanatory power of initial college share does the mitigation effect become statistically insignificant, though it retains the expected negative sign. Given the fundamental importance of college share to housing demand, an omitted confounder of this magnitude represents an exceptionally severe and conservative bounding scenario. Altogether, these coefficient stability tests reinforce that looser geographic constraints mitigate the price effects of rising demand.

Lastly, Column (4) in Table 3 adds state-by-time fixed effects. The estimates of β_3 attenuate only slightly, meaning that the mitigating effects of looser supply constraints on local housing outcomes are robust to time-varying state-level shocks.

Next, we examine the effects of constraints within cities using the county- and zip-code-level LU-ML Supply Indices. Figure 14 plots the intra-city mitigation effect using the LU-ML Index, alongside the Saiz and BSH proxies, by estimating Equation 4 using county- or zip-code-level data.[16] The top row in each panel shows estimates using CBSA fixed effects, yielding a comparison across counties or zip codes within each city. The bottom row instead controls for local, geographic fixed effects (county or zip code) to account for differential local trends and thus compares each county or zip code to its long-run growth.

For counties, Panels 1A and 2A first document that only the LU-ML Index captures the correct theoretical direction of the mitigation effect for both housing units and prices. These results are large and statistically significant and imply that even within cities, counties act as distinct,

[16]We omit the Saiz proxy from zip-code-level specifications as it cannot be meaningfully disaggregated to those delineations.

internally integrated housing markets where looser constraints alleviate the price pressures of demand growth. This finding is congruent with crucial public goods (e.g., school districts and tax regimes) varying markedly between counties even within the same city, creating rigidities in household mobility.

For zip codes, we plot the mitigation effect for the LU-ML Supply Index and the BSH elasticity proxy. The two measures produce a similar mitigation effect for housing units (Panel 1B) when comparing across zip codes (CBSA + Time FE). But the mitigation effect for housing units using BSH falls to zero when including zip code fixed effects, implying that the BSH proxy captures local trends rather than the supply response to income growth. Conversely, the LU-ML Index remains positive and significant, revealing that when effective constraints bind less severely, construction expands locally in response to rising demand, even after accounting for heterogeneous neighborhood-level trends.

Panel 2B, however, documents that less restrictive zip-code-level constraints have no impact on local house prices. In contrast to counties, households can easily substitute between zip codes, so that spatial arbitrage diffuses the price-dampening benefits of new supply across neighborhoods.

Lastly, Figure 15 lets the estimates of income elasticities and the baseline growth premiums from Equation 4 vary across decades. As above, we standardize the LU-ML Supply Index over the pooled sample to have zero mean and unit variance so that a one standard deviation increase in the LU-ML Index represents the same absolute change in effective constraints across decades.

For housing units in Panel 1A, the baseline elasticity ($\beta_{1,t}$) has fallen noticeably over time. Evaluated at the pooled historical average of the LU-ML Supply Index (a value of zero) and comparing across cities using only time fixed effects (red), a 1% increase in income corresponded to 0.48% more housing units in the 1990s but only 0.21% in the 2010s. The sizable difference of 0.27 is statistically significant ($t = 5.35$) and highlights how U.S. housing markets have become less responsive to income growth overall (Glaeser and Gyourko, 2025). Yet Panel 1B also shows the stability of the growth premium, both across and within cities, over time: conditional on income growth, less constrained markets exhibited a positive supply growth premium in every decade. Finally, Panel 1C reveals that the differential supply response in less constrained markets (the mitigation effect for housing units) fluctuated significantly, surging during the 2000s boom before

collapsing in the 2010s.

To reconcile Panels 1A and 1B—that supply responsiveness has declined conditional on constraints, while supply has persistently grown at a faster rate in less constrained markets given income—note that aggregate housing supply has shifted markedly. Figure 16 shows a severe leftward shift and compression of the distribution of the LU-ML Supply Index over time.

Since the time-varying Supply Index targets the expectation of decadal housing unit growth given land unavailability, $\mathbb{E}(\text{HU growth}_{i,t} \mid \text{LU})$, it captures two distinct economic phenomena. First, the leftward shift in the distribution reflects a macroeconomic decline in aggregate supply responsiveness (Glaeser and Gyourko, 2025). The magnitude of this effective tightening is stark: the mean of the LU-ML Supply Index (standardized over the pooled sample) fell from 0.63 in the 1990s to -0.94 in the 2010s. Moreover, the price capitalization of income growth simultaneously increased (across cities, the baseline capitalization elasticity rose from 0.40 in the 1990s to 0.88 in the 2010s; Panel 2A), indicating that this decline in quantities reflects a contraction in supply rather than weakening demand. Note that our time-varying regressions in Figure 15 account for this aggregate change, with decade-specific baseline elasticities and time fixed effects.

Second, the pronounced compression of the distribution directly reveals the changing impact of supply constraints. As the variance in the Supply Index narrowed, the gap in supply responsiveness between unconstrained and constrained markets shrank. A one standard deviation loosening of effective constraints for the pooled sample represented a move from the 25th percentile to the 75th percentile in the 1990s, but that same absolute shift required a leap from the 25th to the 97th percentile in the 2010s.[17] This compression surfaces in Panel 1C: as geography's impact on construction weakened, the additional supply response for less constrained markets vanished.

Thus, although the structural growth advantage of physically unconstrained land persists (Panel 1B), aggregate supply responsiveness has diminished (Panel 1A), demand growth is increasingly capitalized into house prices (Panel 2A), and highly elastic local markets have become exceedingly rare (Figure 16). In macroeconomic terms, the leftward shift and narrowing of the constraints distribution corresponds to an inward shift and steepening of the aggregate national housing supply

[17] Figure 16 displays the unweighted distribution across CBSAs. Appendix Figure B3 shows the densities weighted by 1990 housing units. Weighting by initial housing stock places more mass in the constrained (left) tail during the 1990s and 2000s, reflecting the historical constraints of large coastal markets. Yet the same aggregate tightening and compression trends persist through the 2010s.

curve.

Consequently, looser supply constraints failed to mitigate the price impacts of demand growth in the 2010s (Panel 2C). Indeed, only during the 2000s was the differential supply expansion in looser markets large enough to significantly allay the price effects of income growth.

Lastly, note that Panel 2B shows that the house price growth premium documented in Table 3 is a modern phenomenon. In the 1990s, given average income growth, prices in less constrained cities grew more slowly. This trend reversed in the 2000s, especially in cross-city comparisons (red), as prices in less constrained cities grew faster conditional on demand (Davidoff, 2013). In the 2010s, the growth premium accelerated amid rising demand for traditionally elastic markets (Glaeser and Gyourko, 2025).

Overall, Figure 15 summarizes the evolution of U.S. housing markets since 1990 and reconciles the apparent discordance in the housing and urban literature. In the 1990s, the U.S. was relatively unconstrained: for a city with average constraints, income growth spurred new construction that dampened price appreciation (Glaeser et al., 2014). During this period, conditional on demand, less constrained cities produced more housing units and experienced slower price growth. The massive 2000s boom then pushed up both price and housing unit growth in less constrained cities, leading to a surprising price growth premium in these housing markets compared to their more constrained counterparts. Viewed in isolation, the premiums in Panels 1B and 2B suggest that supply constraints had no impact on housing markets during the 2000s (Davidoff, 2013). Yet the 2000s-era mitigation effect was extremely large, both within and across cities, highlighting the vital role of supply constraints in alleviating price pressures during this period. By the 2010s, the distribution of geographic supply constraints across cities narrowed and moved leftward as construction in traditionally elastic markets slowed (Glaeser and Gyourko, 2025), with profound consequences for the price capitalization of income and the affordability of U.S. cities. Finally, the dynamic evolution of constraints in U.S. housing markets cautions against relying solely on static, time-invariant measures to assess the impact of housing supply (Saiz, 2010; Baum-Snow and Han, 2024; Louie et al., 2026).

5 Application: The Housing Collateral Channel and Entrepreneurship

In a novel application, we use the LU-ML Price Pressure Index to revisit a central question in macro-finance: Does rising housing wealth fuel entrepreneurship? Theory suggests that higher collateral values relax borrowing constraints, enabling small business formation (Adelino et al., 2015).[18] However, causally identifying this channel is challenging. Omitted local factors like gentrification or agglomeration shocks can simultaneously drive house prices and firm entry. Furthermore, firm formation itself can raise local prices (reverse causality), while measurement error in price indices biases naive OLS estimates.

We overcome this identification challenge by leveraging the high spatial resolution of the LU-ML Price Pressure Index. Unlike coarse county-level instruments, our zip-code-level proxy allows us to control for granular fixed effects and trends. We estimate the following 2SLS specification using zip-code-level data from the ZIP Code Business Patterns for 1994–2022:

$$\Delta \log(\text{Firms}_{i,d,t}^{1-4}) = \alpha_i + \delta_{d,t} + \beta \cdot \Delta \log(\text{HP}_{i,d,t}) + \varepsilon_{i,d,t} \tag{6}$$

where $\Delta \log(\text{Firms}_{i,d,t}^{1-4})$ is the year-over-year log change in the number of firms with 1–4 employees in zip code i and Census division d at year t. Crucially, time-invariant differences across zip codes are differenced away via the panel structure, and the inclusion of zip code fixed effects (α_i) in our growth equation accounts for heterogeneous zip-code-level trends. Further, the division-year fixed effects ($\delta_{d,t}$) non-parametrically control for time-varying regional shocks. We instrument house price growth ($\Delta \log(\text{HP}_{i,d,t})$) using the time-varying LU-ML Price Pressure Index at the zip code level. The first-stage Kleibergen-Paap F-statistic (using standard errors clustered at the Commuting Zone level) is 164.73, confirming that local supply constraints are powerful predictors of prices even at this high level of disaggregation. At first glance, this strong first stage using the Price Pressure Index appears at odds with the diffusion of prices across zip codes using the Supply Index in Figure 14. Yet this highlights the core advantage of our dual index approach. Over the short run, the housing stock is largely stagnant; elevated demand and expectations capitalize into higher prices in constrained zip codes, expanding the collateral of incumbent homeowners. It is only over the long run that mobile homebuyers can spatially arbitrage new construction, driving

[18] See also Corradin and Popov (2015); Bahaj et al. (2020); Campello et al. (2022); Haltiwanger (2022); Jensen et al. (2022).

the price-mitigation effects captured by the Supply Index.

Figure 17 presents the 2SLS results. Panel A documents a strong causal link between collateral values and micro-entrepreneurship. We find a house price elasticity of 0.113 for firms with 1–4 employees, significant at the 1% level. These results remain highly robust even when using state-by-year fixed effects to control for time-varying, state-level macroeconomic and policy shocks (Appendix Table C1).

To validate that this effect is driven by credit constraints rather than general local demand, we run a placebo test using slightly larger firms (5–9 employees). Consistent with the collateral channel hypothesis—that home equity matters most for the smallest, most capital-constrained entrants—the elasticity for 5–9 employee firms is statistically indistinguishable from zero.[19] The difference between the two estimates is large and statistically significant, confirming that the housing wealth effect is concentrated at the extensive margin of small firm formation. Of course, entrepreneurs may use their home equity to establish businesses outside of their residential zip code. This geographic mismatch is particularly relevant given our zip-code-level Price Pressure instrument: zip codes with extreme physical land constraints (high LU-ML) often lack dense commercial zoning, perhaps incentivizing founders to start their firms in neighboring commercial areas. Because our specification only captures firm formation within the treated zip code, such spatial spillovers likely attenuate our results, meaning that the measured elasticity corresponds to a conservative estimate of the true housing wealth effect.

Panel B investigates the dynamics of the housing collateral–entrepreneurship relationship over the last three decades. We find that the elasticity of entrepreneurship has evolved significantly. The collateral channel was statistically insignificant with wide confidence intervals during the 1990s, likely due to relatively modest cross-sectional differences during that period. The impact of housing collateral on entrepreneurship then spiked during the 2000s boom—congruent with Adelino et al. (2015)—before falling during the bust. Our estimates for the 2000s are the most precise over our sample due to the large cross-sectional and dynamic changes in house prices over that period. The

[19] A potential concern with our zip-code data is that firms with 5–9 employees may require more commercial space, leading to spatial leakage across zip-code boundaries compared to 1–4 employee firms. However, using broader county data that should largely absorb such spatial spillovers, Adelino et al. (2015, Table 2, Column 5) document the same attenuation in the housing wealth effect for 5–9 employee firms, congruent with our results for 5–9 employee firms reflecting limited housing collateral constraints.

elasticity hovered below 0.1 from 2013–2018 but was statistically insignificant, before elevating to 0.17 during the pandemic era (2019–2022). This recent structural shift suggests that housing collateral has become an increasingly binding constraint for entrepreneurs in the modern economy, consistent with rising home equity providing a financial runway for entrepreneurs in the wake of COVID-19.

6 Conclusion

This paper revisits the fundamental relationship between physical geography and housing markets. By combining high-resolution satellite imagery with machine learning, we construct LU-ML Indices that capture the heterogeneity, nonlinearities, and spatial layout of physical barriers at a highly granular level. The LU-ML Price Pressure Index captures the capitalization of geographic constraints into house prices, while the Supply Index reveals how geographic constraints restrict construction.

Our analysis yields four primary insights. First, we establish that physical constraints are a key determinant of cross-sectional price dynamics. Unlike existing proxies, the LU-ML Price Pressure Index has strong explanatory power for the differences in house price growth across cities. This explanatory power is robust to local housing demand, highlighting the key role of supply in housing markets.

Further decomposing the Price Pressure Index reveals that constraints on the intensive margin—specifically steep slopes that sharply raise the marginal cost of density—are the dominant driver of price pressure. We also uncover critical interactions between slope and water, as well as nonlinearities: while moderate ruggedness steepens the supply curve and raises prices, extreme topography acts as a hard barrier that eventually hampers agglomeration and dampens price growth, creating a non-monotonic relationship missed by linear models.

Second, we validate the theoretical prediction that looser constraints actively mitigate the price effects of demand growth by facilitating supply expansion. Recent literature (e.g., Louie et al., 2026) uses standard proxies to conclude that unconstrained markets fail to dampen the price effects of income shocks. We show that this "null result" is an artifact of previous supply constraint proxies. Using the LU-ML Supply Index, we document a statistically significant and economically large divergence in elasticities: less constrained cities mitigate the price impacts of demand growth

by facilitating new construction.

Third, our decadal and highly disaggregated analyses reveal that this affordability channel is highly nonlinear, time-varying, and scale-dependent. Price mitigation effectively materializes only in the least constrained quintile of cities. Furthermore, we document a severe aggregate tightening of U.S. housing markets over the last three decades. While the supply responsiveness of less constrained markets surged during the 2000s boom, this responsiveness collapsed entirely in the 2010s, weakening the affordability channel nationwide. Within cities, we also highlight the role of spatial arbitrage: while looser constraints facilitate new construction at both the county and hyper-local zip-code levels, localized price mitigation only materializes for counties, as demand diffuses across neighborhoods when examining the microgeography of supply constraints.

Finally, the unprecedented spatial resolution of the LU-ML approach unlocks new identification strategies for applied research. In a novel application to the housing collateral channel for entrepreneurship, we exploit the Price Pressure Index's zip-code-level variation to control for heterogeneous local trends and regional shocks—a level of precision unattainable with coarser approaches. We find that the elasticity of small business formation to housing wealth was reduced during the 1990s, strengthened markedly during the 2000s boom, fell during the subsequent bust, and surged following the COVID-19 pandemic, suggesting an increasing reliance on home equity for small business formation in the modern economy.

These findings have important implications for urban economic theory, empirical research, and housing policy. By accurately quantifying the price and supply impacts of physical geography across multiple spatial scales and time periods, the LU-ML Indices allow for a more precise measurement of market frictions. As policymakers increasingly target land-use reform to combat the affordability crisis, the ability to cleanly distinguish between the immutable constraints of physical geography and the malleable constraints of regulation will be essential for designing effective interventions.

References

M. Adelino, A. Schoar, and F. Severino. House prices, collateral, and self-employment. *Journal of Financial Economics*, 117(2):288–306, 2015.

A. Aladangady. Housing wealth and consumption: Evidence from geographically linked microdata. *American Economic Review*, 107(11):3415–3446, 2017.

S. Bahaj, A. Foulis, and G. Pinter. Home values and firm behavior. *American Economic Review*, 110(7):2225–2270, 2020.

M. Baldauf, L. Garlappi, and C. Yannelis. Does climate change affect real estate prices? only if you believe in it. *The Review of Financial Studies*, 33(3):1256–1295, 2020.

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

M. Burchfield, H. G. Overman, D. Puga, and M. A. Turner. Causes of sprawl: A portrait from space. *The Quarterly Journal of Economics*, 121(2):587–633, 2006.

Bureau of Economic Analysis. Personal Income. URL https://www.bea.gov/data/income-saving/personal-income.

M. Campello, R. A. Connolly, G. Kankanhalli, and E. Steiner. Do real estate values boost corporate borrowing? evidence from contract-level data. *Journal of Financial Economics*, 144(2):611–644, 2022.

K. E. Case and R. J. Shiller. Is there a bubble in the housing market? *Brookings papers on economic activity*, 2003(2):299–362, 2003.

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

T. Chaney, D. Sraer, and D. Thesmar. The collateral channel: How real estate shocks affect corporate investment. *American Economic Review*, 102(6):2381–2409, 2012.

T. Chen and C. Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, page 785–794, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450342322. doi: 10.1145/2939672.2939785. URL https://doi.org/10.1145/2939672.2939785.

I.-H. Cheng, S. Raina, and W. Xiong. Wall street and the housing bubble. *American Economic Review*, 104(9):2797–2829, 2014.

V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018. ISSN 1368-4221. doi: 10.1111/ectj.12097. URL https://doi.org/10.1111/ectj.12097.

R. Chetty, L. Sándor, and A. Szeidl. The effect of housing on portfolio choice. *The Journal of Finance*, 72(3):1171–1212, 2017.

F. Chopra, C. Roth, and J. Wohlfart. Home price expectations and spending: Evidence from a field experiment. *American Economic Review*, 115(7):2267–2305, 2025.

C. Cinelli and C. Hazlett. Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(1):39–67, 2020.

S. Corradin and A. Popov. House prices, home equity borrowing, and entrepreneurship. *The Review of Financial Studies*, 28(8):2399–2428, 08 2015.

T. E. Dahl. *Status and trends of wetlands in the conterminous United States 2004 to 2009*. U.S. Department of the Interior; Fish and Wildlife Service, Washington, D.C., 2011.

T. Davidoff. Supply elasticity and the housing cycle of the 2000s. *Real Estate Economics*, 41(4): 793–813, 2013.

T. Davidoff. Supply constraints are not valid instrumental variables for home prices because they are correlated with many demand factors. *Critical Finance Review*, 5(2):177–206, 2016.

R. Diamond. The determinants and welfare implications of us workers' diverging location choices by skill: 1980–2000. *American Economic Review*, 106(3):479–524, 2016.

J. Favilukis, S. C. Ludvigson, and S. Van Nieuwerburgh. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. *Journal of Political Economy*, 125(1):140–223, 2017.

Federal Housing Finance Agency. FHFA House Price Index (HPI). URL https://www.fhfa.gov/data/hpi.

J. Fonseca, L. Liu, and P. Mabille. Unlocking mortgage lock-in: Evidence from a spatial housing ladder model. *Working Paper*, 2025.

Freddie Mac. Freddie Mac House Price Index (FMHPI). URL https://www.freddiemac.com/research/indices/house-price-index.

P. Ganong and D. Shoag. Why has regional income convergence in the us declined? *Journal of Urban Economics*, 102:76–90, 2017.

Z. Gao, M. Sockin, and W. Xiong. Learning about the neighborhood. *The Review of Financial Studies*, 34(9):4323–4372, 11 2020. ISSN 0893-9454. doi: 10.1093/rfs/hhaa130. URL https://doi.org/10.1093/rfs/hhaa130.

E. L. Glaeser and J. Gyourko. America's housing supply problem: The closing of the suburban frontier? Technical report, National Bureau of Economic Research, 2025.

E. L. Glaeser, J. Gyourko, and R. E. Saks. Urban growth and housing supply. *Journal of economic geography*, 6(1):71–89, 2006.

E. L. Glaeser, J. Gyourko, and A. Saiz. Housing supply and housing bubbles. *Journal of Urban Economics*, 64(2):198–217, 2008.

E. L. Glaeser, J. Gyourko, E. Morales, and C. G. Nathanson. Housing dynamics: An urban approach. *Journal of Urban Economics*, 81:45–56, 2014.

Goldman Sachs. The outlook for US housing supply and affordability, Oct. 2025. URL https://www.goldmansachs.com/insights/articles/the-outlook-for-us-housing-supply-and-affordability. Macroeconomics Insights.

Google Maps. Google Maps Platform. URL https://developers.google.com/maps.

J. M. Griffin, S. Kruger, and G. Maturana. What drove the 2003–2006 house price boom and subsequent collapse? disentangling competing explanations. *Journal of Financial Economics*, 141(3):1007–1035, 2021.

A. Gupta, V. Mittal, J. Peeters, and S. Van Nieuwerburgh. Flattening the curve: pandemic-induced revaluation of urban real estate. *Journal of Financial Economics*, 146(2):594–636, 2022.

A. Gupta, H. Sapriza, and V. Yankov. The collateral channel and bank credit. *Working Paper*, 2025.

A. M. Guren, A. McKay, E. Nakamura, and J. Steinsson. Housing wealth effects: The long view. *The Review of Economic Studies*, 88(2):669–707, 2021.

J. Gyourko, C. Mayer, and T. Sinai. Superstar cities. *American Economic Journal: Economic Policy*, 5(4):167–199, 2013.

J. C. Haltiwanger. Entrepreneurship during the covid-19 pandemic: Evidence from the business formation statistics. *Entrepreneurship and Innovation Policy and the Economy*, 1(1):9–42, 2022.

L. Han. Understanding the puzzling risk-return relationship for housing. *The Review of Financial Studies*, 26(4):877–928, 2013.

G. Howard and J. Liebersohn. Why is the rent so darn high? the role of growing demand to live in housing-supply-inelastic cities. *Journal of Urban Economics*, 124:103369, 2021.

T. Hudson. 'The risk levels are higher': Miami is the number one housing market bubble in the world – again. *The Miami Times*, Oct. 2025.

T. L. Jensen, S. Leth-Petersen, and R. Nanda. Financing constraints, home equity and selection into entrepreneurship. *Journal of Financial Economics*, 145(2):318–337, 2022.

G. Kaplan, K. Mitman, and G. L. Violante. Non-durable consumption and housing net worth in the great recession: Evidence from easily accessible data. *Journal of Public Economics*, 189: 104176, 2020a.

G. Kaplan, K. Mitman, and G. L. Violante. The housing boom and bust: Model meets evidence. *Journal of Political Economy*, 128(9):3285–3345, 2020b.

M. P. Keane and T. Neal. A practical guide to weak instruments. *Annual Review of Economics*, 16, 2024.

S. Louie, J. A. Mondragon, and J. Wieland. Supply constraints do not explain house price and quantity growth across us cities. *Working Paper*, 2026.

S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.

S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee. From local explanations to global understanding with explainable ai for trees. *Nature machine intelligence*, 2(1):56–67, 2020.

C. Lutz. Bridging the geographic divide: Crosswalks across space and time. *Working Paper*, 2025a.

C. Lutz. Natural amenity demand index construction using google places. https://github.com/ChandlerLutz/gmaps-natural-amenities-2025, 2025b.

A. Mian and A. Sufi. The consequences of mortgage credit expansion: Evidence from the us mortgage default crisis. *The Quarterly journal of economics*, 124(4):1449–1496, 2009.

A. Mian and A. Sufi. House prices, home equity–based borrowing, and the us household leverage crisis. *American Economic Review*, 101(5):2132–2156, 2011.

A. Mian and A. Sufi. What explains the 2007–2009 drop in employment? *Econometrica*, 82(6): 2197–2223, 2014.

A. Mian, K. Rao, and A. Sufi. Household balance sheets, consumption, and the economic slump. *The Quarterly Journal of Economics*, 128(4):1687–1726, 2013.

R. Molloy. The effect of housing supply regulation on housing affordability: A review. *Regional science and urban economics*, 80(C):1–5, 2020.

A. Murphy. A dynamic model of housing supply. *American Economic Journal: Economic Policy*, 10(4):243–267, 2018.

K. M. Murphy and R. H. Topel. Estimation and inference in two-step econometric models. *Journal of Business & Economic Statistics*, 3(4):370–379, 1985. ISSN 07350015.

C. G. Nathanson and E. Zwick. Arrested development: Theory and evidence of supply-side speculation in the housing market. *The Journal of Finance*, 73(6):2587–2633, 2018.

Office of the Governor of California. Governor newsom signs into law groundbreaking reforms to build more housing, boost affordability, June 2025. URL https://www.gov.ca.gov/2025/06/30/governor-newsom-signs-into-law-groundbreaking-reforms-to-build-more-housing-affordabilit Press Release.

Office of the Governor of New York. The new york housing compact, Mar. 2023. URL https://www.governor.ny.gov/sites/default/files/2023-03/NY_Housing_Compact_info_deck_March_2023.pdf. Information Deck.

H. Oh and C. Yoon. Time to build and the real-options channel of residential investment. *Journal of Financial Economics*, 135(1):255–269, 2020.

A. Paciorek. Supply constraints and housing market dynamics. *Journal of Urban Economics*, 77: 11–26, 2013.

A. Pagan. Econometric issues in the analysis of regressions with generated regressors. *International Economic Review*, pages 221–247, 1984.

J. R. Pierce and P. K. Schott. The surprisingly swift decline of us manufacturing employment. *American Economic Review*, 106(7):1632–1662, 2016.

A. Saiz. The geographic determinants of housing supply. *The Quarterly Journal of Economics*, 125 (3):1253–1296, 2010.

A. Saiz. The global housing affordability crisis: Policy options and strategies. *MIT Center for Real Estate Research Paper*, (23/01), 2023.

J. Schroeder, D. Van Riper, S. Manson, K. Knowles, T. Kugler, F. Roberts, and S. Ruggles. IPUMS National Historical Geographic Information System: Version 20.0 [dataset], 2025. URL http://doi.org/10.18128/D050.V20.0.

L. S. Shapley. A value for n-person games. 1953.

S. A. Sherman, C. S. Cheong, D. Banerjee, A. Kim, and A. Yang. The 2025 State of Housing in Harris County and Houston. Report, Kinder Institute for Urban Research, Rice University, 2025. URL https://kinder.rice.edu/research/2025-state-housing-harris-county-and-houston.

T. Sinai and N. S. Souleles. Owner-occupied housing as a hedge against rent risk. *The Quarterly Journal of Economics*, 120(2):763–789, 2005.

D. Staiger and J. H. Stock. Instrumental variables regression with weak instruments. *Econometrica*, 65(3):557–586, 1997.

J. Stroebel and J. Vavra. House prices, local demand, and retail prices. *Journal of Political Economy*, 127(3):1391–1436, 2019.

The New York Times. The housing shortage isn't just a coastal crisis anymore. *The New York Times*, July 2022. URL https://www.nytimes.com/2022/07/14/upshot/housing-shortage-us.html. The Upshot.

The White House. Fact sheet: President donald j. trump delivers emergency price relief for american families to defeat the cost-of-living crisis, Jan. 2025. URL https://www.whitehouse.gov/fact-sheets/2025/01/fact-sheet-president-donald-j-trump-delivers-emergency-price-relief-for-american-familie Fact Sheet.

S. Titman. Urban land prices under uncertainty. *The American Economic Review*, 75(3):505–514, 1985.

United States Senate Committee on Banking, Housing, and Urban Affairs. ROAD to housing act of 2025: Section-by-section, July 2025. URL https://www.banking.senate.gov/imo/media/doc/road_to_housing_act_of_2025_section_by_section.pdf. Committee Document.

U.S. Census Bureau. TIGER/Line Shapefiles, a. URL https://www.census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html.

U.S. Census Bureau. County Business Patterns (CBP) Data, b. URL https://www.census.gov/programs-surveys/cbp/data.html.

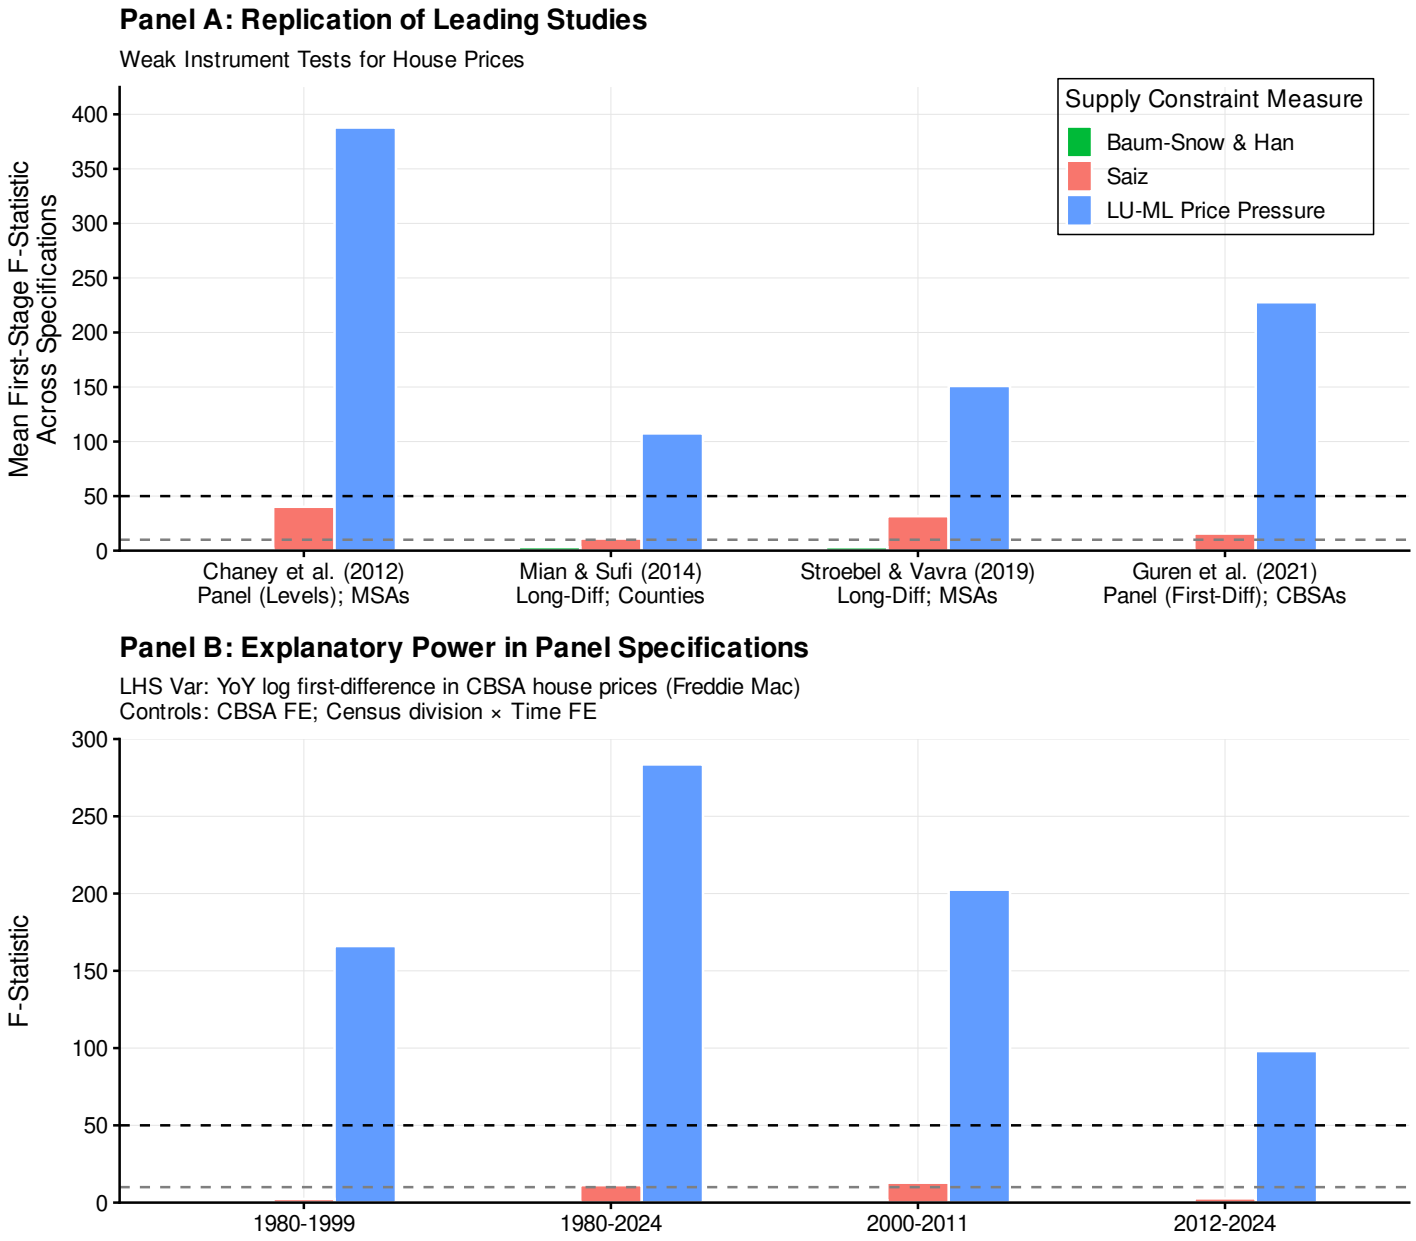
U.S. Geological Survey. The National Map - Data Delivery. URL https://www.usgs.gov/the-national-map-data-delivery.

XGBoost Documentation. XGBoost Documentation: Notes on Parameter Tuning, 2025. URL https://xgboost.readthedocs.io/en/stable/tutorials/param_tuning.html.

Zillow. Zillow Housing Data. URL https://www.zillow.com/research/data/.

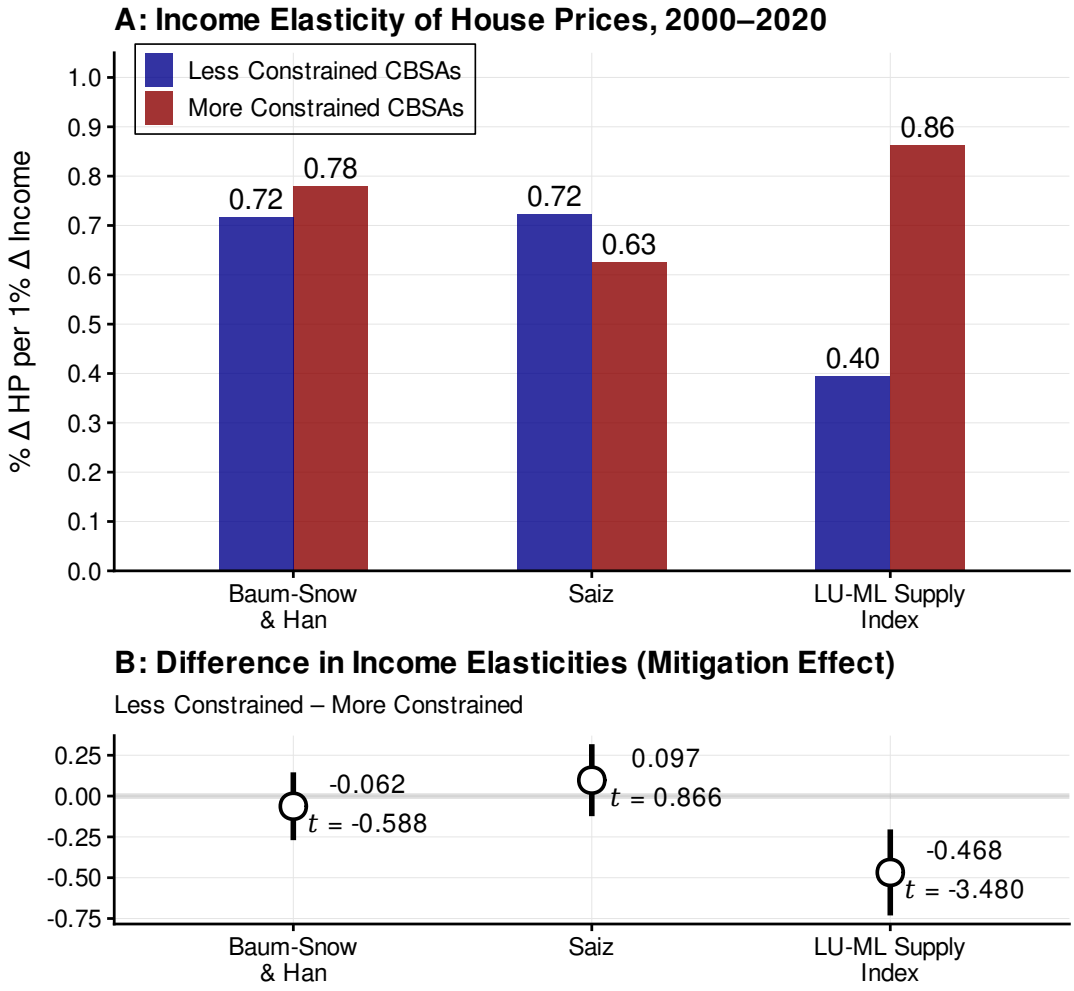
7 Tables and Figures

Figure 1: The Explanatory Power of Supply Constraint Measures for House Prices



Notes: This figure assesses the explanatory power of supply constraint measures for house prices using weak instrument tests (first-stage F-statistics). We compare the Saiz (2010) and Baum-Snow and Han (BSH, 2024) elasticity proxies against the LU-ML Price Pressure Index. The LU-ML Price Pressure Index is constructed using house prices as the target variable (in levels or differences, matching the respective specification). The gray-dashed line marks the traditional weak instrument threshold of 10 (Staiger and Stock, 1997); the black-dashed line marks a modern threshold of 50 suggested by Keane and Neal (2024). Panel A replicates leading studies. Chaney et al. (2012) use a panel (1993–2007) of MSA-level log house prices, instrumenting with the interaction of the Saiz elasticity and national mortgage rates. Mian and Sufi (2014) use a long-difference specification (2006–2009) for counties. Stroebel and Vavra (2019) use long differences (2001–2007, 2007–2011) for MSAs. Guren et al. (2021) use a quarterly panel (1978–2017) of year-over-year (YoY) log price changes for CBSAs, interacting the Saiz proxy with the change in national house prices. See Appendix D for full replication output. Panel B reports results for panel models using Freddie Mac house price indices. The panel specification is defined in Equation (1), where the dependent variable is the YoY log first-difference in house prices. We interact the BSH and Saiz elasticity proxies with the YoY log first-difference in national house prices, following Guren et al. (2021), to generate time-varying proxies. Standard errors are clustered by CBSA and time.

Figure 2: Income Elasticities of House Prices by Supply Measure



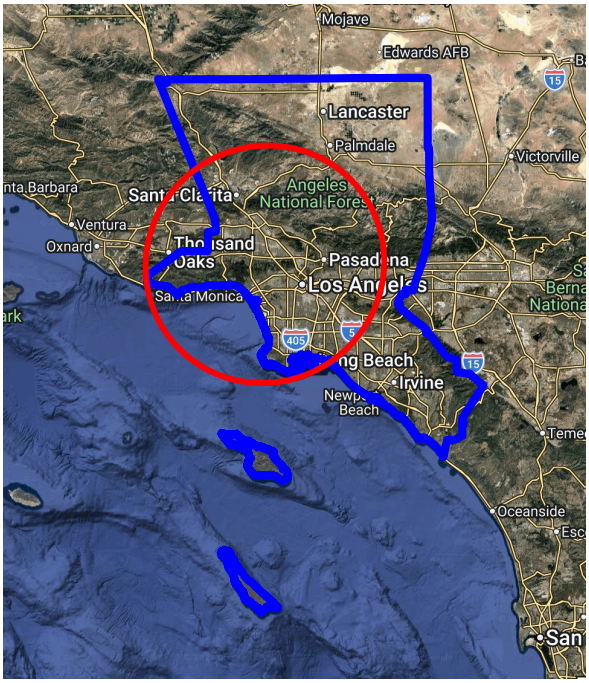
Notes: This figure tests whether looser supply constraints mitigate house price appreciation in response to income growth across major metropolitan CBSAs using a long-differenced specification from 2000–2020. We compare results using the Baum-Snow and Han (2024) and Saiz (2010) proxies against the LU-ML Supply Index. The LU-ML Supply Index is constructed using the log change in CBSA-level housing units from 2000–2020 as the target variable. Estimates are from the following specification:

$$\Delta \log(HP_i) = \alpha + \beta_1 \Delta \log(Income_i) + \beta_2 LessConstrained_i + \beta_3 [\Delta \log(Income_i) \times LessConstrained_i] + \epsilon_i$$

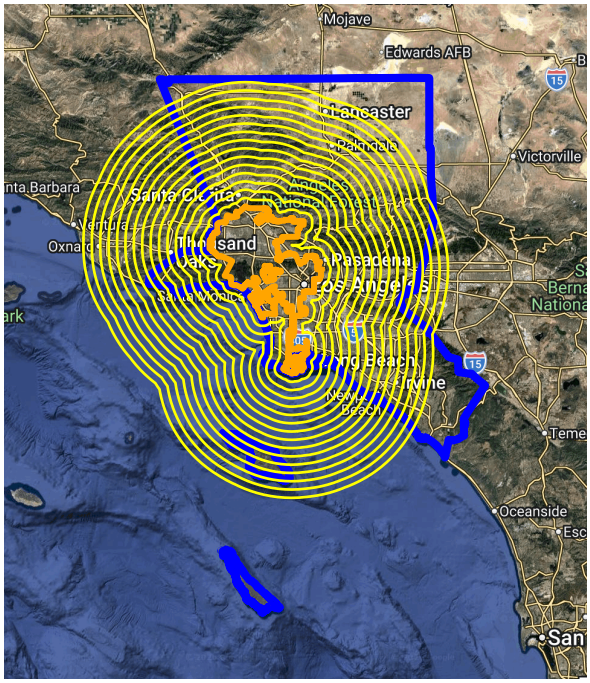
$LessConstrained_i$ is an indicator equal to 1 for CBSAs with looser-than-average supply constraints according to the respective proxy. Panel A plots the income elasticity of house prices for *More* versus *Less Constrained* CBSAs (β_1 versus $\beta_1 + \beta_3$). Panel B plots the difference in the income elasticity of house prices (β_3), representing the mitigation effect of looser constraints. HP_i is the Freddie Mac house price index. *Income* is total BEA income. All data are aggregated to 2020 CBSA delineations. See Appendix Tables B1 and B2, and Appendix Figure B1 for robustness to alternative house price datasets, spatial coverage, and county delineations.

Figure 3: Land Unavailability Circles and Buffers for Los Angeles

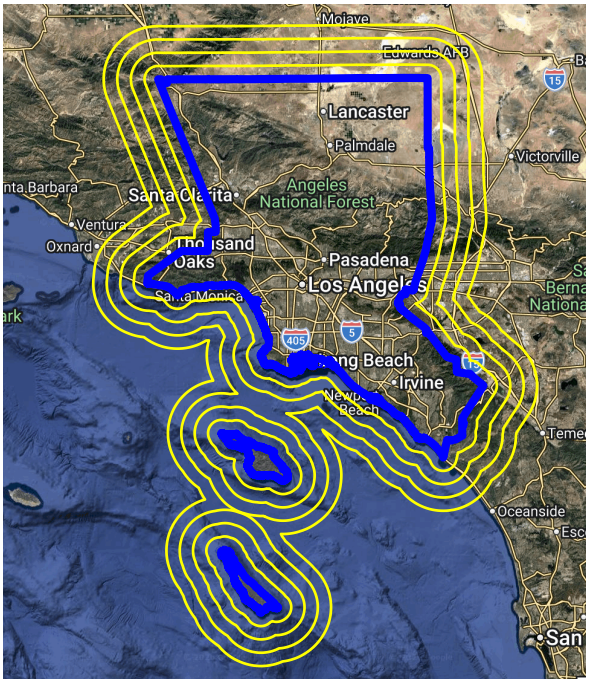
A: LA-Long Beach-Anaheim CBSA with Saiz 50km Circle



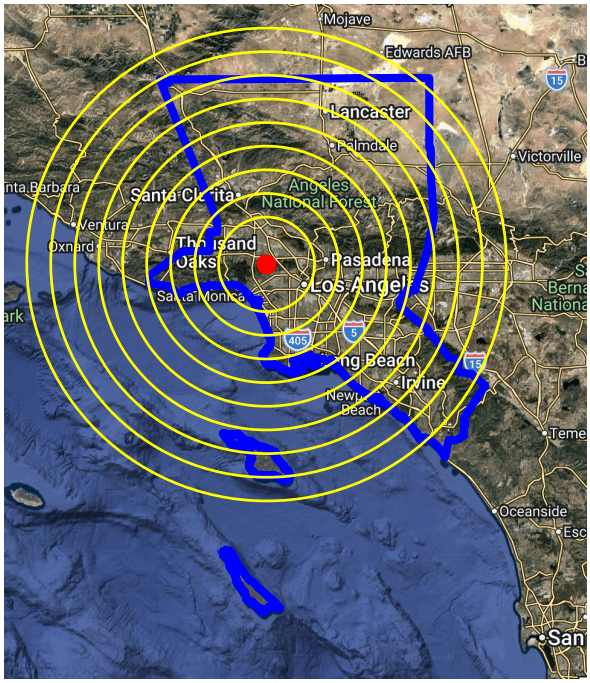
B: LU-ML First Principal City Buffers



C: LU-ML CBSA Polygon Buffers

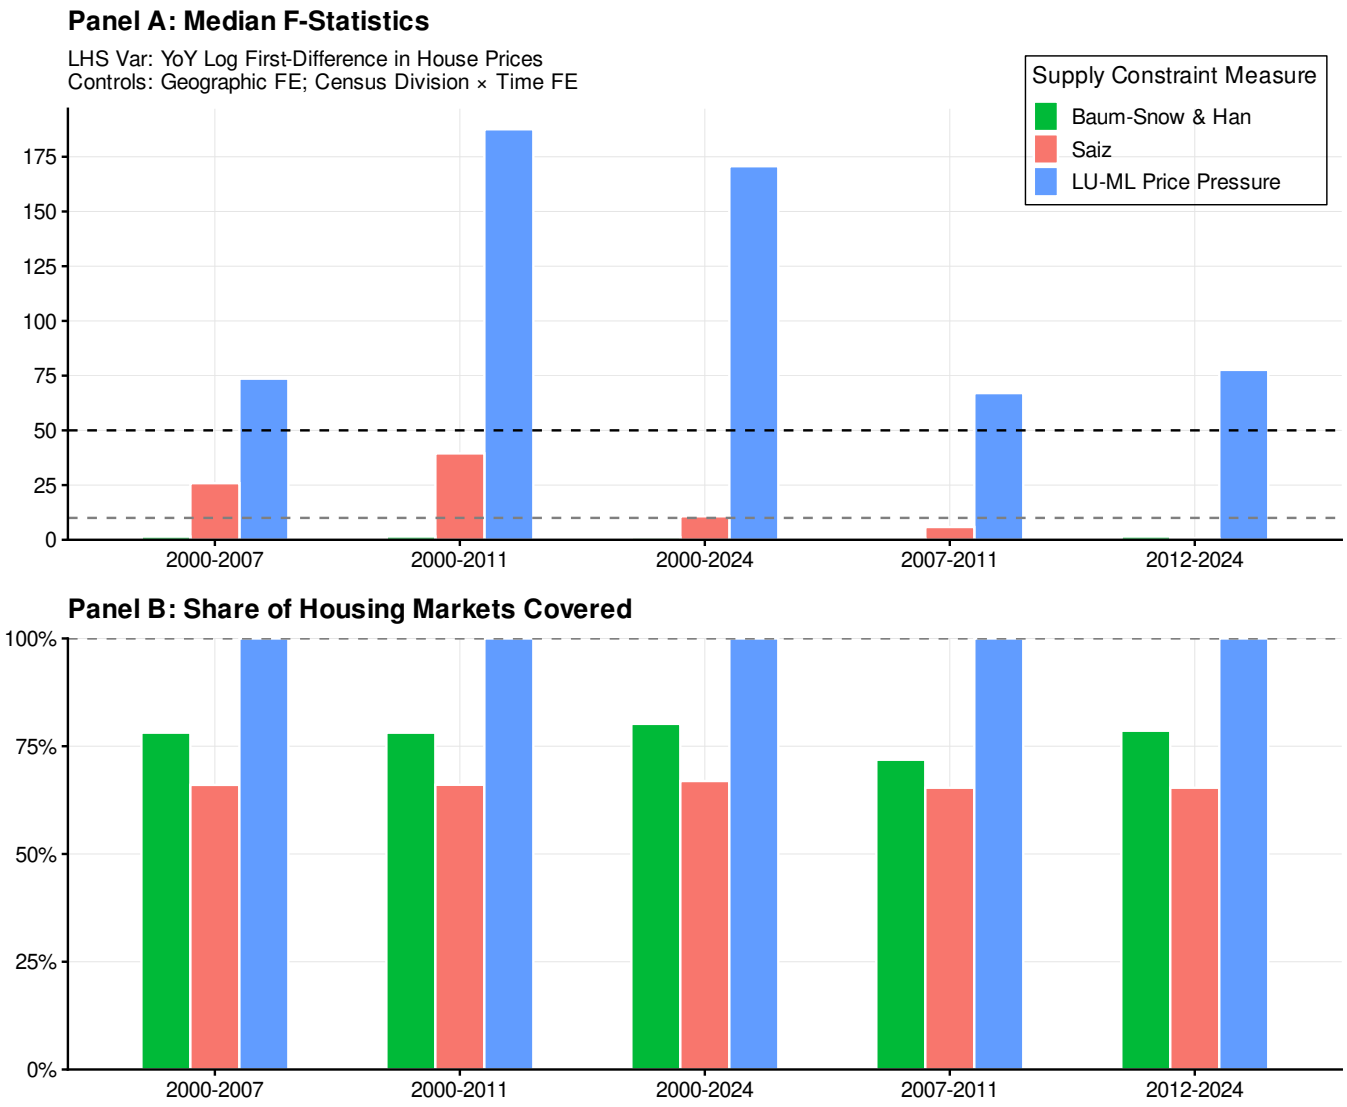


D: LU-ML First Principal City Centroid Circles



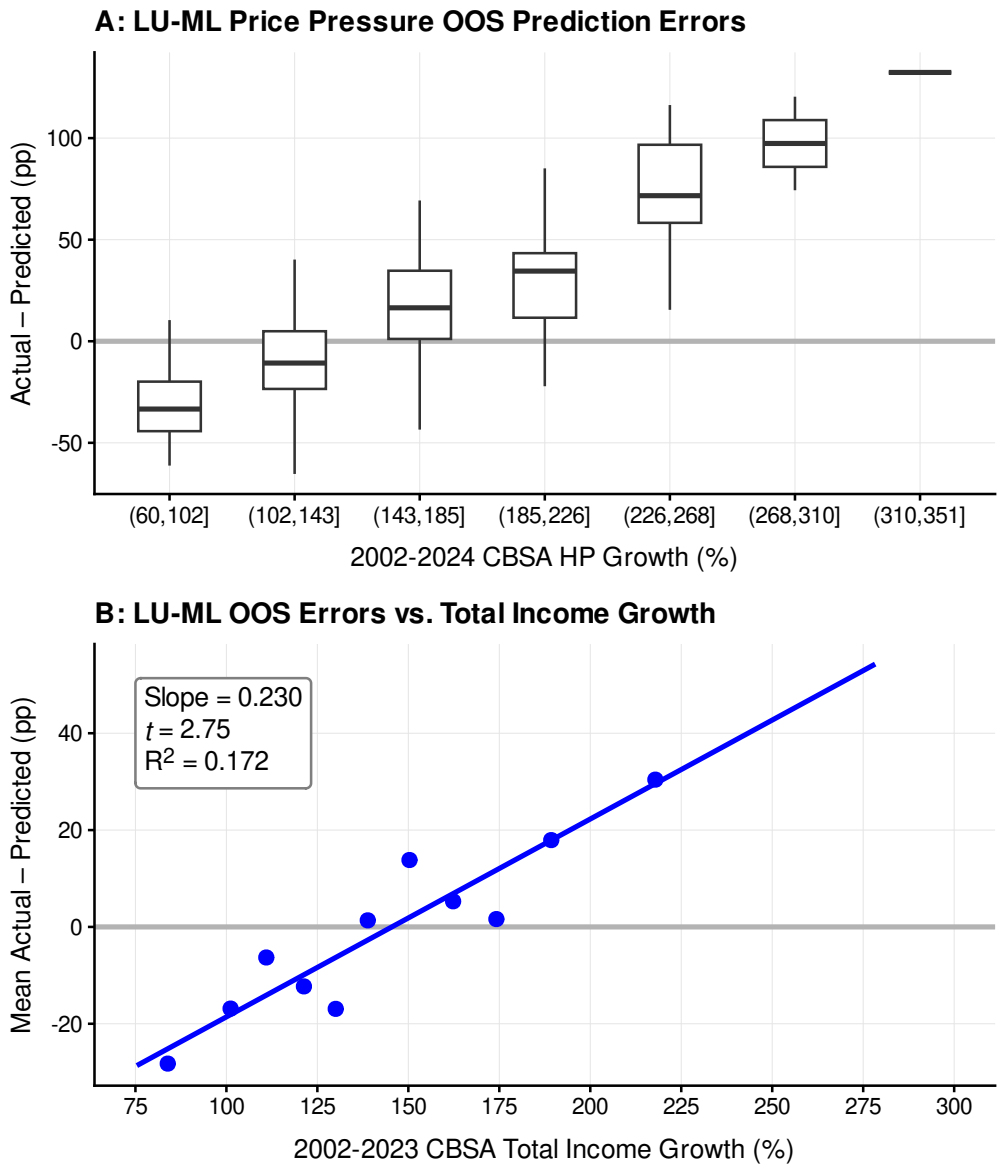
Notes: This figure illustrates the spatial definitions used to calculate Land Unavailability. In each panel, the blue boundary represents the Los Angeles-Long Beach-Anaheim CBSA (2023 delineations). Panel A displays the standard definition from Saiz (2010) and Baum-Snow and Han (2024): a 50km radius circle (red) centered on the centroid of the first principal city (Los Angeles). Panels B–D illustrate the varying spatial features used in the LU-ML indices. Panel B displays the first principal city polygon (orange) and its surrounding buffers (yellow) ranging from 10 to 150 percent (in 10pp increments). Panel C displays the CBSA polygon (blue) and its surrounding buffers (yellow) ranging from 5 to 20 percent (in 5pp increments). Panel D displays concentric circles with radii ranging from 20km to 100km (in 10km increments) centered on the first principal city centroid (red dot).

Figure 4: Explanatory Power of Supply Constraint Measures in House Price Panels



Notes: This figure summarizes the explanatory power of supply constraint measures across house price panels. We compare the Saiz (2010) and Baum-Snow and Han (2024) elasticity proxies against the LU-ML Price Pressure Index. The summary includes panel datasets from Freddie Mac and the FHFA, aggregated to various geographic levels including CBSAs, counties, zip codes, and census tracts. Estimates are derived from the following specification: $\Delta \log HP_{i,d,t} = \psi_i + \gamma_{d,t} + \beta \cdot SC_{i,d,t} + \varepsilon_{i,d,t}$. $\Delta \log HP_{i,d,t}$ is the year-over-year (YoY) log first-difference in house prices for geographic unit i (e.g., a zip code or CBSA) in Census Division d at month t. ψ_i denotes geographic unit fixed effects, and $\gamma_{d,t}$ denotes Census division-by-time fixed effects. $SC_{i,d,t}$ represents the supply constraint measure. To generate time-varying supply constraints for the static Saiz and BSH proxies, we interact them with the YoY log change in national house prices (Guren et al., 2021). The LU-ML Price Pressure Index is constructed for each specific dataset using the corresponding YoY price change as the target variable. Panel A reports weak instrument tests: the median first-stage F-statistic across all specifications within each estimation period. The gray-dashed line marks the weak instrument threshold of 10 (Staiger and Stock, 1997); the black-dashed line marks the threshold of 50 (Keane and Neal, 2024). Panel B reports the coverage rate measured as the share of available cross-sectional geographic units for which a supply constraint measure exists, averaged across specifications in that period. See Appendix Table A1 for the partial R^2 and number of cross-sectional units for each individual specification.

Figure 5: LU-ML Prediction Errors and Housing Demand Growth



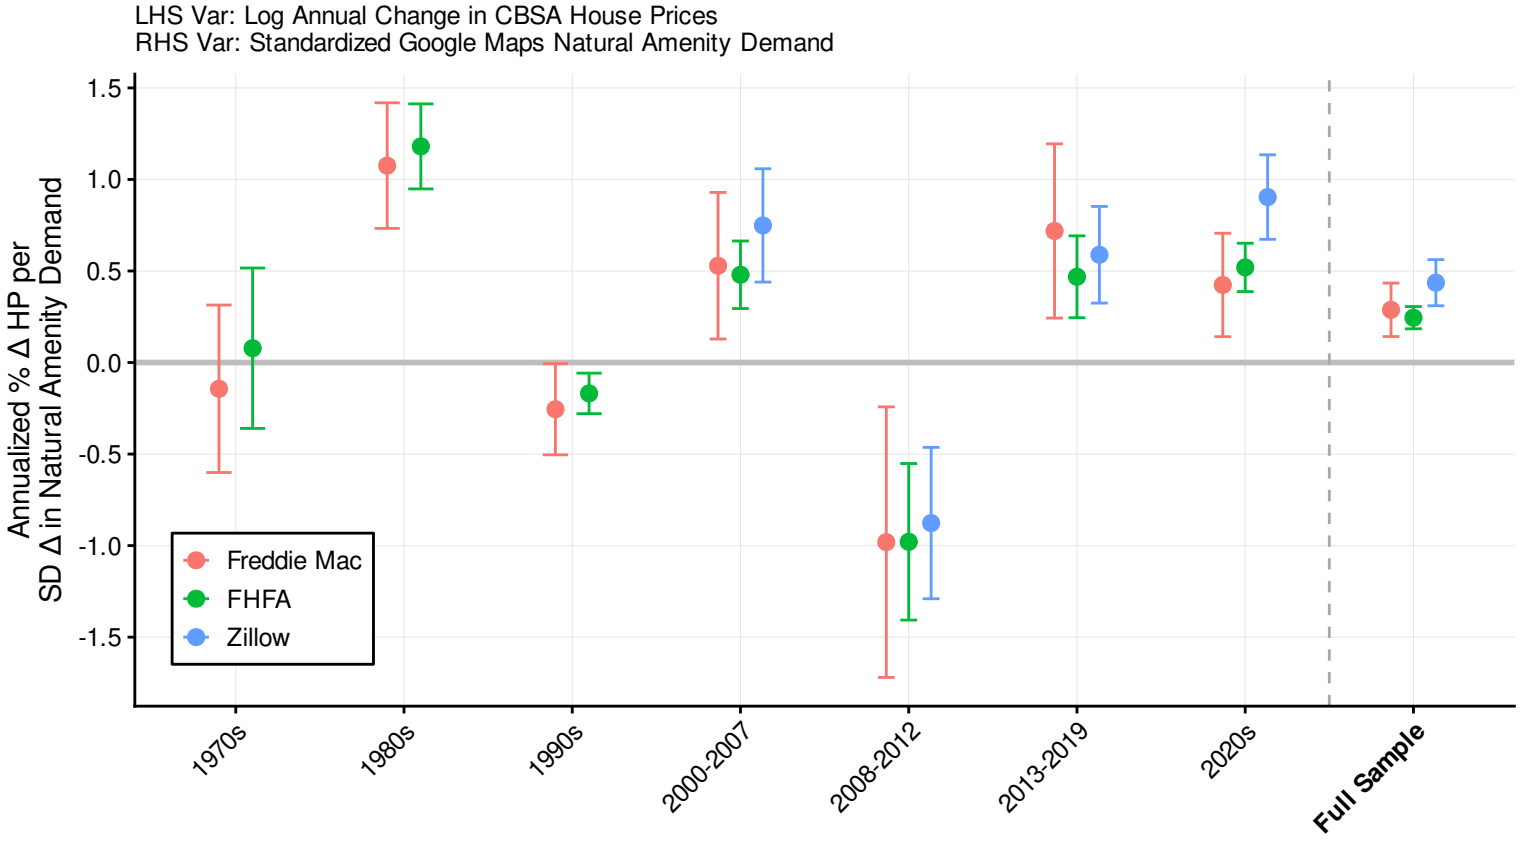
Notes: This figure analyzes the out-of-sample (OOS) prediction errors of the LU-ML Price Pressure Index. Panel A displays boxplots of the OOS prediction errors (vertical axis) sorted by cumulative Freddie Mac CBSA house price growth (horizontal axis) from 2002–2024. We construct the LU-ML Price Pressure Index using 2002–2024 Freddie Mac cumulative house price growth as the target variable. Panel B presents a binscatter of the OOS prediction errors against 2002–2023 BEA total income growth. We select the number of bins following Cattaneo et al. (2024). The regression line and statistics are estimated on the underlying data and use robust standard errors. See Appendix Figures A1 and A2 for the same figures using the FHFA and Zillow house price data.

Table 1: LU-ML Price Pressure and House Prices (2000–2023)

	Dependent variable: YoY Log Change in House Prices								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
LU-ML Price Pressure Index	0.695***	0.661***	0.642***	0.634***	0.612***	0.593***	0.583***	0.508***	0.201***
	(0.048)	(0.045)	(0.045)	(0.044)	(0.045)	(0.045)	(0.044)	(0.043)	(0.031)
Total Income Growth		0.328***	0.315***	0.314***	0.320***	0.314***	0.316***	0.310***	0.229***
		(0.056)	(0.054)	(0.054)	(0.056)	(0.055)	(0.055)	(0.050)	(0.030)
Natural Amenity Quintile 2 × YoY National HP Growth			0.004						
			(0.065)						
Natural Amenity Quintile 3 × YoY National HP Growth			0.058						
			(0.065)						
Natural Amenity Quintile 4 × YoY National HP Growth			0.089						
			(0.065)						
Natural Amenity Quintile 5 × YoY National HP Growth			0.177***						
			(0.068)						
CBSA Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Census Division × Date FE	✓	✓	✓	✓	✓	✓	✓	✓	
Natural Amenity Quintile × Date FE				✓	✓	✓	✓	✓	✓
2000 College Share × Date FE					✓	✓	✓	✓	✓
2000 Median Income × Date FE						✓	✓	✓	✓
2000 Housing Density × Date FE							✓	✓	✓
2000 Emp Industry Shares × Date FE								✓	✓
State × Date FE									✓

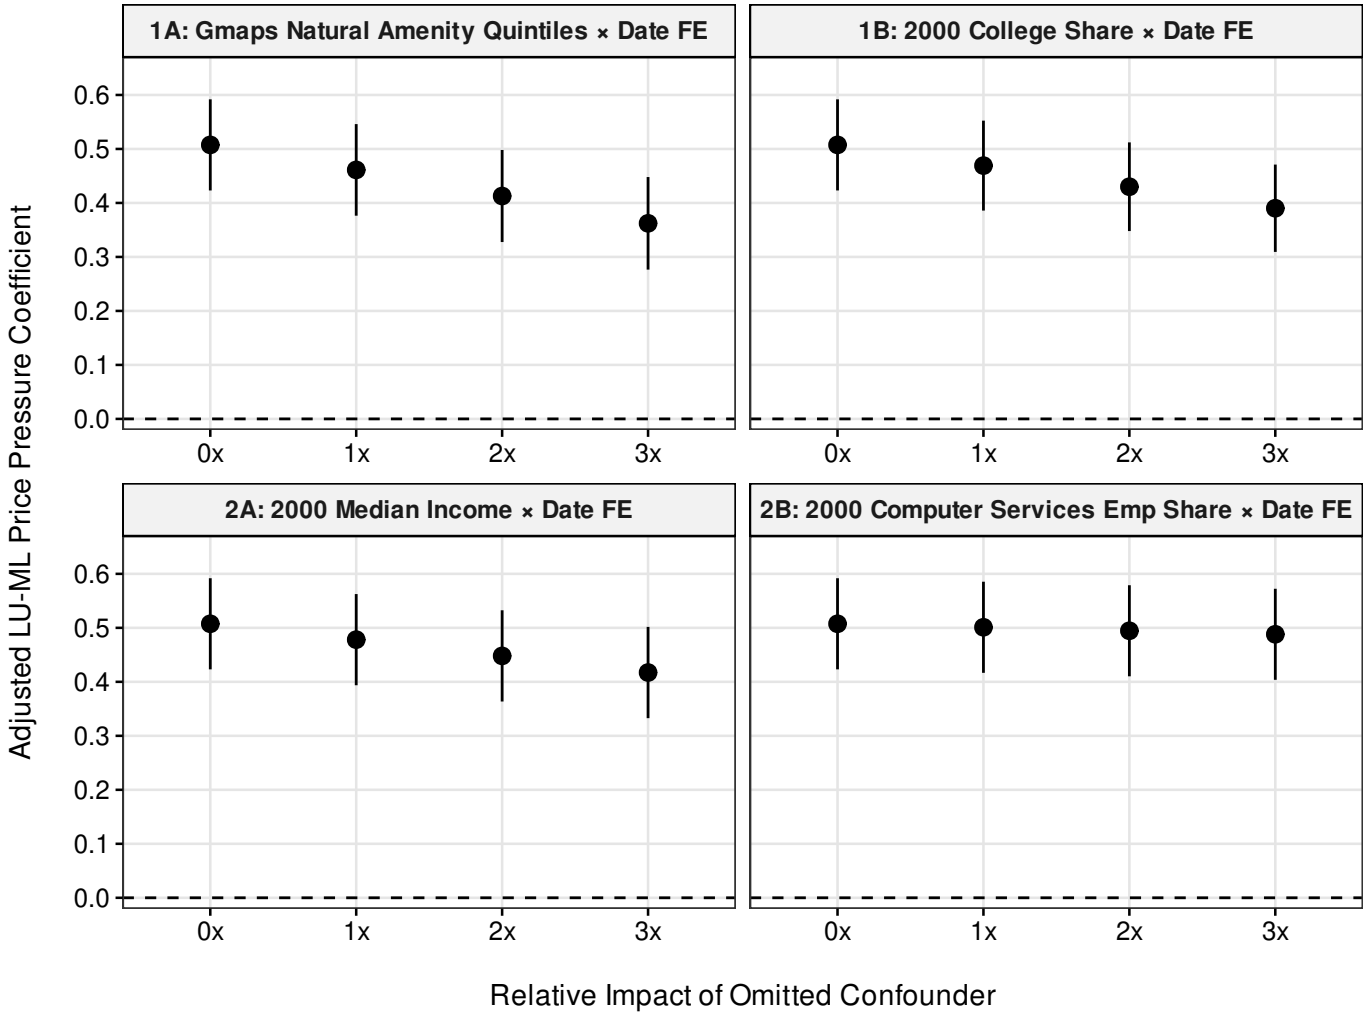
Notes: This table reports panel regressions estimating the effect of the LU-ML Price Pressure Index on year-over-year (YoY) CBSA-level house price growth from 2000–2023, controlling for local factors. Column (1) presents the baseline specification: $\Delta \log HP_{i,d,t} = \psi_i + \gamma_{d,t} + \beta \cdot LU\text{-}ML_{i,t} + \varepsilon_{i,d,t}$, where $\Delta \log HP_{i,d,t}$ is the YoY log change in Freddie Mac house prices for CBSA i in Census division d at time t. ψ_i and $\gamma_{d,t}$ are CBSA and Census division-by-date fixed effects, respectively. We construct the LU-ML Price Pressure Index using YoY house price growth as the target variable. Subsequent columns progressively add local demand and supply controls. Column (2) controls for total CBSA income growth. Column (3) adds Natural Amenity Quintiles interacted with YoY national house price growth to allow the housing amenity premium to vary with the aggregate housing cycle. Column (4) instead interacts Natural Amenity Quintiles with date fixed effects. Columns (5) to (8) sequentially saturate the model with year-2000 CBSA-level characteristics (college share, median income, housing unit density, and local industry employment shares). We interact these year-2000 CBSA-level characteristics with date fixed effects to allow their associated house price premium to vary over time. Employment shares are based on 1998–2002 NAICS codes and span computer services; construction; finance, insurance, and real estate (FIRE); manufacturing; retail; and services less computer services. Finally, Column (9) uses state-by-date fixed effects. Robust standard errors clustered by CBSA and time are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 6: The Natural Amenity House Price Growth Premium



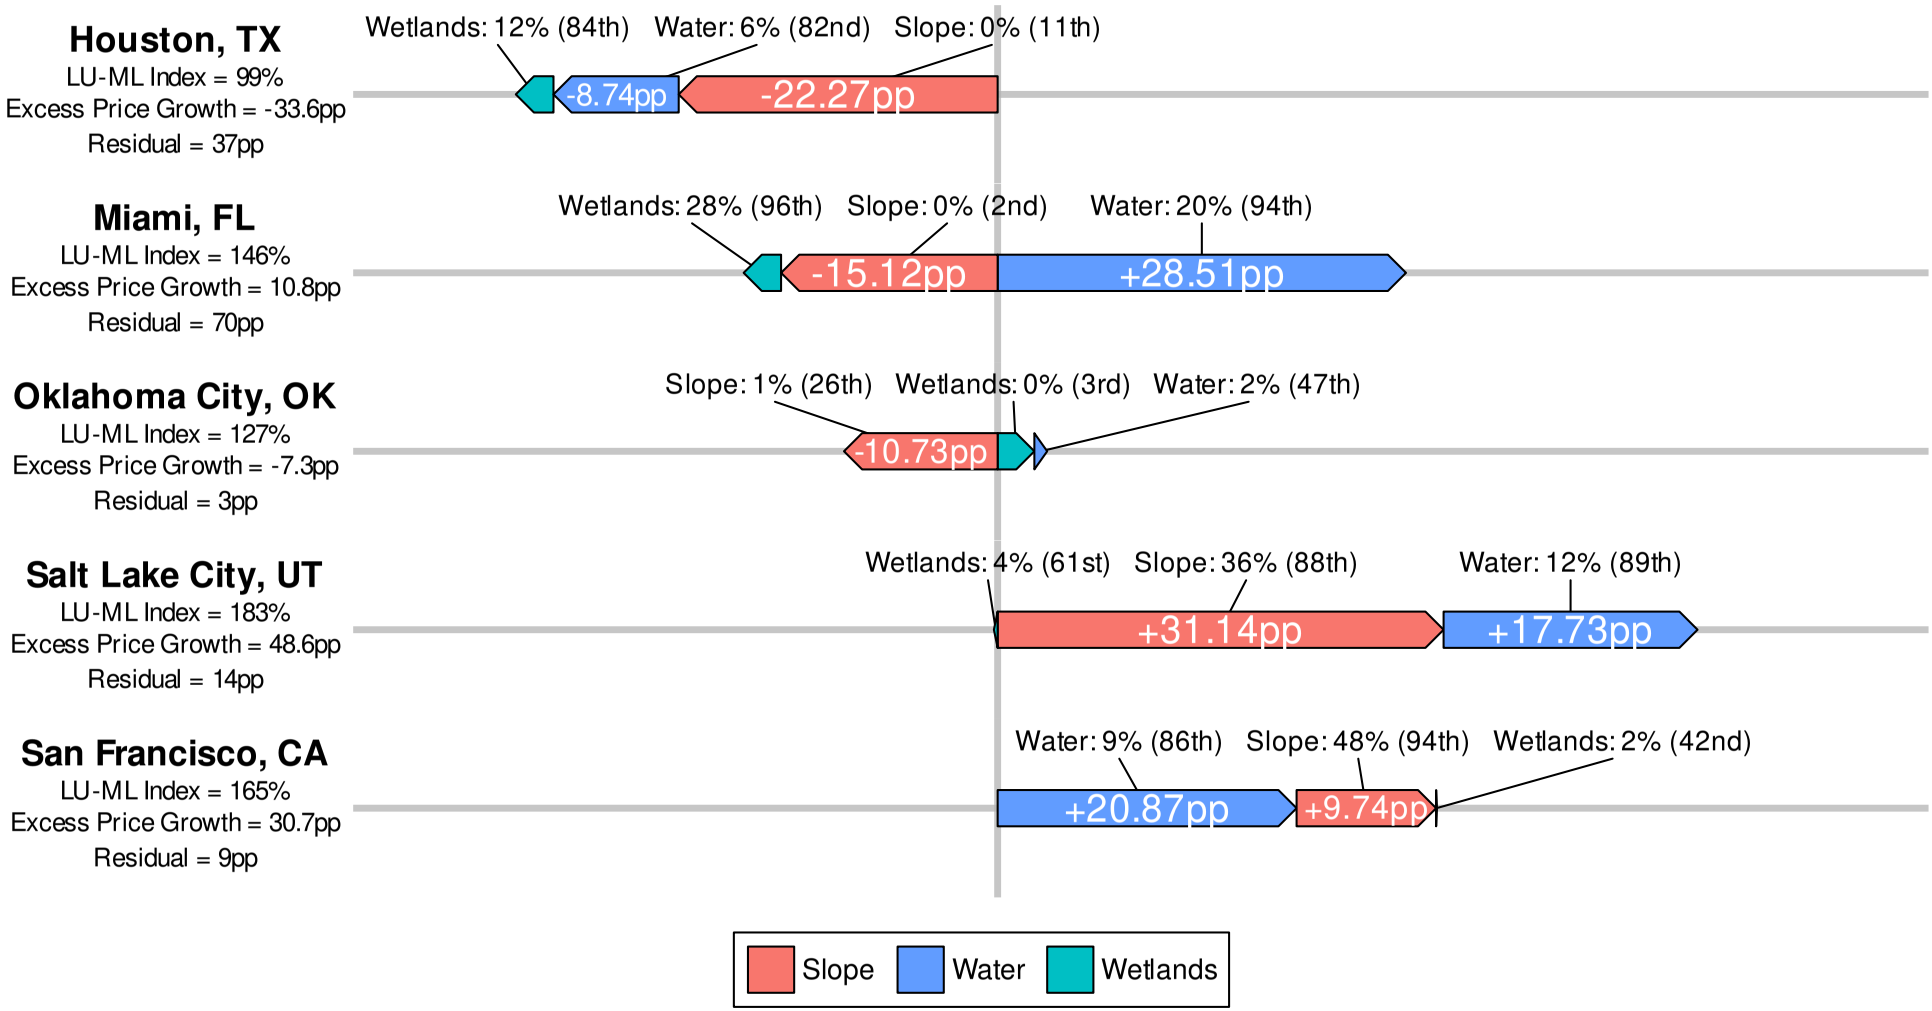
Notes: This figure plots the natural amenity house price growth premium. For each time period on the horizontal axis and house price dataset, we run separate regressions using the following specification: $\Delta \log(HP_i) = \beta_0 + \beta_1 Amenities_i + \varepsilon_i$, where $\Delta \log(HP_i)$ is the house price growth for CBSA i for the specified period and $Amenities_i$ is the Google Maps Natural Amenities Index (Lutz, 2025b) for CBSA i. Confidence intervals correspond to ± 1.96 robust standard errors.

Figure 7: Stability of the LU-ML Price Pressure Index to Unobserved Spatial Demand



Notes: This figure plots the stability of the coefficient of the LU-ML Price Pressure Index to unobserved spatial demand (Cinelli and Hazlett, 2020), evaluated via the following model: $\Delta \log HP_{i,d,t} = \psi_i + \gamma_{d,t} + \beta \cdot LU\text{-}ML_{i,d,t} + \mathbf{X}'_{i,d,t}\lambda_t + \varepsilon_{i,d,t}$, where $\Delta \log HP_{i,d,t}$ is the year-over-year (YoY) log first-difference in house prices for CBSA i in Census Division d at month t. We construct the LU-ML Price Pressure Index using the YoY log first difference in CBSA house prices as the target variable. The vector $\mathbf{X}'_{i,d,t}\lambda_t$ represents CBSA-level controls. Baseline controls are interacted with decade fixed effects. See Table 1 for the full regression results and list of controls. The unadjusted baseline (0x on the horizontal axis) reports the coefficient on the LU-ML Price Pressure Index assuming no unobserved confounding. The bounds simulate an unobserved spatial demand confounder that has 1, 2, or 3 times the explanatory power of the chosen benchmark covariates (listed in the sub-panel titles). Because the CBSA fixed effects absorb the base-period (2000M01), the benchmarks reflects the explanatory power of the time-varying slopes for all remaining periods. To account for spatial clustering by CBSAs, we scale the adjusted IID standard errors by the ratio of the baseline clustered standard error to the baseline IID standard error. This assumes that the intra-cluster correlation does not change with the addition of the omitted variable. Confidence intervals correspond to ± 1.96 adjusted standard errors. See Appendix Table A2 for the full estimation output.

Figure 8: Feature Contributions to the LU-ML Price Pressure Index in Representative Cities (2002–2024)



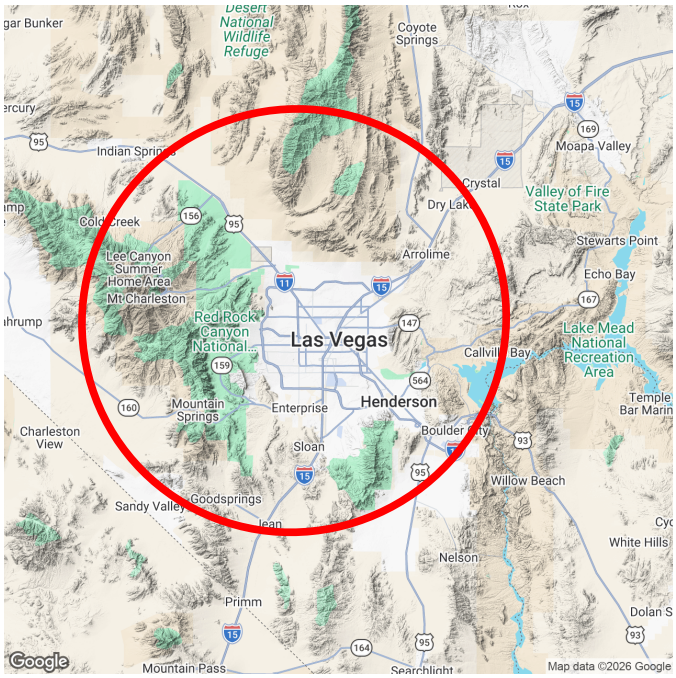
Notes: This figure decomposes the LU-ML Price Pressure Index for representative CBSAs using SHAP values (Lundberg and Lee, 2017; Lundberg et al., 2020). The LU-ML Price Pressure Index is trained using 2002–2024 Zillow cumulative house price growth as the target variable and employs 90 total features (slope, water, and wetlands LU measured using 30 different spatial definitions). The SHAP value, $\varphi_i^{(j)}$, measures the marginal contribution of feature j to the LU-ML Index for CBSA i relative to a coalition of other features ($\neq j$), according to a fair distribution of payouts from cooperative game theory (Shapley, 1953). This allows for an additive decomposition of the index: LU-ML Index$_i = \mu + \sum_{j=1}^{M} \varphi_i^{(j)}$, where μ represents the unconditional mean prediction across all CBSAs. Text below each city label reports the total predicted growth (LU-ML Index), the predicted growth in excess of the national mean ($\sum_j \varphi^{(j)}$), and the prediction error (Residual). The arrows visualize the SHAP contributions: we aggregate the SHAP values from the 90 underlying features into three categories—Slope (red), Water (blue), and Wetlands (teal). The length of each arrow represents the feature set's contribution to the deviation from the mean prediction (μ), measured in percentage points (pp). Labels on the arrows report the underlying land unavailability percentages for that category, with national percentiles in parentheses.

Figure 9: Comparing Unweighted Land Unavailability to LU-ML Feature Contributions for Chicago and Las Vegas

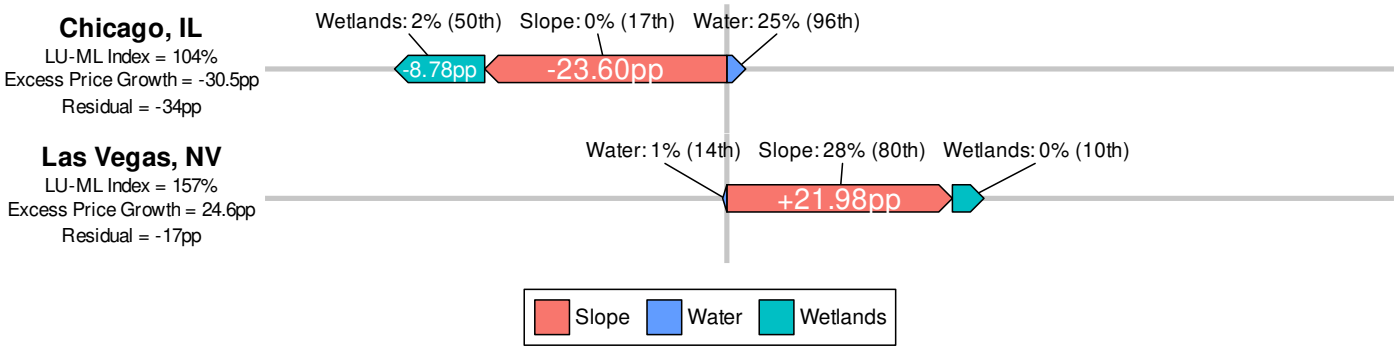
A: Chicago with a 50km Radius Circle



B: Las Vegas with a 50km Radius Circle

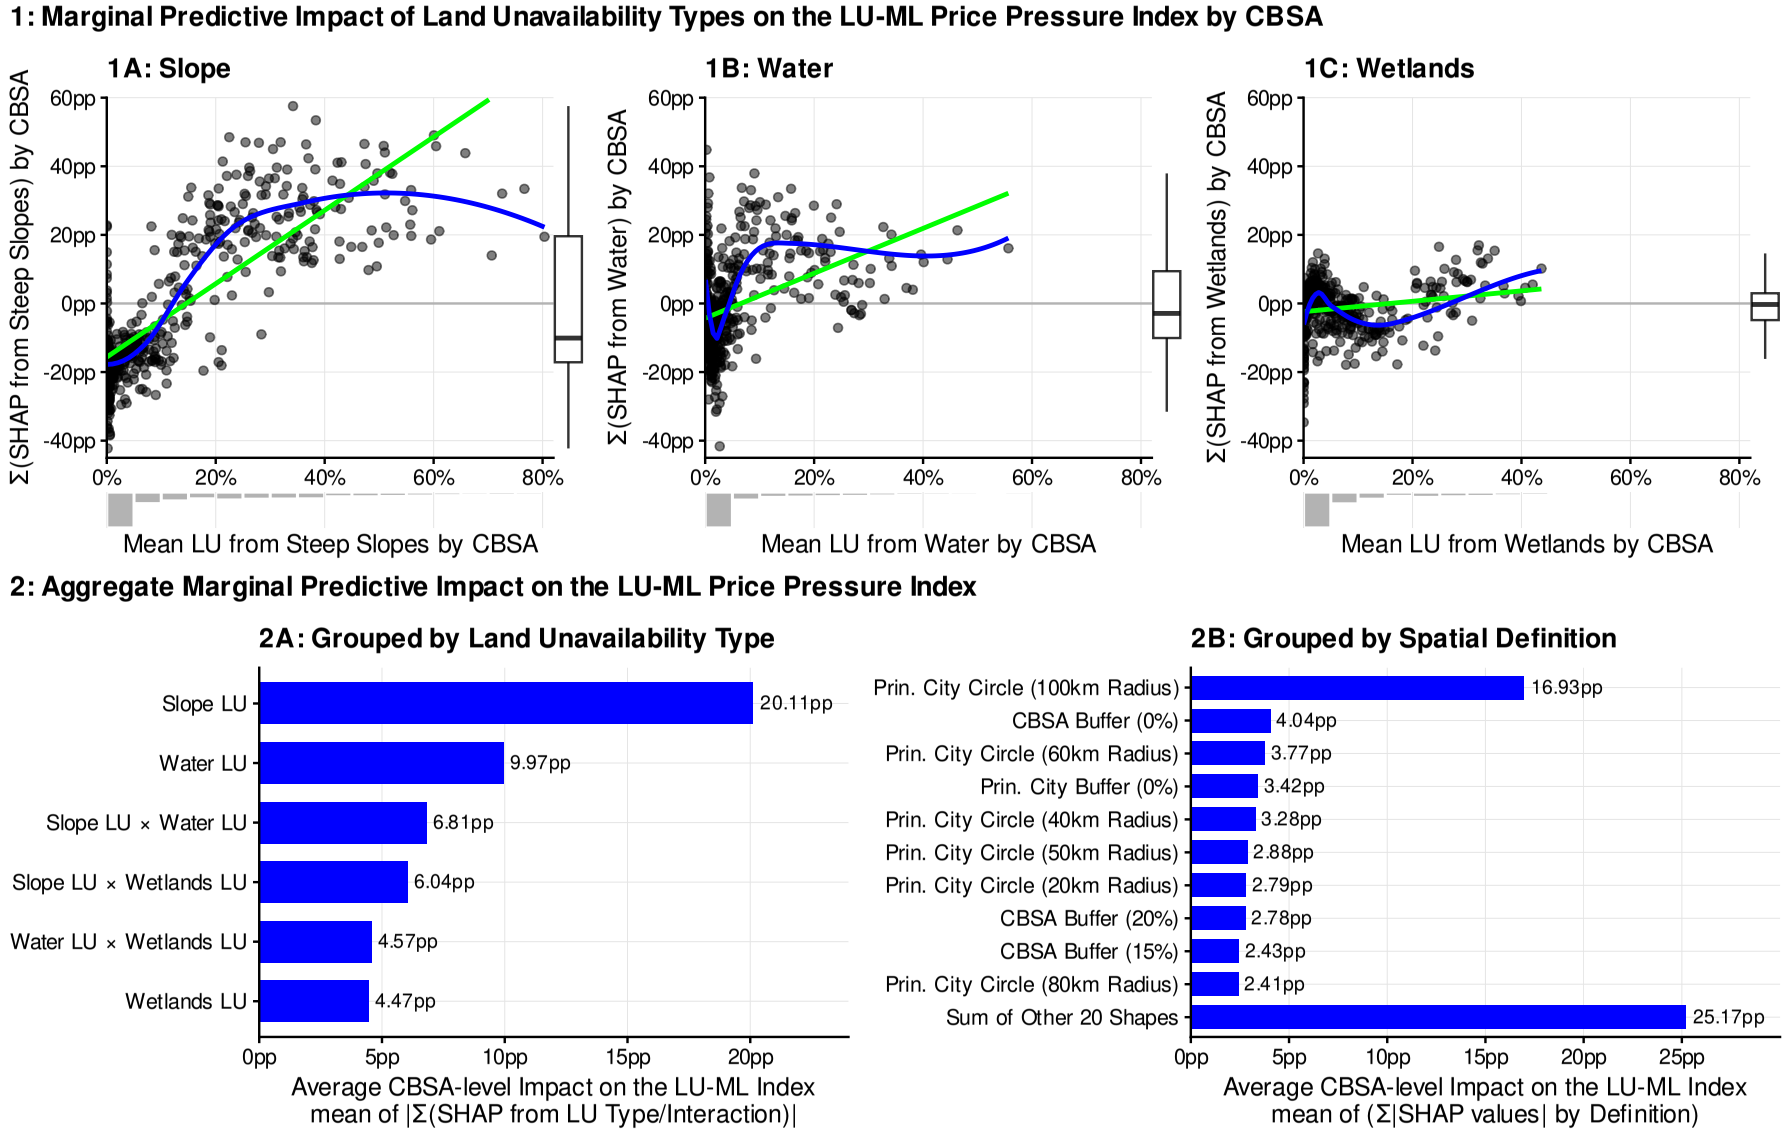


C: Land Unavailability Contributions to the LU-ML Price Pressure Index



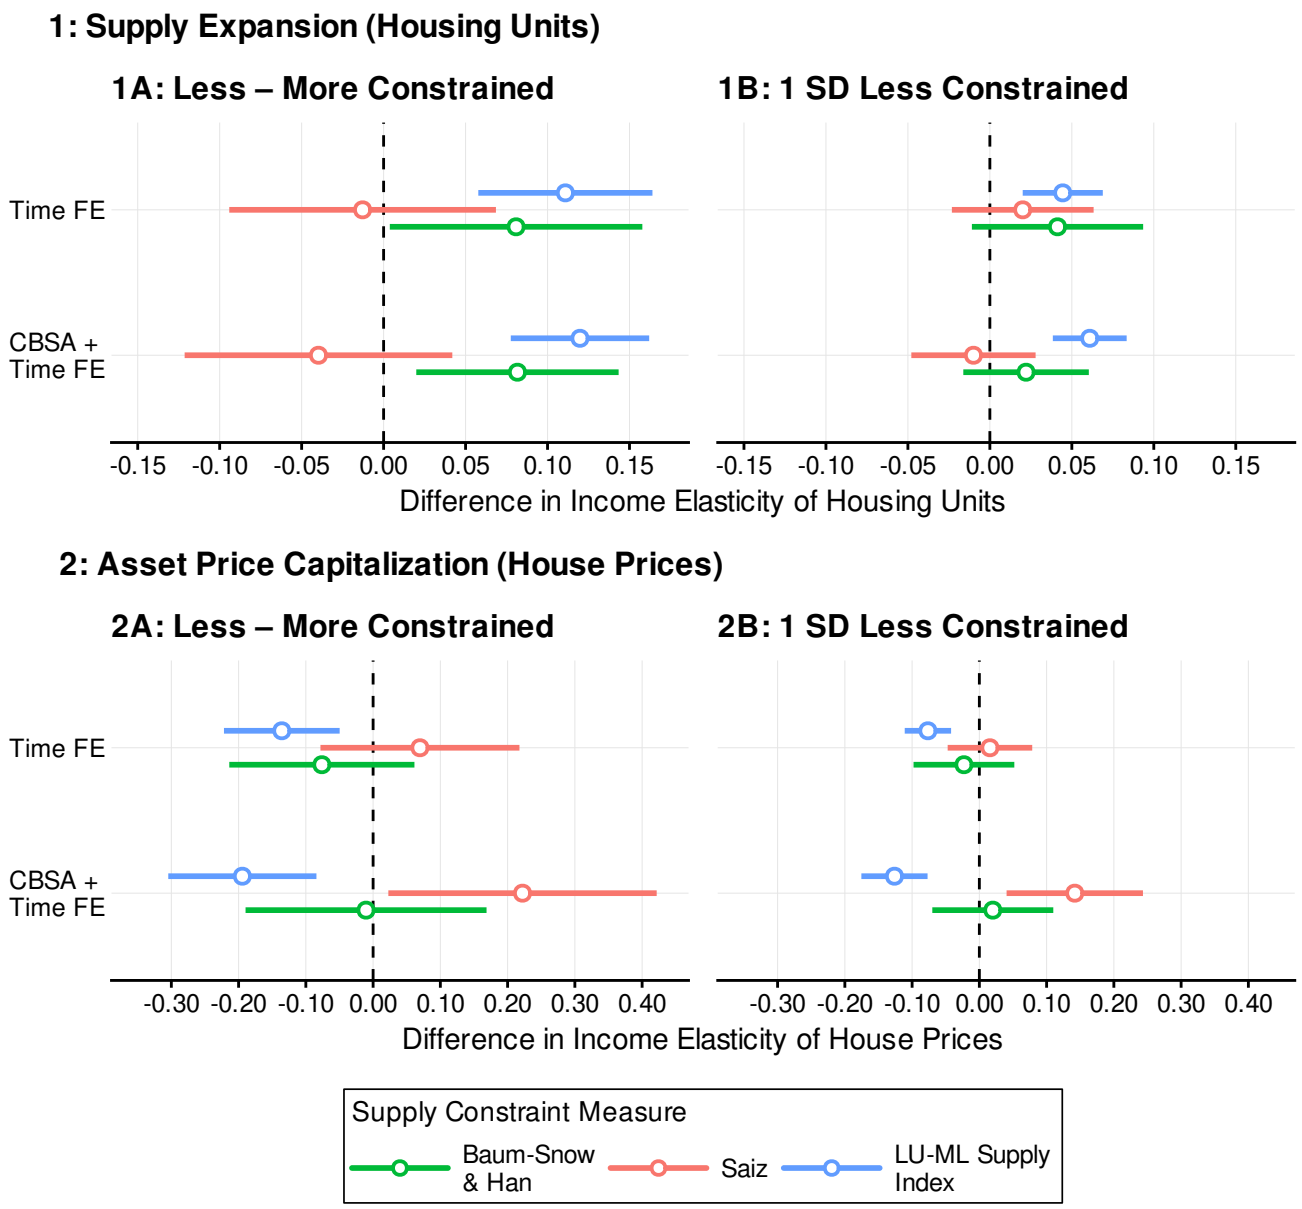
Notes: Panels A and B show Google Maps imagery for Chicago and Las Vegas along with the 50km radius circle around the principal city used by Saiz (2010) and Baum-Snow and Han (BSH, 2024). Saiz and BSH calculate land unavailability as the unweighted sum of the share of undevelopable land due to steep slopes, water, and wetlands within the 50km radius circle. Using this method, Chicago and Las Vegas have nearly identical Land Unavailability (39.5% and 38.2%). Panel C decomposes the LU-ML Price Pressure Index for Chicago and Las Vegas using SHAP values (Lundberg and Lee, 2017; Lundberg et al., 2020). The LU-ML Price Pressure Index is trained using 2002–2024 Zillow cumulative house price growth as the target variable and employs 90 total features (slope, water, and wetlands LU measured using 30 different spatial definitions). The SHAP value, $\varphi_i^{(j)}$, measures the marginal contribution of feature j to the LU-ML Index for CBSA i relative to a coalition of other features ($\neq j$), according to a fair distribution of payouts from cooperative game theory (Shapley, 1953). This allows for an additive decomposition of the index: LU-ML Index$_i = \mu + \sum_{j=1}^{M} \varphi_i^{(j)}$, where μ represents the unconditional mean prediction across all CBSAs. Text below each city label reports the total predicted growth (LU-ML Index), the predicted growth in excess of the national mean ($\sum_j \varphi^{(j)}$), and the prediction error (Residual). The arrows visualize the SHAP contributions: we aggregate the SHAP values from the 90 underlying features into three categories—Slope (red), Water (blue), and Wetlands (teal). The length of each arrow represents the feature set's contribution to the deviation from the mean prediction (μ), measured in percentage points (pp). Labels on the arrows report the underlying land unavailability percentages for that category, with national percentiles in parentheses.

Figure 10: Global Feature Contributions and Nonlinearities in the LU-ML Price Pressure Index (2002–2024)



Notes: The LU-ML Price Pressure Index is trained using 2002–2024 Zillow cumulative house price growth for CBSAs as the target variable and employs 90 total features (slope, water, and wetlands LU measured using 30 different spatial definitions). We calculate marginal predictive impacts using SHAP values (Lundberg and Lee, 2017; Lundberg et al., 2020). The SHAP value, $\varphi_i^{(j)}$, measures the marginal contribution of LU feature j to the LU-ML Index for CBSA i relative to a coalition of other features ($\neq j$), according to a fair distribution of payouts from cooperative game theory (Shapley, 1953). Panels 1A–1C illustrate aggregate nonlinearities. For each CBSA, we sum the SHAP values associated with a specific LU type (Slope, Water, or Wetlands) and plot this net contribution (vertical axis) against the mean LU share for that type (horizontal axis). The fitted lines show linear (green) and LOESS (blue) relationships in the average marginal impact across the distribution of land unavailability. Panel 2 compares feature importance using two different aggregation methods. Panel 2A reports the average magnitude of the *net* impact for each LU type or interaction. Specifically, we sum the SHAP values for all features within a category (e.g., Slope) to find the net impact for each CBSA, then average the absolute values of these sums across all CBSAs. Panel 2B reports the *gross* predictive contribution of each spatial definition (polygon shape). Specifically, we sum the absolute SHAP values for all features derived from a specific polygon (e.g., Principal City 100km Radius) within each CBSA, then average these sums across all CBSAs. For the Freddie Mac and FHFA data, see Appendix Figures A4 and A5.

Figure 11: Mitigation Effects (β_3) by Supply Proxy, CBSA Decadal Panel (1990–2020)



Notes: This figure displays the estimated differences in income elasticities due to looser supply constraints (the mitigation effect) for housing units (Panel 1) and FHFA house prices (Panel 2) across CBSAs by supply constraint measure. Plotted estimates correspond to β_3 from the following panel specification:

$$\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 Constraint_i + \beta_3 [\Delta \log(Income_{it}) \times Constraint_i] + \epsilon_{it}$$

α_i represents CBSA fixed effects (in the bottom row of each plot) and δ_t represents time fixed effects. *Income* is total BEA income at the CBSA level. In Panels 1A and 2A, the constraint is a binary indicator equal to 1 for CBSAs with looser-than-average supply constraints, so that β_3 represents the difference in elasticity between less and more constrained markets. In Panels 1B and 2B, the constraint measures are continuous and standardized across the pooled sample to have zero mean and unit variance, meaning β_3 represents the change in elasticity for a one-standard-deviation relaxation in supply constraints. The LU-ML Supply Index is constructed using the log change in housing units from 1990–2020 (long-difference) as the target variable. Error bars correspond to ± 1.96 standard errors clustered at the CBSA level. See Appendix Table B3 for the full regression output.

Table 2: Income Elasticities using the LU-ML Index – Across vs Within CBSAs

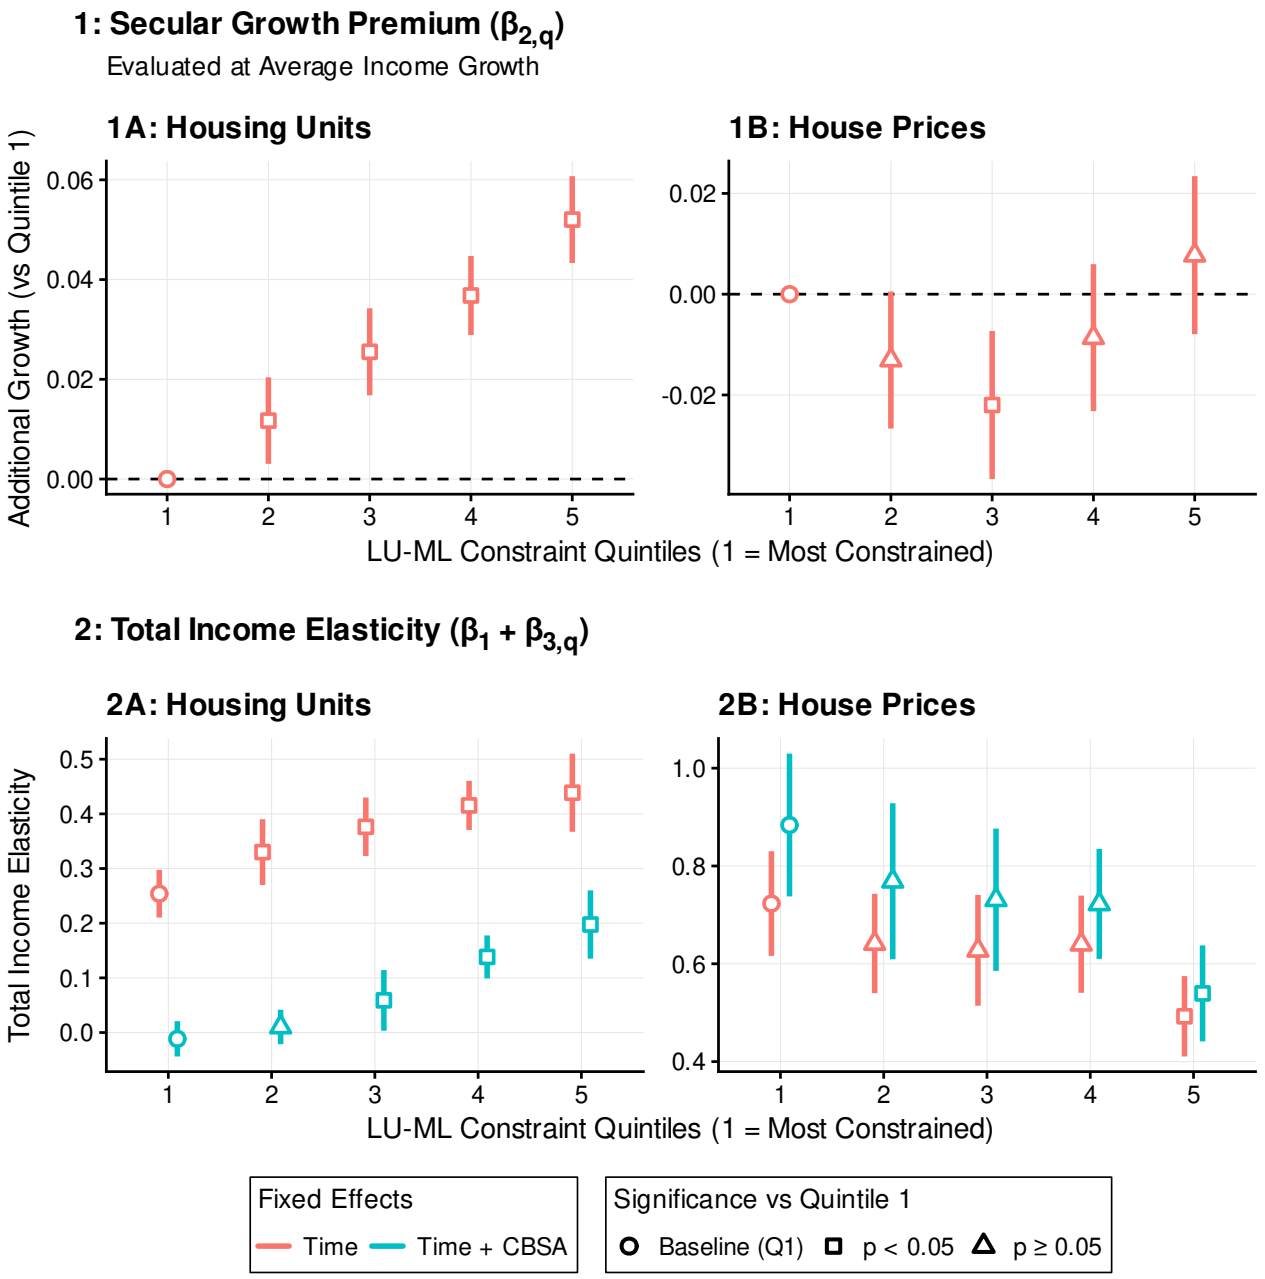
	Fixed Effects:	
	Time	Time + CBSA
	(1)	(2)
Panel A: LHS = Housing Unit Growth		
Total Income Growth (Demeaned)	0.360***	0.084***
	(0.015)	(0.014)
LU-ML Supply Index (Standardized)	0.018***	
	(0.001)	
Total Income Growth × LU-ML Supply Index	0.044***	0.061***
	(0.012)	(0.011)
Panel B: LHS = House Price Growth		
Total Income Growth (Demeaned)	0.631***	0.736***
	(0.026)	(0.038)
LU-ML Supply Index (Standardized)	0.005*	
	(0.003)	
Total Income Growth × LU-ML Supply Index	−0.076***	−0.126***
	(0.018)	(0.025)

Notes: This table reports the baseline income elasticity (β_1), the secular growth premium (β_2), and the mitigation effect (β_3) of looser supply constraints on housing units (Panel A) and FHFA house prices (Panel B), across CBSAs (Column 1) or within CBSAs (Column 2). Estimates are from decadal CBSA panel regressions (1990–2020) using the following specification:

$$\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 LU\text{-}ML_i + \beta_3[\Delta \log(Income_{it}) \times LU\text{-}ML_i] + \epsilon_{it}$$

where Y_{it} is housing units or house prices, α_i represents CBSA fixed effects (included in Column 2), and δ_t signifies decade fixed effects. The LU-ML Supply Index is constructed using the log change in housing units from 1990–2020 (long-difference) as the target variable. We standardize the LU-ML Index to have zero mean and unit variance, such that larger values signal looser constraints. $Income_{it}$ is total BEA income. We demean income growth across all U.S. CBSAs in the pooled BEA dataset so that β_2 measures the secular growth premium: the increase in housing unit or house price growth associated with looser constraints given average long-run city-level income growth. Because the constraint indicator is time-invariant, the CBSA fixed effects (α_i) absorb the secular growth premium (β_2) in Column (2). Robust standard errors clustered at the CBSA level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 12: The Secular Growth Premium and Mitigation Effect by LU-ML Quintile



Notes: This figure plots the secular growth premium (Panel 1) and total income elasticities (Panel 2) across quintiles of the LU-ML Supply Index, estimated using the following CBSA panel regression:

$$\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \sum_{q=2}^{5} \beta_{2,q} LU\text{-}ML_{iq} + \sum_{q=2}^{5} \beta_{3,q} [\Delta \log(Income_{it}) \times LU\text{-}ML_{iq}] + \epsilon_{it}$$

where α_i represents CBSA fixed effects (included in the blue estimates) and δ_t represents time fixed effects. $LU\text{-}ML_{iq}$ is an indicator equal to 1 if CBSA i falls into the q-th quintile of the LU-ML Supply Index, with Quintile 1 (most constrained) serving as the omitted baseline. The LU-ML Supply Index is constructed using the log change in housing units from 1990–2020 (long-difference) as the target variable. $Income_{it}$ is total BEA income at the CBSA level. We demean income growth across all U.S. CBSAs in the pooled BEA dataset so that the main effects ($\beta_{2,q}$) measure the secular growth premium: the baseline difference in housing unit or house price growth associated with looser constraints given average long-run national income growth. Because the LU-ML constraint quintiles are time-invariant, the CBSA fixed effects (α_i) absorb the secular growth premiums ($\beta_{2,q}$). Point shapes denote statistical significance relative to the Quintile 1 baseline ($\beta_{2,q}$ in Panel 1 and the mitigation effect, $\beta_{3,q}$, in Panel 2): squares indicate $p < 0.05$, and triangles indicate $p \geq 0.05$. Confidence intervals correspond to ± 1.96 robust standard errors clustered at the CBSA level. See Appendix Table B6 for the full regression output.

Table 3: Average Income Elasticities using the Time-Varying LU-ML Supply Index

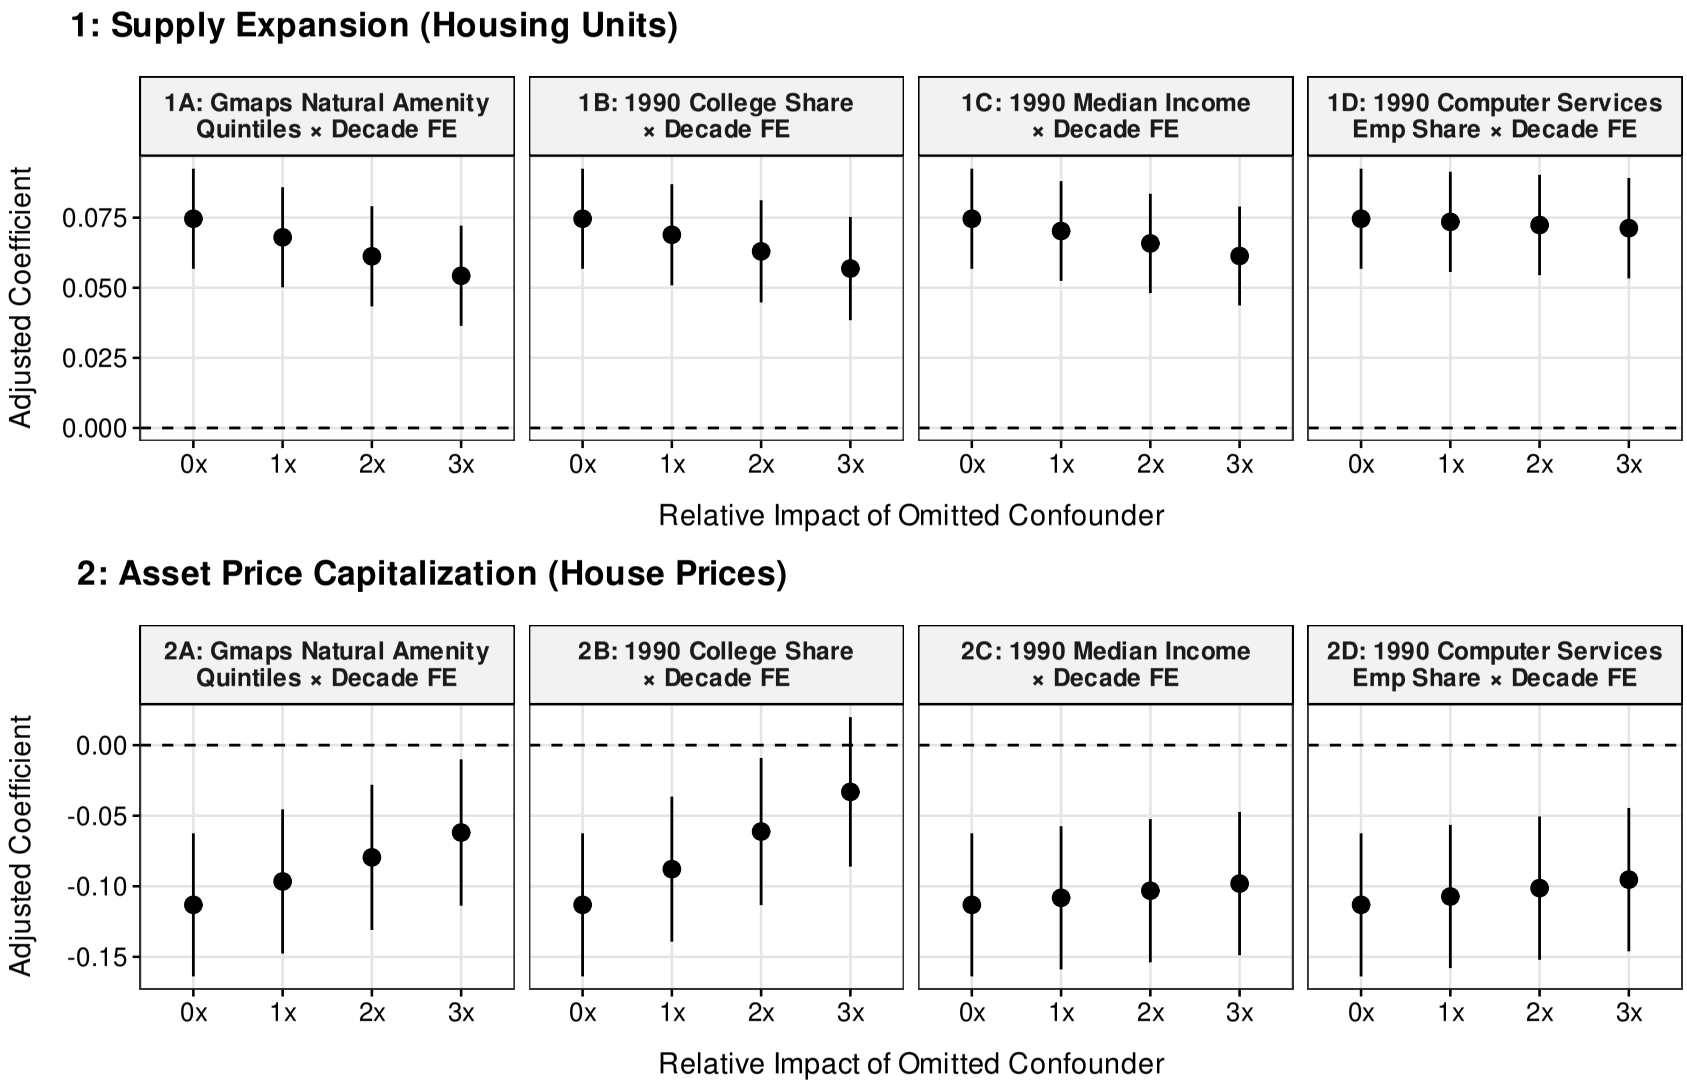
	(1)	(2)	(3)	(4)
Panel A: LHS = Housing Unit Growth				
Total Income Growth (Demeaned)	0.355***	0.070***	0.105***	0.118***
	(0.015)	(0.014)	(0.015)	(0.017)
LU-ML Supply Index (Standardized)	0.029***	0.032***	0.023***	0.010***
	(0.002)	(0.003)	(0.003)	(0.003)
Total Income Growth × LU-ML Supply Index	0.061***	0.071***	0.075***	0.071***
	(0.011)	(0.009)	(0.009)	(0.010)
Panel B: LHS = House Price Growth				
Total Income Growth (Demeaned)	0.642***	0.755***	0.811***	0.504***
	(0.024)	(0.035)	(0.043)	(0.035)
LU-ML Supply Index (Standardized)	0.012***	0.006	−0.001	0.003
	(0.004)	(0.008)	(0.008)	(0.006)
Total Income Growth × LU-ML Supply Index	−0.168***	−0.177***	−0.113***	−0.094***
	(0.022)	(0.026)	(0.026)	(0.019)
Decade Fixed Effects	✓	✓	✓	
CBSA Fixed Effects		✓	✓	✓
CBSA Controls × Decade FE			✓	✓
State × Decade Fixed Effects				✓

Notes: This table reports the results of CBSA decadal panel regressions (1990–2020) evaluating the impact of the time-varying LU-ML Supply Index on the income elasticities of housing units (Panel A) and FHFA house prices (Panel B). Estimates are from the following specification:

$$\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 LU\text{-}ML_{it} + \beta_3[\Delta \log(Income_{it}) \times LU\text{-}ML_{it}] + \mathbf{X}_i' \lambda_t + \epsilon_{it}$$

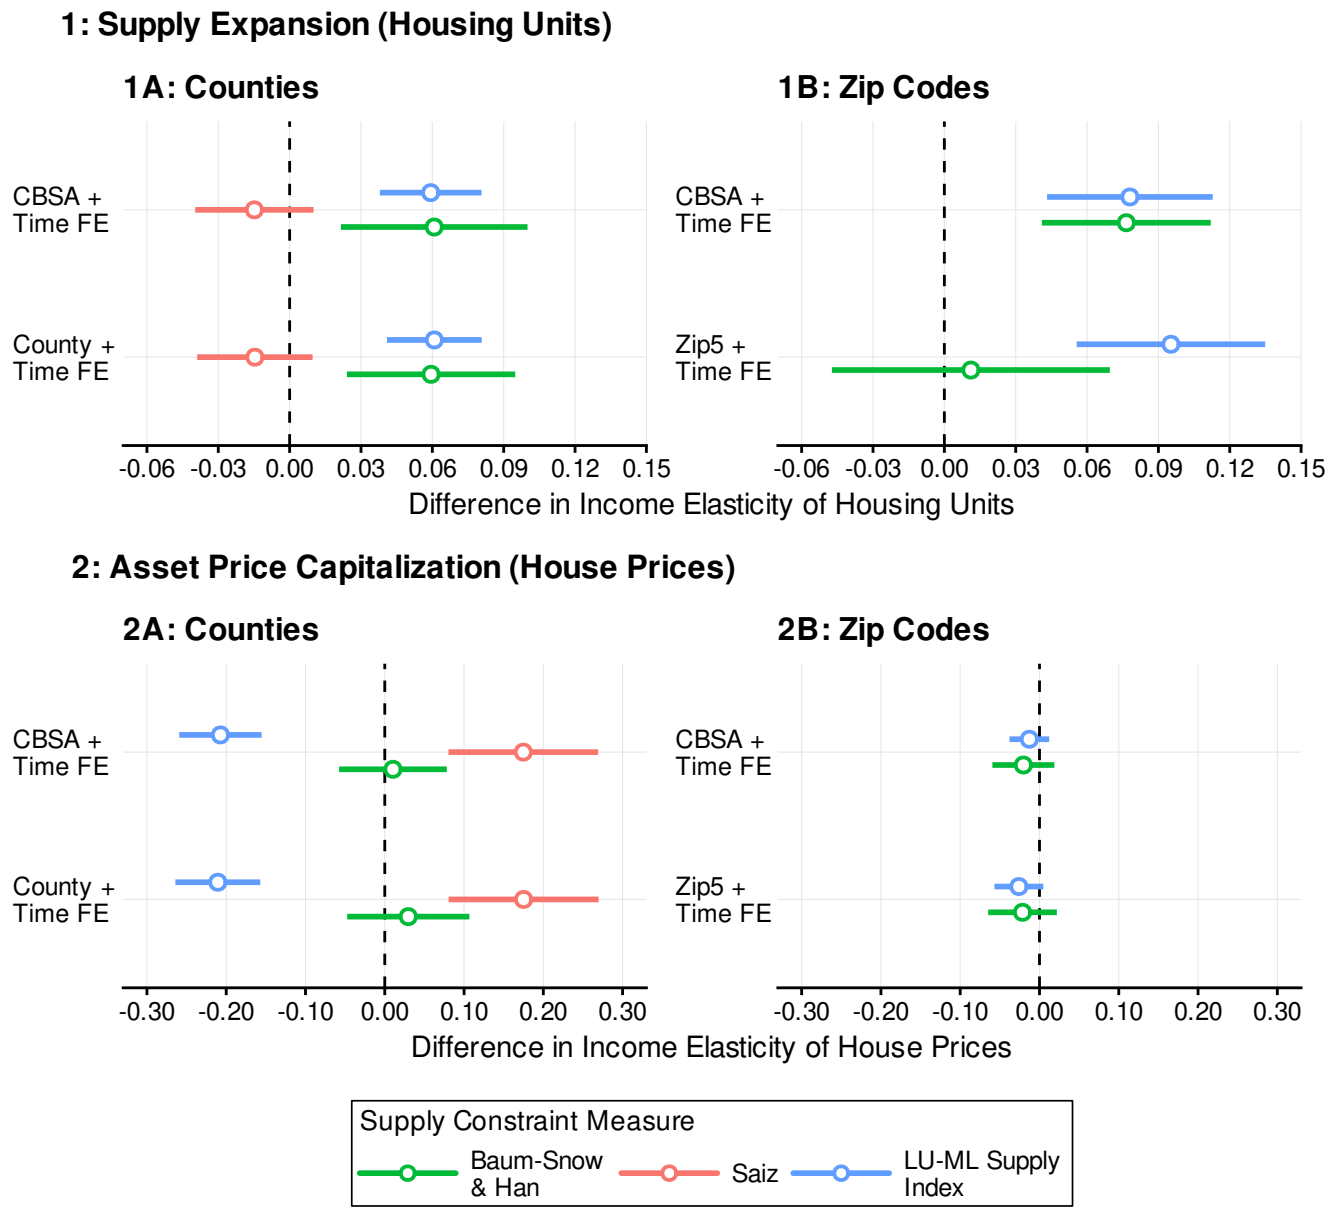
where α_i represents CBSA fixed effects (included in Columns 2–4) and δ_t represents decade fixed effects. We construct the LU-ML Supply Index using the decadal log change in housing units from 1990–2020 as the target variable. It is standardized over the pooled sample to have zero mean and unit variance; higher values indicate looser effective supply constraints. $Income_{it}$ is BEA total income. We demean income growth across all U.S. CBSAs in the pooled BEA dataset so that β_2 measures the secular growth premium: the increase in housing unit or house price growth associated with looser constraints given average long-run city-level income growth. The vector $\mathbf{X}_i' \lambda_t$ represents baseline CBSA-level controls interacted with decade fixed effects in Columns (3) and (4). These controls consist of natural amenity quintiles and 1990 levels of college share, median income, housing unit density, and local industry employment shares. Employment shares are based on 1986–1997 SIC codes and span computer services; construction; finance, insurance, and real estate (FIRE); manufacturing; retail; and services less computer services. Column (4) replaces decade fixed effects with state-by-decade fixed effects. Robust standard errors clustered at the CBSA level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 13: Stability of the LU-ML Mitigation Effect (β_3) to Unobserved Spatial Demand



Notes: This figure plots the stability of the LU-ML Supply Index mitigation effect (β_3) on housing units (Panel 1) and FHFA house prices (Panel 2) to unobserved spatial demand (Cinelli and Hazlett, 2020), evaluated via the following model: $\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 LU\text{-}ML_{it} + \beta_3[\Delta \log(Income_{it}) \times LU\text{-}ML_{it}] + \mathbf{X}_i'\lambda_t + \epsilon_{it}$, where α_i represents CBSA fixed effects and δ_t represents decade fixed effects. We construct the LU-ML Supply Index using the decadal log change in housing units from 1990–2020 as the target variable. It is standardized over the pooled sample to have zero mean and unit variance; higher values indicate looser supply constraints. $Income_{it}$ is BEA total income. The vector $\mathbf{X}_i'\lambda_t$ represents baseline CBSA-level controls interacted with decade fixed effects. See Table 3 for the full regression results and list of controls. The unadjusted baseline (0x on the horizontal axis) reports β_3 assuming no unobserved confounding. The bounds simulate an unobserved spatial demand confounder that has 1, 2, or 3 times the explanatory power of the chosen benchmark covariates (listed in the sub-panel titles). Because the CBSA fixed effects absorb the base-decade (1990s) interaction for each benchmark, the benchmarks reflect the explanatory power of the time-varying slopes for the 2000s and 2010s. To account for spatial clustering by CBSAs, we scale the adjusted IID standard errors by the ratio of the baseline clustered standard error to the baseline IID standard error. This assumes that the intra-cluster correlation does not change with the addition of the omitted variable. Confidence intervals correspond to ±1.96 adjusted standard errors. See Appendix Table B8 for the full estimation output.

Figure 14: Intra-City Mitigation Effects (β_3) by Supply Proxy

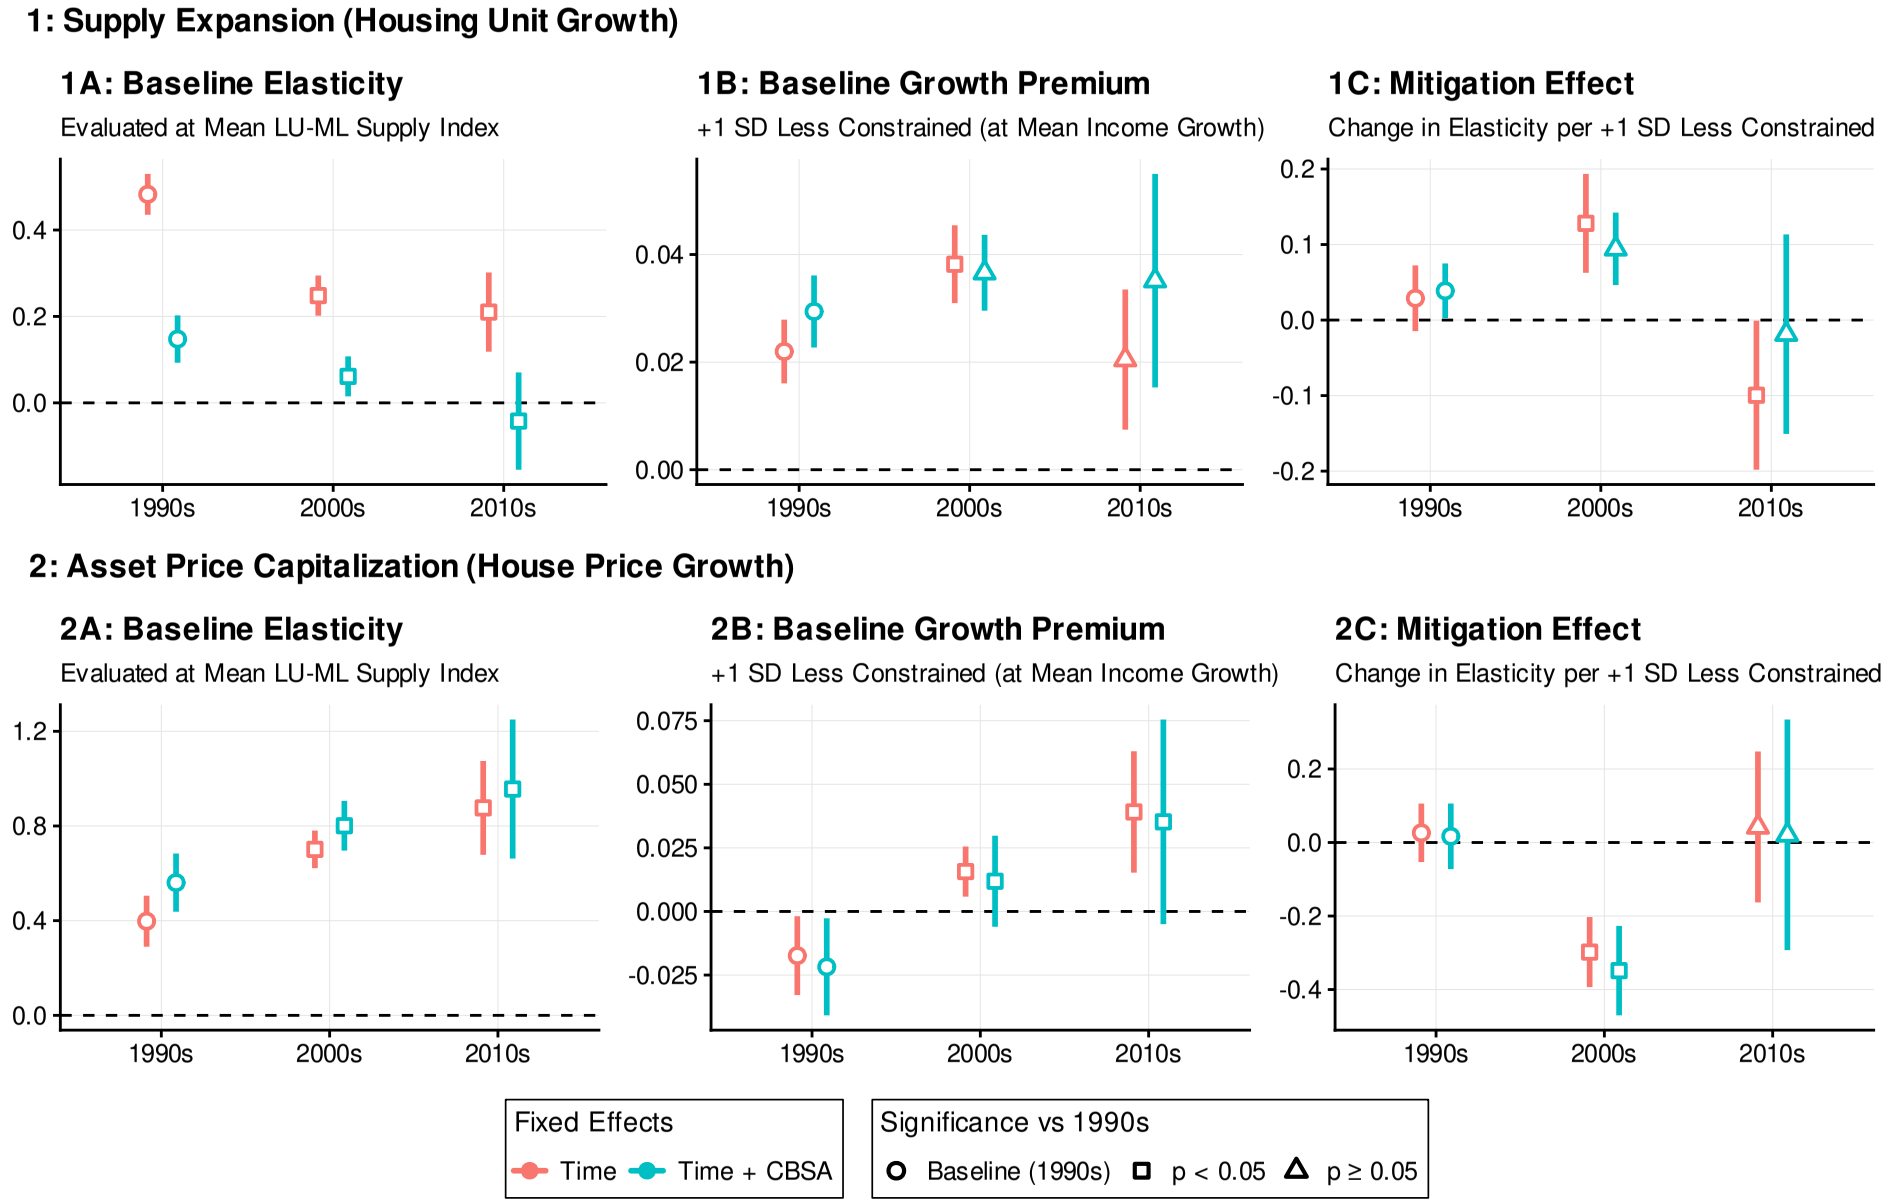


Notes: This figure displays the estimated differences in income elasticities due to looser supply constraints (the mitigation effect) for housing units (Panel 1) and FHFA house prices (Panel 2) within CBSAs across counties or zip codes. Plotted estimates correspond to β_3 from the panel regression:

$$\Delta \log(Y_{i,c,t}) = \alpha_i + \gamma_c + \delta_t + \beta_1 \Delta \log(Income_{c,t}) + \beta_2 Constraint_{i,c,t} + \beta_3[\Delta \log(Income_{c,t}) \times Constraint_{i,c,t}] + \epsilon_{i,c,t}$$

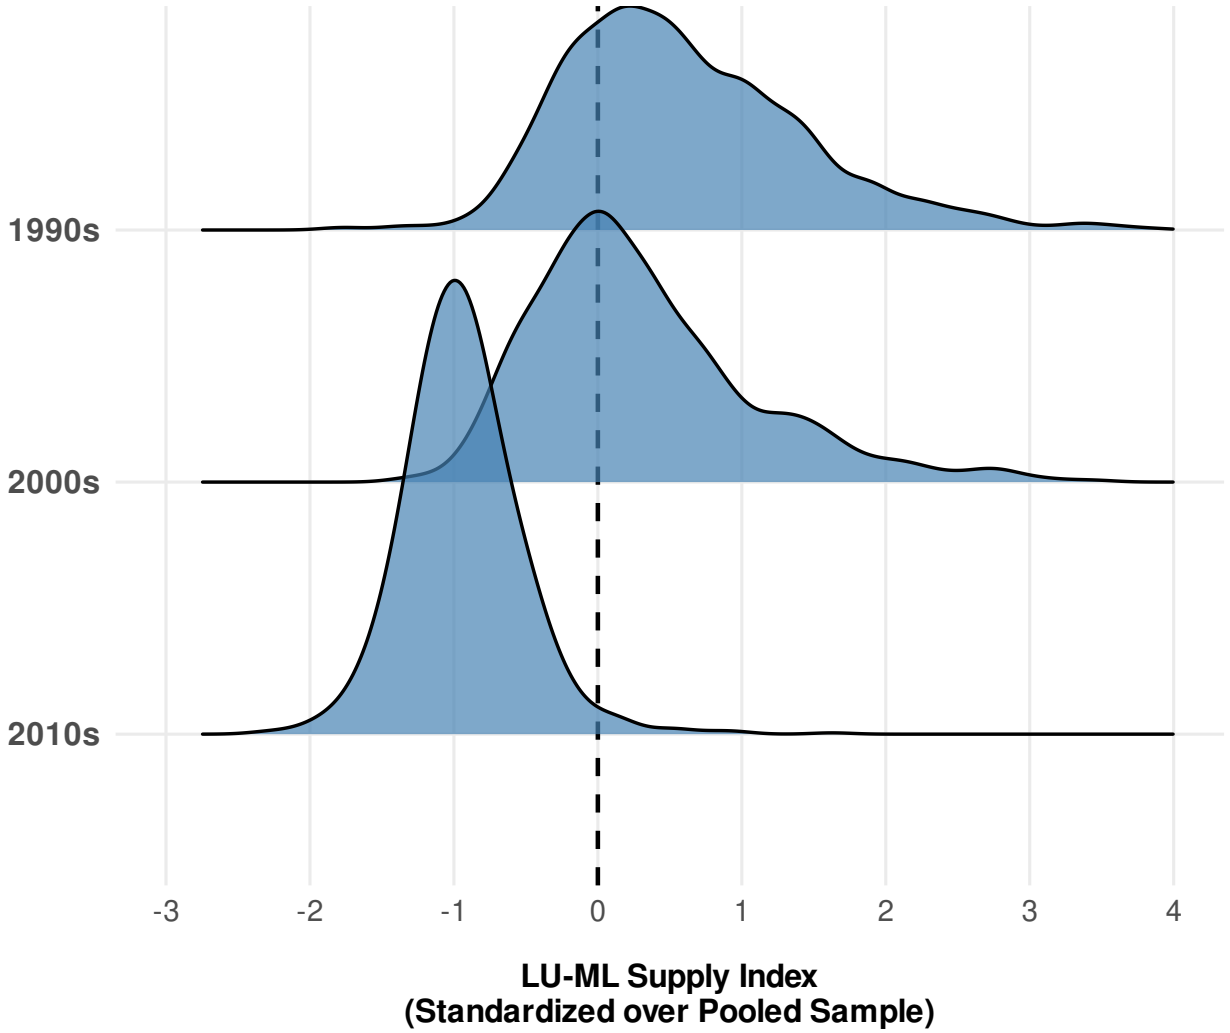
$Y_{i,c,t}$ is housing units or house prices in zip code or county i in CBSA c measured at the end of decade t. The specification includes either CBSA fixed effects (γ_c, top row of each panel) or local fixed effects (α_i for county or zip code, bottom row of each panel), along with decade fixed effects (δ_t). $Income_{c,t}$ is BEA total income for CBSA c in decade t. $Constraint_{i,c,t}$, the constraint proxy for zip code or county i in CBSA c, is either the Baum-Snow and Han (2024, BSH) measure, the Saiz (2010) measure, or our time-varying LU-ML Supply Index. All constraint indices are standardized to zero mean and unit variance, and signed such that larger values indicate looser constraints. The county-level LU-ML Supply Index targets county-level housing units and employs a spatial cross-fitting design by training exclusively on counties outside of the target county's 2020 Commuting Zone to prevent contamination from local labor markets. The zip-code-level LU-ML Index similarly targets the decadal log change in zip-code housing units. Because the BSH and Saiz proxies are time-invariant, the county or zip code fixed effects (α_i) absorb their secular growth premiums (β_2) when applied. Error bars correspond to ± 1.96 robust standard errors clustered at the CBSA level. See Appendix Table B9 for the full regression output.

Figure 15: Time-Varying Income Elasticities, Secular Growth Premiums, and Mitigation Effects, 1990–2020



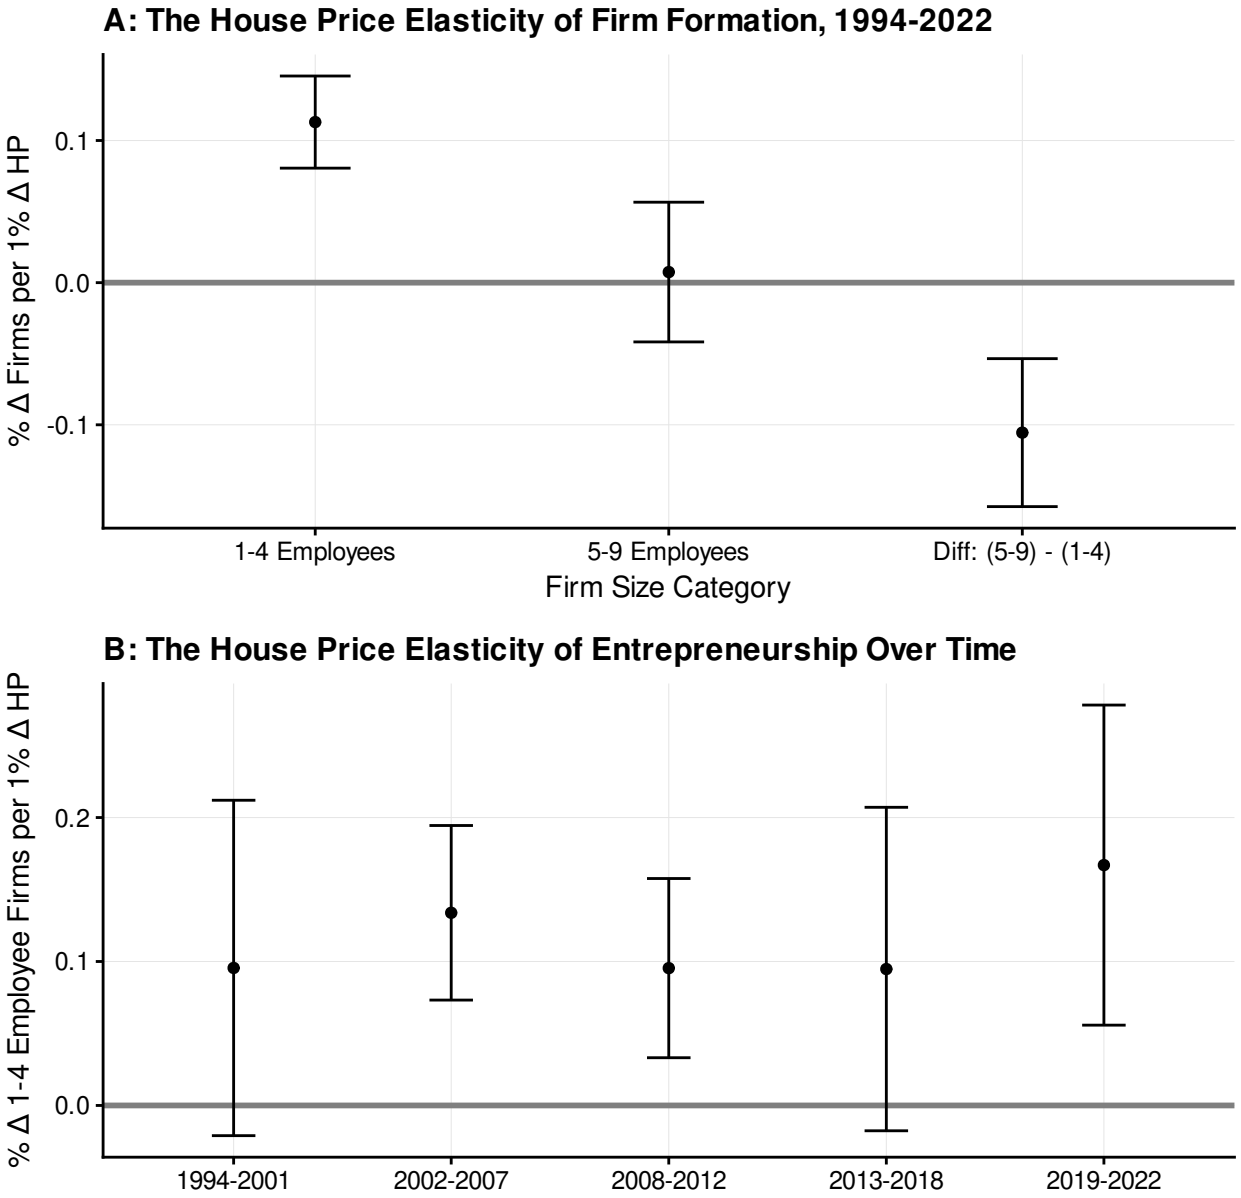
Notes: This figure displays the results of CBSA decadal panel regressions evaluating the time-varying impact of the LU-ML Supply Index on the income elasticities of housing units and FHFA house prices using the following specification: $\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_{1t}\Delta \log(Income_{it}) + \beta_{2t} LU\text{-}ML_{it} + \beta_{3t}[\Delta \log(Income_{it}) \times LU\text{-}ML_{it}] + \epsilon_{it}$, where α_i and δ_t represent CBSA and decade fixed effects, respectively. The time-varying parameters (β_{kt}) are estimated via interactions with decade indicator variables, where the baseline coefficients correspond to the omitted 1990s decade. The LU-ML Supply Index is standardized over the pooled sample (zero mean, unit variance); higher values indicate looser supply constraints. Because $\Delta \log(Income_{it})$ is demeaned across all U.S. CBSAs, β_{2t} isolates the baseline growth premium associated with looser constraints given average income growth. Without CBSA fixed effects (red), this measures the cross-sectional secular growth premium; with CBSA fixed effects (blue), it isolates the transitory growth premium—the deviation from a city's secular growth trend when its constraints loosen. Confidence intervals correspond to ±1.96 robust standard errors clustered at the CBSA level. See Table B10 for the full regression output.

Figure 16: Densities of the Time-Varying LU-ML Supply Index by Decade



Notes: This figure displays empirical densities of the time-varying LU-ML Supply Index by decade. The time-varying LU-ML Supply Index targets log decadal change in housing units by CBSA. We standardize the LU-ML Supply Index over the pooled sample (zero mean, unit variance). Higher values indicate looser supply constraints. For these densities weighted by 1990 housing units, see Appendix Figure B3.

Figure 17: The Housing Collateral Channel For Entrepreneurship (2SLS Estimates, 1994–2022)



Notes: This figure shows the impact of house price growth on entrepreneurship, proxied by the growth in firms with 1–4 employees. The estimates are derived from the following two-stage least squares (2SLS) specification:
$\Delta \log(\text{Firms}^{1-4}_{i,d,t}) = \alpha_i + \delta_{d,t} + \beta \cdot \Delta \log(\text{HP}_{i,d,t}) + \epsilon_{i,d,t}$
where $\Delta \log(\text{Firms}^{1-4}_{i,d,t})$ is the year-over-year (YoY) log first-difference in the number of firms with 1–4 employees in zip code i and Census division d at year t. α_i represents zip code fixed effects, $\delta_{d,t}$ signifies Census division-time fixed effects, and $\Delta \log(\text{HP}_{i,d,t})$ is the YoY log difference in zip code FHFA house prices. We instrument house price growth with the LU-ML Price Pressure Index. Panel A plots the baseline estimate for β (see Appendix Table C1 for the full results), alongside a placebo test using the growth in firms with 5–9 employees ($\Delta \log(\text{Firms}^{5-9}_{i,d,t})$), and the fully stacked difference between the two elasticities. Panel B plots the time-varying elasticity, estimated by interacting both the endogenous house price variable and the LU-ML instrument with time-period indicators. All confidence bands represent ±1.96 standard errors clustered at the commuting zone level.

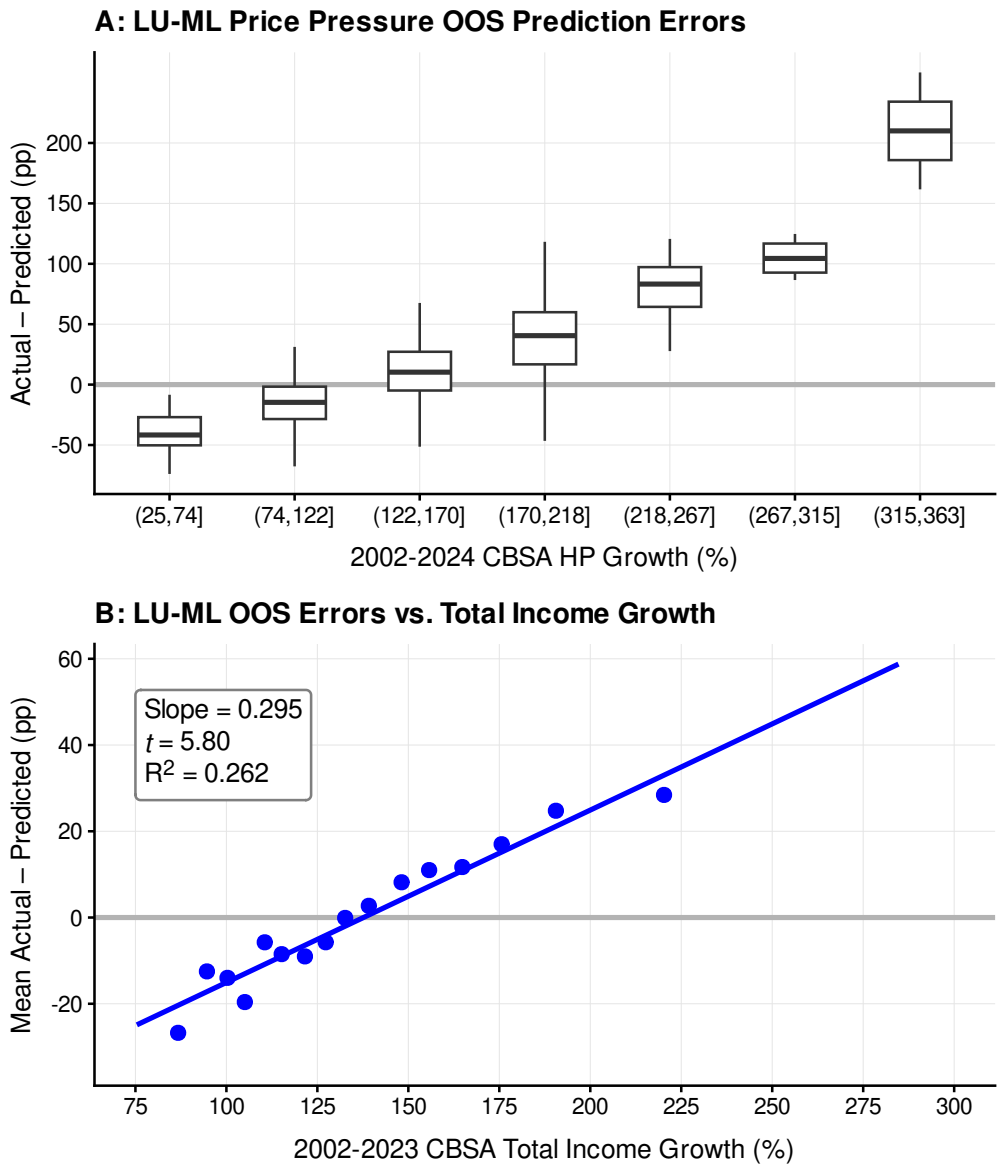
A Appendix: The LU-ML Price Pressure Index

Table A1: Explanatory Power of Supply Constraint Measures in Panel Regressions

Geographic Level	House Price Dataset	Baum-Snow & Han	Saiz	LU-ML Price Pressure Index
Panel A: 2000–2007				
CBSA	FHFA	$R^2 = 0.00$, N = 332	$R^2 = 0.03$, N = 256	$R^2 = 0.14$, N = 910
CBSA	Freddie Mac	$R^2 = 0.01$, N = 314	$R^2 = 0.03$, N = 251	$R^2 = 0.21$, N = 380
Census Tract	FHFA	$R^2 = 0.00$, N = 41806	$R^2 = 0.02$, N = 40659	$R^2 = 0.10$, N = 50172
County	FHFA	$R^2 = 0.00$, N = 711	$R^2 = 0.02$, N = 1582	$R^2 = 0.11$, N = 2396
Zip5	FHFA	$R^2 = 0.00$, N = 11678	$R^2 = 0.03$, N = 11406	$R^2 = 0.08$, N = 14870
Panel B: 2007–2011				
CBSA	FHFA	$R^2 = 0.00$, N = 332	$R^2 = 0.01$, N = 256	$R^2 = 0.09$, N = 915
CBSA	Freddie Mac	$R^2 = 0.00$, N = 314	$R^2 = 0.02$, N = 251	$R^2 = 0.15$, N = 380
Census Tract	FHFA	$R^2 = 0.00$, N = 45624	$R^2 = 0.01$, N = 44911	$R^2 = 0.07$, N = 57720
County	FHFA	$R^2 = 0.00$, N = 714	$R^2 = 0.01$, N = 1681	$R^2 = 0.04$, N = 2707
Zip5	FHFA	$R^2 = 0.00$, N = 13000	$R^2 = 0.01$, N = 12810	$R^2 = 0.06$, N = 18008
Panel C: 2000–2011				
CBSA	FHFA	$R^2 = 0.00$, N = 332	$R^2 = 0.05$, N = 256	$R^2 = 0.19$, N = 910
CBSA	Freddie Mac	$R^2 = 0.00$, N = 314	$R^2 = 0.05$, N = 251	$R^2 = 0.24$, N = 380
Census Tract	FHFA	$R^2 = 0.00$, N = 40373	$R^2 = 0.05$, N = 39215	$R^2 = 0.12$, N = 48636
County	FHFA	$R^2 = 0.00$, N = 711	$R^2 = 0.05$, N = 1582	$R^2 = 0.17$, N = 2395
Zip5	FHFA	$R^2 = 0.01$, N = 11633	$R^2 = 0.06$, N = 11356	$R^2 = 0.12$, N = 14813
Panel D: 2000–2024				
CBSA	FHFA	$R^2 = 0.00$, N = 332	$R^2 = 0.03$, N = 256	$R^2 = 0.13$, N = 908
CBSA	Freddie Mac	$R^2 = 0.00$, N = 314	$R^2 = 0.03$, N = 251	$R^2 = 0.20$, N = 380
Census Tract	FHFA	$R^2 = 0.00$, N = 29638	$R^2 = 0.02$, N = 28871	$R^2 = 0.08$, N = 35554
County	FHFA	$R^2 = 0.00$, N = 710	$R^2 = 0.02$, N = 1562	$R^2 = 0.11$, N = 2334
Zip5	FHFA	$R^2 = 0.00$, N = 10453	$R^2 = 0.02$, N = 10195	$R^2 = 0.08$, N = 12979
Panel E: 2012–2024				
CBSA	FHFA	$R^2 = 0.00$, N = 332	$R^2 = 0.00$, N = 256	$R^2 = 0.04$, N = 911
CBSA	Freddie Mac	$R^2 = 0.00$, N = 314	$R^2 = 0.00$, N = 251	$R^2 = 0.10$, N = 380
Census Tract	FHFA	$R^2 = 0.00$, N = 31544	$R^2 = 0.00$, N = 30913	$R^2 = 0.01$, N = 38454
County	FHFA	$R^2 = 0.00$, N = 713	$R^2 = 0.00$, N = 1636	$R^2 = 0.02$, N = 2526
Zip5	FHFA	$R^2 = 0.00$, N = 10884	$R^2 = 0.00$, N = 10651	$R^2 = 0.02$, N = 13790

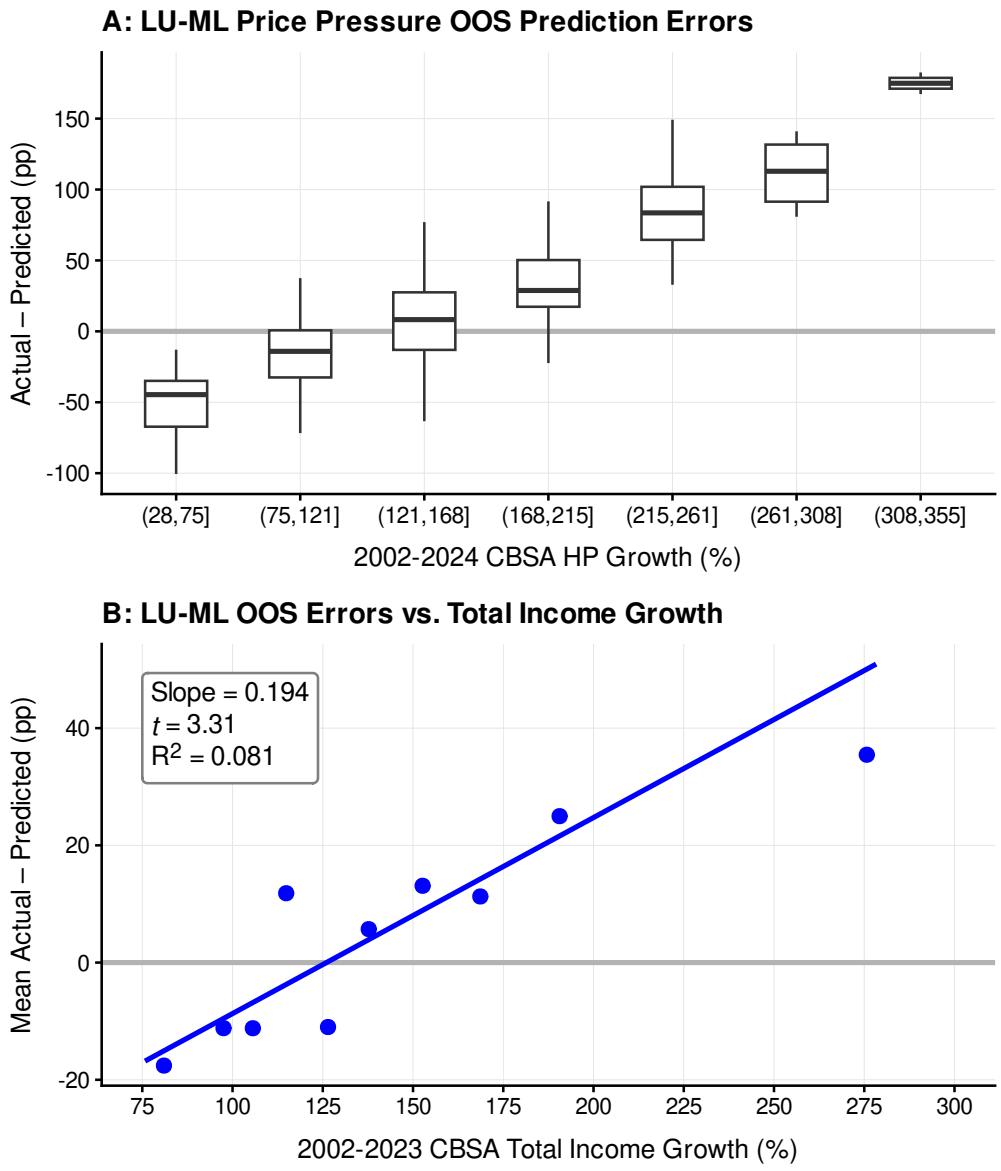
Notes: This table shows partial R^2 statistics and the number of geographic (cross-sectional) observations using supply constraint measures in house price panels. We compare the Saiz and BSH elasticity proxies against the LU-ML Price Pressure Index. House prices panels are from Freddie Mac and the FHFA, aggregated to various geographic levels including CBSAs, counties, zip codes, and census tracts. Estimates are derived from: $\Delta \log HP_{i,d,t} = \psi_i + \gamma_{d,t} + \beta \cdot SC_{i,d,t} + \varepsilon_{i,d,t}$. $\Delta \log HP_{i,d,t}$ is the YoY log first-difference in house prices for geographic unit i (e.g., a zip code or CBSA) in Census Division d at month t. ψ_i denotes geographic unit fixed effects, and $\gamma_{d,t}$ denotes Census division-by-time fixed effects. $SC_{i,d,t}$ represents the supply constraint measure. To generate time-varying variables for the static Saiz and BSH proxies, we interact them with the YoY log change in national house prices, following Guren et al. (2021). The LU-ML Price Pressure Index is constructed for each specific dataset using the corresponding YoY price change as the target variable. For each regression, we report the partial R^2 on the supply constraint measure and the number of geographic units included in the regression.

Figure A1: LU-ML Prediction Errors and Housing Demand Growth (FHFA Data)



Notes: This figure analyzes the out-of-sample (OOS) prediction errors of the LU-ML Price Pressure Index. Panel A displays boxplots of the OOS prediction errors (vertical axis) sorted by cumulative FHFA CBSA house price growth (horizontal axis) from 2002–2024. We construct the LU-ML Price Pressure Index using 2002–2024 FHFA cumulative house price growth as the target variable. Panel B presents a binscatter of the OOS prediction errors against 2002–2023 BEA total income growth. We select the number of bins following Cattaneo et al. (2024). The regression line and statistics are estimated on the underlying data.

Figure A2: LU-ML Prediction Errors and Housing Demand Growth (Zillow Data)



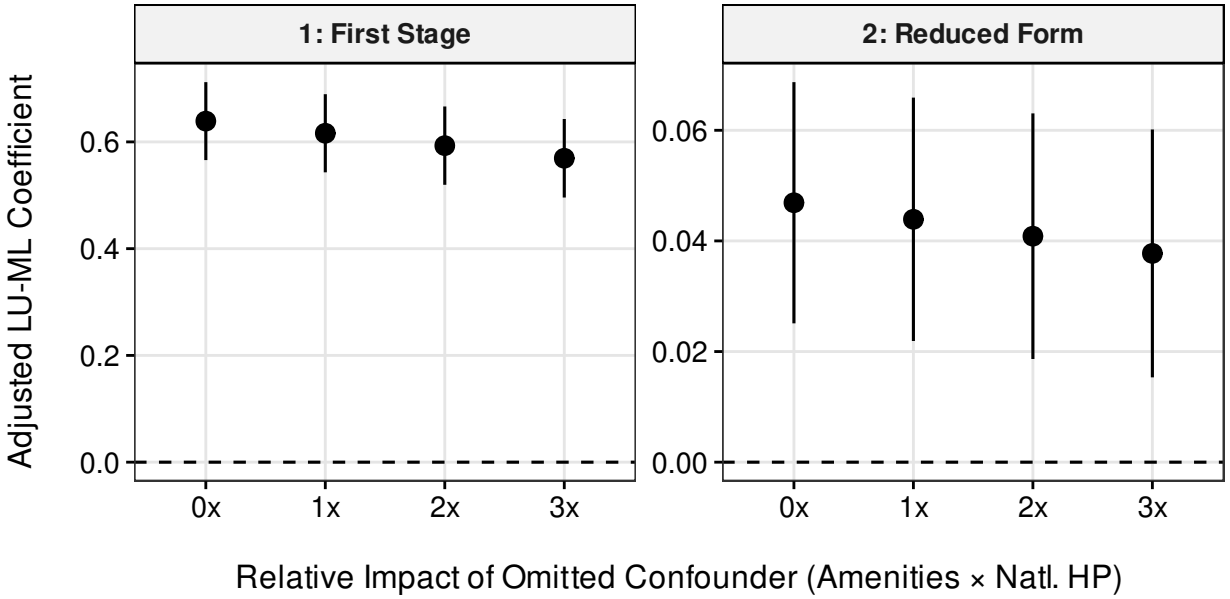
Notes: This figure analyzes the out-of-sample (OOS) prediction errors of the LU-ML Price Pressure Index. Panel A displays boxplots of the OOS prediction errors (vertical axis) sorted by cumulative Zillow CBSA house price growth (horizontal axis) from 2002–2024. We construct the LU-ML Price Pressure Index using 2002–2024 Zillow cumulative house price growth as the target variable. Panel B presents a binscatter of the OOS prediction errors against 2002–2023 BEA total income growth. We select the number of bins following Cattaneo et al. (2024). The regression line and statistics are estimated on the underlying data.

Table A2: Stability of the LU-ML Price Pressure Index to Unobserved Spatial Demand

Bound Scenario	$R^2_{D\sim Z\|X}$	$R^2_{Y\sim Z\|D,X}$	Adjusted Estimate	Robust SE	Adjusted t-stat
Unadjusted Baseline	–	–	0.508	(0.043)	11.79
1x Natural Amenities (Gmaps) × Date FE	0.038	0.029	0.461	(0.043)	10.66
2x Natural Amenities (Gmaps) × Date FE	0.076	0.058	0.413	(0.043)	9.49
3x Natural Amenities (Gmaps) × Date FE	0.115	0.087	0.362	(0.044)	8.28
1x College Share (2000) × Date FE	0.018	0.043	0.469	(0.042)	11.04
2x College Share (2000) × Date FE	0.035	0.087	0.430	(0.042)	10.26
3x College Share (2000) × Date FE	0.053	0.131	0.390	(0.041)	9.46
1x Median Income (2000) × Date FE	0.022	0.021	0.478	(0.043)	11.10
2x Median Income (2000) × Date FE	0.043	0.042	0.448	(0.043)	10.40
3x Median Income (2000) × Date FE	0.065	0.062	0.417	(0.043)	9.68
1x Computer Services Share (2000) × Date FE	0.004	0.005	0.501	(0.043)	11.64
2x Computer Services Share (2000) × Date FE	0.009	0.010	0.495	(0.043)	11.49
3x Computer Services Share (2000) × Date FE	0.013	0.015	0.488	(0.043)	11.35

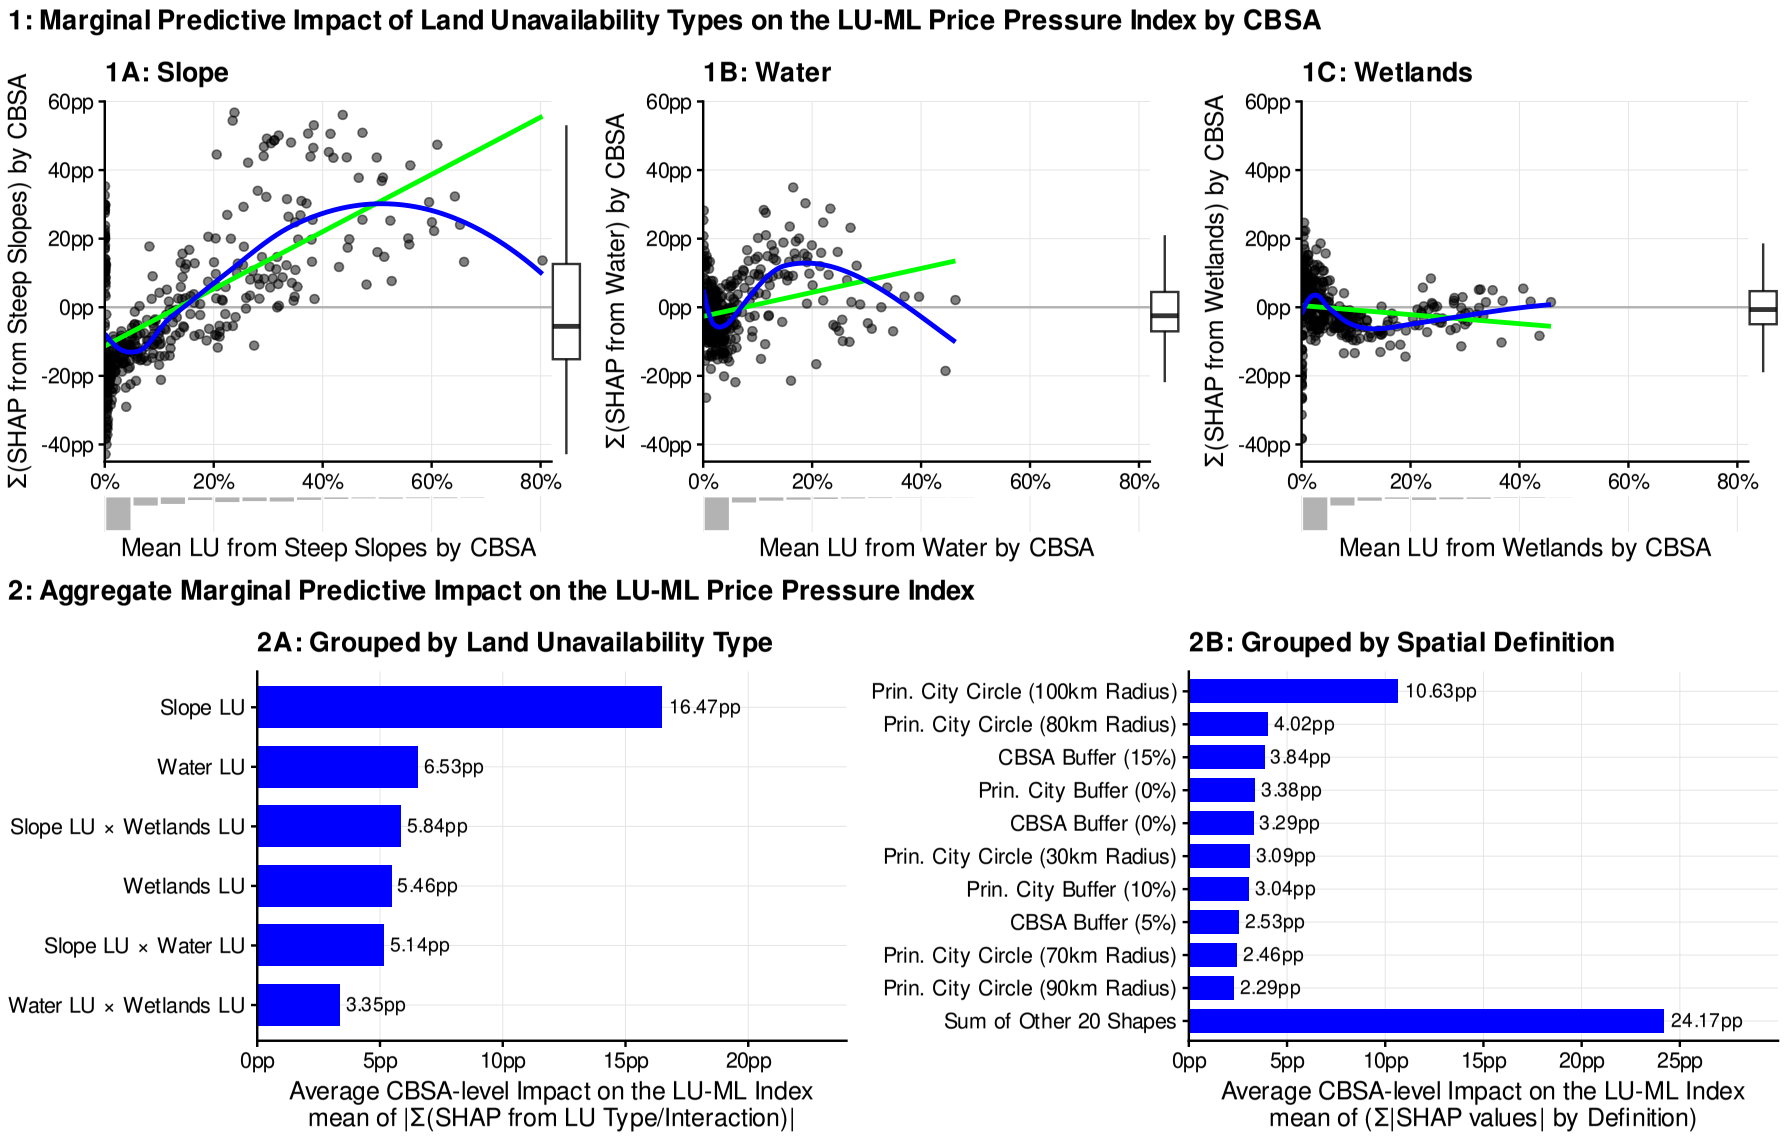
Notes: This table reports the stability of the LU-ML Price Pressure Index coefficient to unobserved spatial demand (Cinelli and Hazlett, 2020) using the model: $\Delta \log HP_{i,d,t} = \psi_i + \gamma_{d,t} + \beta \cdot LU\text{-}ML_{i,t} + \mathbf{X}_i'\lambda_t + \varepsilon_{i,d,t}$, where $\Delta \log HP_{i,d,t}$ is the YoY log change in Freddie Mac house prices for CBSA i in Census division d at time t. ψ_i and $\gamma_{d,t}$ are CBSA and Census division-by-date fixed effects, respectively. $\mathbf{X}_i'\lambda_t$ represents CBSA-level controls. We construct the LU-ML Price Pressure Index using YoY house price growth as the target variable. See Table 1 for the full regression results and the list of controls. The bounds simulate an unobserved spatial demand shock that has 1, 2, or 3 times the explanatory power of the chosen benchmark covariates; the benchmarks reflect the explanatory power of the time-varying slopes. $R^2_{D\sim Z|X}$ and $R^2_{Y\sim Z|D,X}$ represent the hypothetical partial R^2 of the unobserved confounder with the treatment and the outcome, respectively. To account for two-way spatial and temporal clustering, the adjusted IID standard errors are scaled by the ratio of the baseline clustered standard error to the baseline IID standard error. This assumes that the clustering penalty is approximately unchanged by the addition of the hypothetical omitted confounder.

Figure A3: Stability of the LU-ML Price Pressure Index In the First Stage and Reduced Form of Guren et al. (2021)



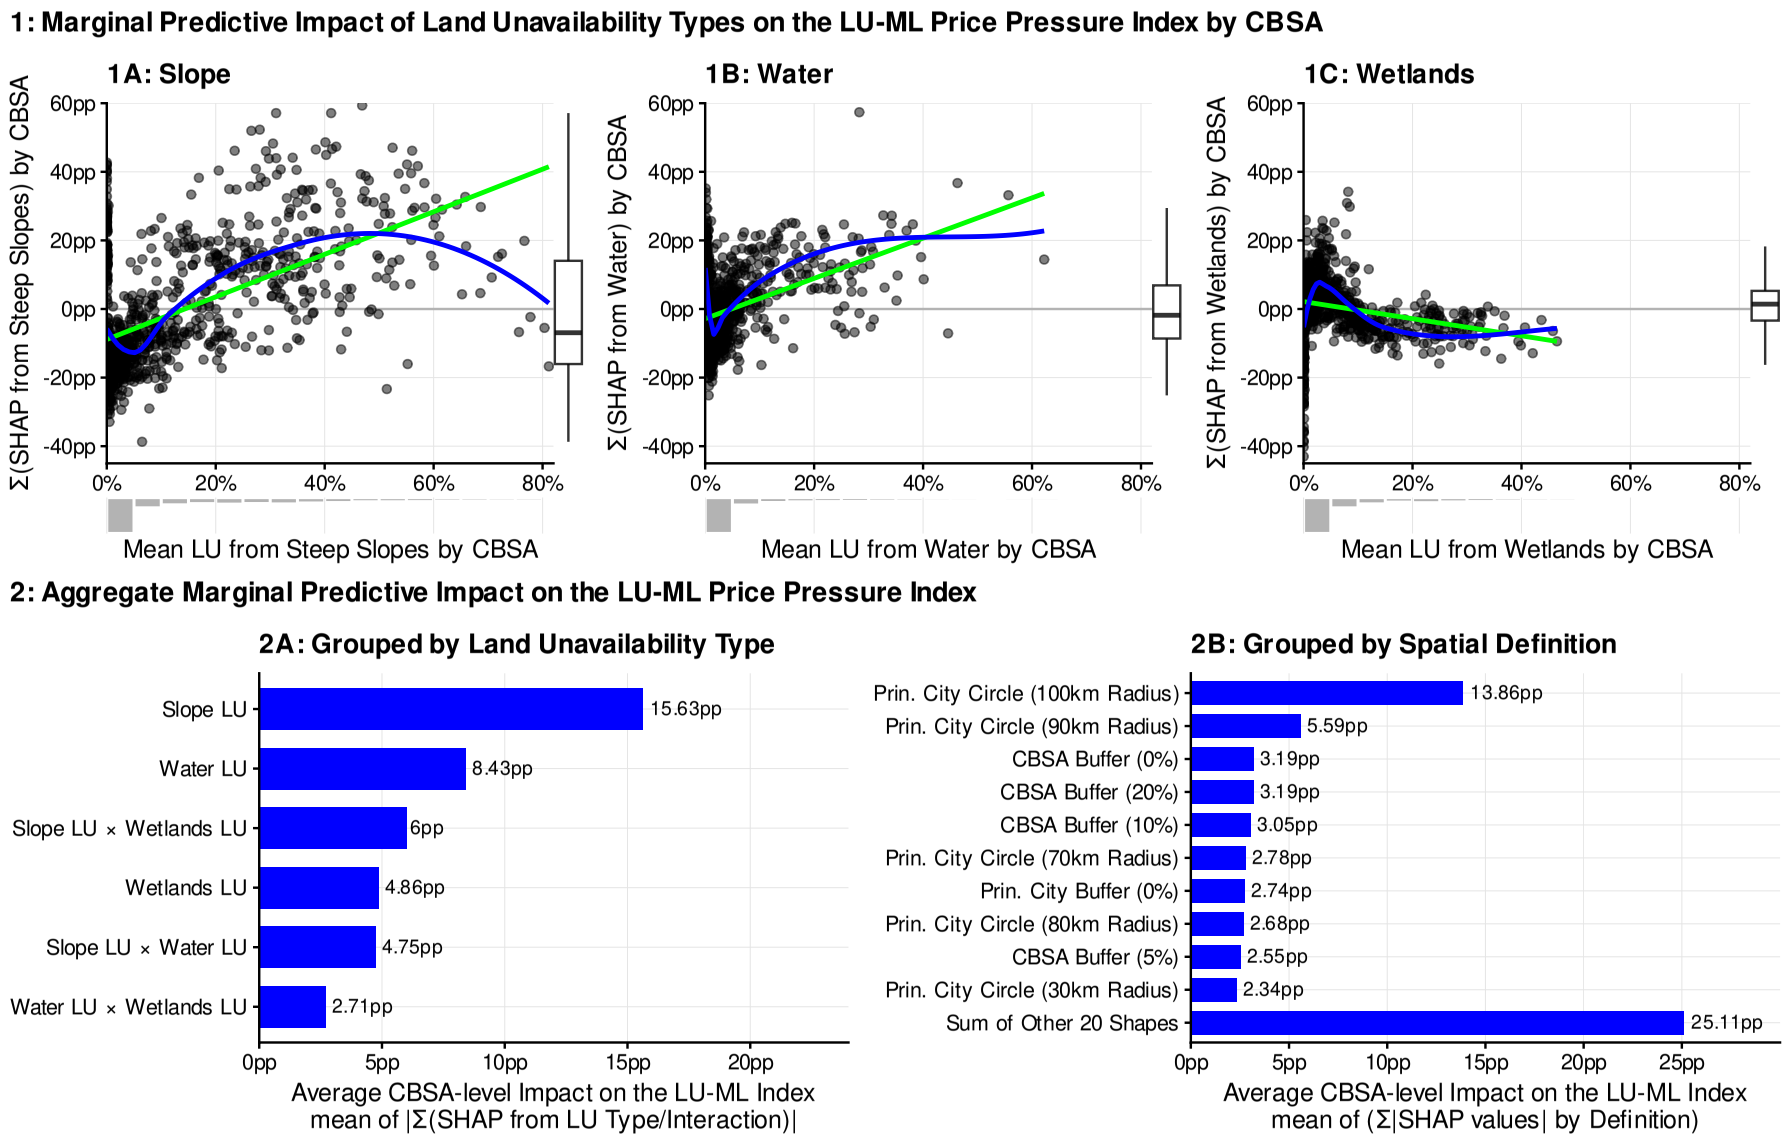
Notes: This figure plots the stability of the coefficient of the LU-ML Price Pressure Index to unobserved spatial demand (Cinelli and Hazlett, 2020), evaluated in the main first stage and reduced form equations in Guren et al. (2021). See Appendix Table D5. The unadjusted baseline (0x on the horizontal axis) reports the coefficient on the LU-ML Price Pressure Index assuming no unobserved confounding. The bounds simulate an unobserved spatial demand confounder that has 1, 2, or 3 times the explanatory power of natural amenity quintiles interacted with date fixed effects. To account for spatial clustering by CBSAs, we scale the adjusted IID standard errors by the ratio of the baseline clustered standard error to the baseline IID standard error. This assumes that the intra-cluster correlation does not change with the addition of the omitted variable. Confidence intervals correspond to ±1.96 adjusted standard errors.

Figure A4: Global Feature Contributions and Nonlinearities in the LU-ML Price Pressure Index (Freddie Mac Data, 2002–2024)



Notes: See the notes for Figure 10.

Figure A5: Global Feature Contributions and Nonlinearities in the LU-ML Price Pressure Index (FHFA Data, 2002–2024)



Notes: See the notes for Figure 10.

B Appendix: The LU-ML Supply Index and Income Elasticities

Table B1: Income Elasticity of House Prices by Supply Measure, 2000–2020

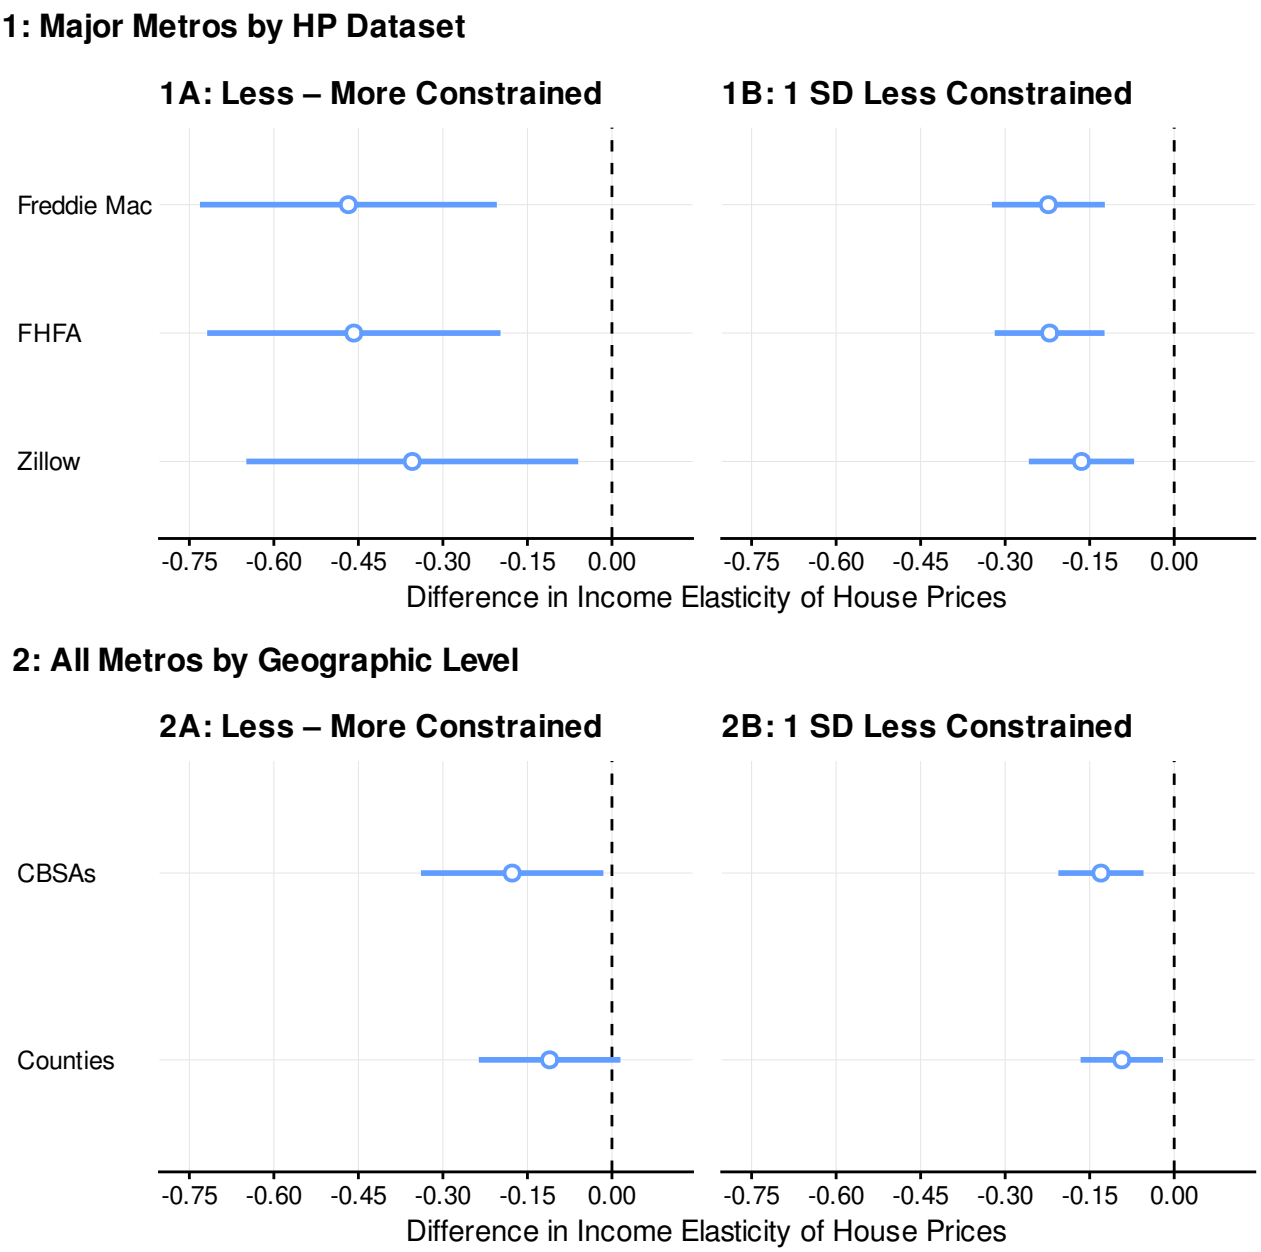
	Supply Constraint Measure:		
	BSH	Saiz	LU-ML
	(1)	(2)	(3)
Total Income Growth	0.779***	0.626***	0.863***
	(0.062)	(0.084)	(0.073)
LessConstrained	−0.011	−0.117***	0.051**
	(0.016)	(0.017)	(0.021)
Total Income Growth × *LessConstrained*	−0.062	0.097	−0.468***
	(0.106)	(0.112)	(0.134)
Number of CBSAs	314	251	380

Notes: This table tests whether looser supply constraints mitigate house price appreciation in response to income growth across major metropolitan CBSAs using the following long-differenced specification from 2000–2020:

$$\Delta \log(HP_i) = \alpha + \beta_1 \Delta \log(Income_i) + \beta_2 LessConstrained_i + \beta_3 [\Delta \log(Income_i) \times LessConstrained_i] + \epsilon_i$$

HP_i is the Freddie Mac house price index for CBSA i. *Income* is BEA total income. $LessConstrained_i$ is an indicator equal to 1 for CBSAs with looser-than-average supply constraints. We demean income growth across all U.S. CBSAs in the BEA dataset so that β_2 measures the secular price growth premium—the additional house price appreciation for less-constrained cities over the 20-year sample—evaluated at the average city-level income growth rate. Columns (1), (2), and (3) define the constraint indicator using the Baum-Snow and Han (2024, BSH) elasticity proxy, the Saiz (2010) measure, and our LU-ML Supply Index. The LU-ML Supply Index is constructed using the log change in CBSA-level housing units from 2000–2020 as the target variable. The regression intercept is not shown. All data are aggregated to 2020 CBSA delineations. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure B1: The Mitigation Effect (β_3) Using the LU-ML Supply Index, Alternative HP Datasets, Geographic Coverage, or County Delineations (2000–2020)



Notes: This figure tests whether looser supply constraints mitigate house price appreciation in response to income growth using a long-differenced specification from 2000–2020. Plotted estimates correspond to β_3 from the following specification:

$$\Delta \log(HP_i) = \alpha + \beta_1 \Delta \log(Income_i) + \beta_2 LU\text{-}ML_i + \beta_3[\Delta \log(Income_i) \times LU\text{-}ML_i] + \epsilon_i.$$

Income is the BEA total income growth at the CBSA level. In Panels 1A and 2A, $LU\text{-}ML_i$ is an indicator equal to 1 for CBSAs/counties with looser-than-average supply constraints. Panels 1B and 2B use the continuous version of the LU-ML Supply Index standardized to have zero mean and unit variance, such that positive values indicate looser constraints. Panel 1 shows the results for the major metropolitan areas across three house price datasets (Freddie Mac, FHFA, Zillow). Panel 2 shows the results at the CBSA and county levels using all U.S. CBSAs and the FHFA house price dataset. For CBSAs, the LU-ML Supply Index is constructed using the log change in CBSA-level housing units from 2000–2020 as the target variable. The county-level LU-ML Supply Index targets county-level housing units and employs a spatial cross-fitting design by training exclusively on counties outside of the target county's 2020 Commuting Zone to prevent contamination from local labor markets. Confidence intervals correspond to ±1.96 robust standard errors clustered at the CBSA level. Coefficient estimates are in Appendix Table B2.

Table B2: Estimated Income Elasticities of House Prices using the LU-ML Supply Index, 2000–2020

Geographic Level	LU-ML Specification	Sample	HP Dataset	Baseline Elasticity (β_1)	Secular Growth Premium (β_2)	Mitigation Effect (β_3)
CBSA	Binary	Major Metros	Freddie Mac	0.863*** (0.073)	0.051** (0.021)	-0.468*** (0.134)
CBSA	Binary	Major Metros	FHFA	0.834*** (0.070)	0.061*** (0.021)	-0.458*** (0.133)
CBSA	Binary	Major Metros	Zillow	0.649*** (0.103)	0.036 (0.030)	-0.354** (0.150)
CBSA	Binary	All CBSAs	FHFA	0.734*** (0.043)	0.055*** (0.010)	-0.177** (0.083)
CBSA	Continuous	Major Metros	Freddie Mac	0.697*** (0.052)	0.024** (0.010)	-0.223*** (0.051)
CBSA	Continuous	Major Metros	FHFA	0.675*** (0.051)	0.029*** (0.009)	-0.221*** (0.050)
CBSA	Continuous	Major Metros	Zillow	0.531*** (0.071)	0.014 (0.014)	-0.165*** (0.048)
CBSA	Continuous	All CBSAs	FHFA	0.665*** (0.032)	0.033*** (0.005)	-0.130*** (0.039)
County	Binary	All CBSAs	FHFA	0.707*** (0.048)	0.050*** (0.010)	-0.111* (0.064)
County	Continuous	All CBSAs	FHFA	0.638*** (0.035)	0.031*** (0.006)	-0.093** (0.037)

Notes: This table tests whether looser supply constraints mitigate house price appreciation in response to income growth using the LU-ML Supply Index. Estimates are from the following specification: $\Delta \log(HP_i) = \alpha + \beta_1 \Delta \log(Income_i) + \beta_2 LU\text{-}ML_i + \beta_3[\Delta \log(Income_i) \times LU\text{-}ML_i] + \epsilon_i$. $Income_i$ is BEA total income at the CBSA level. We demean income growth across all U.S. CBSAs in the BEA dataset so that β_2 measures the secular price growth premium: the increase in house price growth associated with looser constraints given average city-level income growth rate. For binary specifications, $LU\text{-}ML_i$ is an indicator equal to 1 for CBSAs/counties with looser-than-average supply constraints. For continuous specifications, the LU-ML Supply Index is standardized to have zero mean and unit variance. We report results for the major metropolitan areas across three house price datasets (Freddie Mac, FHFA, Zillow), as well as for all U.S. CBSAs and counties using the comprehensive FHFA house price dataset. For CBSAs, the LU-ML Supply Index is constructed using the log change in CBSA-level housing units from 2000–2020 as the target variable. The county-level LU-ML Supply Index targets county-level housing units and employs a spatial cross-fitting design by training exclusively on counties outside of the target county's 2020 Commuting Zone to prevent contamination from local labor markets. Robust standard errors are clustered at the CBSA level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B3: Estimated Income Elasticities from CBSA Decadal Panel Models, 1990–2020

Constraint Proxy	Specification	Fixed Effects	Baseline Elasticity (β_1)	Secular Growth Premium (β_2)	Mitigation Effect (β_3)
Panel A: LHS = Housing Unit Growth					
LU-ML Supply Index	Indicator	Time	0.328*** (0.018)	0.030*** (0.003)	0.111*** (0.027)
LU-ML Supply Index	Indicator	CBSA + Time	0.034** (0.015)	–	0.120*** (0.022)
LU-ML Supply Index	Continuous	Time	0.360*** (0.015)	0.018*** (0.001)	0.044*** (0.012)
LU-ML Supply Index	Continuous	CBSA + Time	0.084*** (0.014)	–	0.061*** (0.011)
Baum-Snow and Han	Indicator	Time	0.399*** (0.033)	0.006* (0.004)	0.081** (0.039)
Baum-Snow and Han	Indicator	CBSA + Time	0.049** (0.024)	–	0.082*** (0.031)
Baum-Snow and Han	Continuous	Time	0.437*** (0.024)	0.004** (0.002)	0.041 (0.027)
Baum-Snow and Han	Continuous	CBSA + Time	0.086*** (0.019)	–	0.022 (0.020)
Saiz	Indicator	Time	0.483*** (0.037)	-0.003 (0.004)	-0.013 (0.041)
Saiz	Indicator	CBSA + Time	0.116*** (0.042)	–	-0.040 (0.042)
Saiz	Continuous	Time	0.480*** (0.026)	-0.001 (0.002)	0.020 (0.022)
Saiz	Continuous	CBSA + Time	0.102*** (0.030)	–	-0.010 (0.019)
Panel B: LHS = House Price Growth					
LU-ML Supply Index	Indicator	Time	0.674*** (0.034)	0.013*** (0.005)	-0.136*** (0.044)
LU-ML Supply Index	Indicator	CBSA + Time	0.806*** (0.051)	–	-0.194*** (0.056)
LU-ML Supply Index	Continuous	Time	0.631*** (0.026)	0.005* (0.003)	-0.076*** (0.018)
LU-ML Supply Index	Continuous	CBSA + Time	0.736*** (0.038)	–	-0.126*** (0.025)
Baum-Snow and Han	Indicator	Time	0.688*** (0.051)	-0.010 (0.007)	-0.076 (0.070)
Baum-Snow and Han	Indicator	CBSA + Time	0.822*** (0.081)	–	-0.010 (0.091)
Baum-Snow and Han	Continuous	Time	0.653*** (0.040)	-0.007** (0.004)	-0.023 (0.038)
Baum-Snow and Han	Continuous	CBSA + Time	0.822*** (0.060)	–	0.020 (0.046)
Saiz	Indicator	Time	0.667*** (0.054)	-0.013* (0.008)	0.070 (0.075)
Saiz	Indicator	CBSA + Time	0.881*** (0.085)	–	0.222** (0.102)
Saiz	Continuous	Time	0.688*** (0.046)	-0.008** (0.004)	0.016 (0.032)
Saiz	Continuous	CBSA + Time	0.982*** (0.077)	–	0.142*** (0.052)

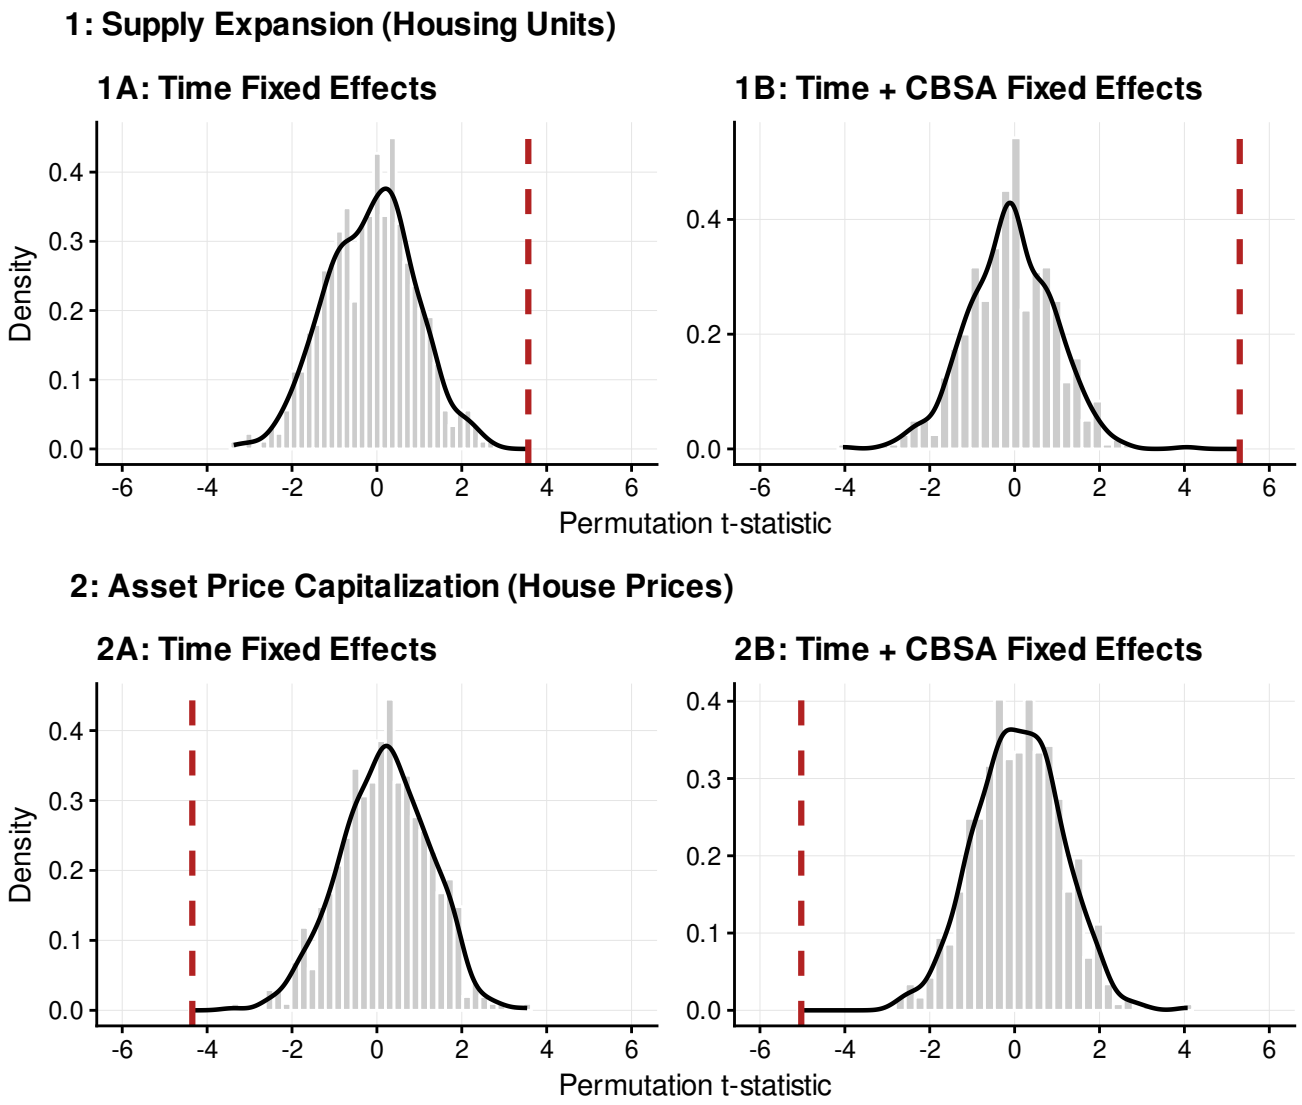
Notes: This table reports the results of CBSA decadal panel regressions evaluating the impact of supply constraints on the income elasticities of housing units (Panel A) and FHFA house prices (Panel B). Estimates are from the following specification: $\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 Constraint_i + \beta_3[\Delta \log(Income_{it}) \times Constraint_i] + \epsilon_{it}$, where α_i represents CBSA fixed effects (when applied) and δ_t represents time fixed effects. The constraint proxy is either the Baum-Snow and Han (2024) measure, the Saiz (2010) measure, or our LU-ML Supply Index. The LU-ML Supply Index is constructed using the log change in housing units from 1990–2020 (long-difference) as the target variable. $Income_{it}$ is total BEA income at the CBSA level. We demean income growth across all U.S. CBSAs in the pooled BEA dataset so that β_2 measures the secular growth premium: the increase in housing quantity or house price growth associated with looser constraints given average long-run city-level income growth. Because the constraint measures are time-invariant, the CBSA fixed effects (α_i) absorb the secular growth premium (β_2) when included. For binary specifications, the constraint is an indicator equal to 1 for CBSAs with looser-than-average supply constraints. The continuous versions are standardized to have zero mean and unit variance. Robust standard errors clustered at the CBSA level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B4: Robustness of LU-ML Supply Index: Alternative Targets and Thresholds

Index Construction	Specification	Fixed Effects	Baseline Elasticity (β_1)	Secular Growth Premium (β_2)	Mitigation Effect (β_3)
Panel A: LHS = Housing Unit Growth					
Main (1990–2020 Target)	Continuous	Time FE	0.360*** (0.015)	0.018*** (0.001)	0.044*** (0.012)
Main (1990–2020 Target)	Continuous	CBSA + Time FE	0.084*** (0.014)	–	0.061*** (0.011)
Main (1990–2020 Target)	Indicator	Time FE	0.328*** (0.018)	0.030*** (0.003)	0.111*** (0.027)
Main (1990–2020 Target)	Indicator	CBSA + Time FE	0.034** (0.015)	–	0.120*** (0.022)
1990–2000 Target	Continuous	Time FE	0.363*** (0.014)	0.017*** (0.002)	0.059*** (0.012)
1990–2000 Target	Continuous	CBSA + Time FE	0.081*** (0.014)	–	0.082*** (0.012)
1990–2000 Target	Indicator	Time FE	0.323*** (0.019)	0.031*** (0.003)	0.120*** (0.028)
1990–2000 Target	Indicator	CBSA + Time FE	0.014 (0.015)	–	0.159*** (0.021)
Winsorized	Continuous	Time FE	0.357*** (0.015)	0.018*** (0.001)	0.063*** (0.013)
Winsorized	Continuous	CBSA + Time FE	0.080*** (0.014)	–	0.081*** (0.012)
Median Split	Indicator	Time FE	0.313*** (0.019)	0.030*** (0.003)	0.129*** (0.027)
Median Split	Indicator	CBSA + Time FE	0.006 (0.012)	–	0.154*** (0.019)
Panel B: LHS = House Price Growth					
Main (1990–2020 Target)	Continuous	Time FE	0.631*** (0.026)	0.005* (0.003)	-0.076*** (0.018)
Main (1990–2020 Target)	Continuous	CBSA + Time FE	0.736*** (0.038)	–	-0.126*** (0.025)
Main (1990–2020 Target)	Indicator	Time FE	0.674*** (0.034)	0.013*** (0.005)	-0.136*** (0.044)
Main (1990–2020 Target)	Indicator	CBSA + Time FE	0.806*** (0.051)	–	-0.194*** (0.056)
1990–2000 Target	Continuous	Time FE	0.632*** (0.024)	0.002 (0.002)	-0.073*** (0.018)
1990–2000 Target	Continuous	CBSA + Time FE	0.725*** (0.036)	–	-0.113*** (0.026)
1990–2000 Target	Indicator	Time FE	0.659*** (0.034)	0.010** (0.005)	-0.104** (0.044)
1990–2000 Target	Indicator	CBSA + Time FE	0.803*** (0.050)	–	-0.192*** (0.056)
Winsorized	Continuous	Time FE	0.630*** (0.025)	0.003 (0.003)	-0.070*** (0.020)
Winsorized	Continuous	CBSA + Time FE	0.728*** (0.037)	–	-0.121*** (0.027)
Median Split	Indicator	Time FE	0.691*** (0.034)	0.004 (0.005)	-0.138*** (0.043)
Median Split	Indicator	CBSA + Time FE	0.834*** (0.053)	–	-0.229*** (0.057)

Notes: This table reports robustness checks for CBSA-level decadal panel regressions evaluating the impact of the LU-ML Supply Index on housing units (Panel A) and FHFA house prices (Panel B). Estimates are from the following specification: $\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 LU\text{-}ML_i + \beta_3 [\Delta \log(Income_{it}) \times LU\text{-}ML_i] + \epsilon_{it}$, where α_i represents CBSA fixed effects (when applied) and δ_t represents time fixed effects. To ensure our baseline results are not artifacts of the chosen target timeframe or extreme topological outliers, we vary the construction of the LU-ML index. We test the index trained on 1990–2000 unit growth as opposed to the baseline 1990–2020 target, a continuous index winsorized at the 5th and 95th percentiles, and a binary specification using the median rather than the mean. Continuous proxies are standardized to have zero mean and unit variance, while indicator variables equal 1 for looser-than-average supply constraints. Because these long-difference indices are time-invariant, the CBSA fixed effects (α_i) perfectly absorb their secular growth premiums (β_2) when included. Robust standard errors clustered at the CBSA level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure B2: Spatial Permutation Test of the LU-ML Supply Index Mitigation Effect (β_3)



Notes: This figure shows the empirical densities from a 500-iteration spatial permutation test evaluating the t-statistics of the supply mitigation effect (β_3) on housing units (Panel 1) and FHFA house prices (Panel 2). In each permutation, we first perform spatial block resampling by randomly shuffling CBSA geographic identifiers; this randomly reassigns all land unavailability features from one city to another. We then re-estimate the LU-ML Supply Index targeting the log change in housing units from 1990–2020 (long-difference) and re-run the following panel specification:

$$\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 LU\text{-}ML_i + \beta_3[\Delta \log(Income_{it}) \times LU\text{-}ML_i] + \epsilon_{it}$$

where α_i represents CBSA fixed effects (in Column B of each panel) and δ_t represents time fixed effects. The LU-ML Supply Index is standardized to have zero mean and unit variance, such that positive values indicate looser constraints. Y_{it} represents CBSA housing units or FHFA house prices. The gray bars correspond to the histogram of the t-statistics from the permutation tests and the black lines represent the empirical densities. The red dashed vertical lines represent the observed t-statistics from the unpermuted model. See also Appendix Table B5.

Table B5: Spatial Permutation Test of the LU-ML Supply Index Mitigation Effect (β_3)

Dependent Variable	Specification	Placebo 1% Bound	Placebo 99% Bound	Observed t-statistic	Empirical p-value
$\Delta \log$(Housing Units)	Time	-2.41	2.21	3.57	< 0.002
$\Delta \log$(House Prices)	Time	-2.31	2.39	-4.35	< 0.002
$\Delta \log$(Housing Units)	CBSA + Time	-2.53	2.06	5.30	< 0.002
$\Delta \log$(House Prices)	CBSA + Time	-2.35	2.31	-5.03	< 0.002

Notes: This table reports the empirical bounds and p-values from a 500-iteration spatial permutation test evaluating the t-statistics of the supply mitigation effect (β_3). In each permutation, we first perform spatial block resampling by randomly shuffling CBSA geographic identifiers; this randomly reassigns all land unavailability features from one city to another. We then re-estimate the LU-ML Supply Index targeting the log change in housing units from 1990–2020 (long-difference) and re-run the following panel specification:

$$\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 LU\text{-}ML_i + \beta_3[\Delta \log(Income_{it}) \times LU\text{-}ML_i] + \epsilon_{it}$$

where Y_{it} represents CBSA housing units or FHFA house prices, α_i represents CBSA fixed effects (bottom 2 rows) and δ_t signifies time fixed effects. The LU-ML Supply Index is standardized to have zero mean and unit variance, such that positive values indicate looser constraints. The placebo bounds represent the 1st and 99th percentiles of the resulting empirical null distribution of t-statistics for β_3. The empirical p-value is calculated as the proportion of permutations yielding an absolute t-statistic greater than or equal to the absolute t-statistic of the unpermuted model, bounded below by $1/501 \approx 0.002$. Across all specifications, the observed mitigation effect falls entirely outside the bounds of spatial and algorithmic noise.

Table B6: Income Elasticities using LU-ML Supply Index Quintiles

	Dependent variable:			
	Housing Unit Growth		House Price Growth	
	(1)	(2)	(3)	(4)
LU-ML Quintile 2	0.012^{***}		-0.013^{*}	
	(0.004)		(0.007)	
LU-ML Quintile 3	0.026^{***}		-0.022^{***}	
	(0.004)		(0.007)	
LU-ML Quintile 4	0.037^{***}		-0.009	
	(0.004)		(0.007)	
LU-ML Quintile 5	0.052^{***}		0.008	
	(0.004)		(0.008)	
Income Growth	0.254^{***}	-0.012	0.723^{***}	0.884^{***}
(Baseline: Quintile 1)	(0.022)	(0.016)	(0.055)	(0.074)
Income Growth ×	0.076^{**}	0.022	-0.082	-0.115
LU-ML Quintile 2	(0.036)	(0.021)	(0.074)	(0.104)
Income Growth ×	0.122^{***}	0.070^{**}	-0.096	-0.153
LU-ML Quintile 3	(0.034)	(0.030)	(0.077)	(0.099)
Income Growth ×	0.162^{***}	0.150^{***}	-0.083	-0.161^{*}
LU-ML Quintile 4	(0.030)	(0.024)	(0.073)	(0.089)
Income Growth ×	0.185^{***}	0.209^{***}	-0.231^{***}	-0.344^{***}
LU-ML Quintile 5	(0.041)	(0.033)	(0.068)	(0.086)
Time Fixed Effects	✓	✓	✓	✓
CBSA Fixed Effects		✓		✓

Notes: This table reports the results of CBSA decadal panel regressions evaluating the impact of quintiles in the LU-ML Supply Index on the income elasticities of housing units (Columns 1 and 2) and FHFA house prices (Columns 3 and 4). Estimates are from the following specification:

$$\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \sum_{q=2}^{5} \beta_{2,q} LU\text{-}ML_{iq} + \sum_{q=2}^{5} \beta_{3,q} [\Delta \log(Income_{it}) \times LU\text{-}ML_{iq}] + \epsilon_{it}$$

where α_i represents CBSA fixed effects (when applied) and δ_t represents time fixed effects. The LU-ML Supply Index is constructed using the log change in housing units from 1990–2020 (long-difference) as the target variable. $LU\text{-}ML_{iq}$ is an indicator equal to 1 if CBSA i falls into the q-th quintile of the LU-ML Supply Index, with Quintile 1 (most constrained) serving as the omitted baseline. $Income_{it}$ is total BEA income at the CBSA level. We demean income growth across all U.S. CBSAs in the pooled BEA dataset so that the main effects ($\beta_{2,q}$) measure the secular growth premium: the baseline difference in housing quantity or house price growth associated with looser constraints given average long-run national income growth. Because the LU-ML constraint quintiles are time-invariant, the CBSA fixed effects (α_i) absorb the secular growth premiums ($\beta_{2,q}$) in Columns (2) and (4). Robust standard errors clustered at the CBSA level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B7: Income Elasticities using Quintiles of the LU-ML Supply Index, Purged of Local Demand Factors

	Dependent variable:			
	Housing Unit Growth		House Price Growth	
	(1)	(2)	(3)	(4)
LU-ML Quintile 2	0.014***		−0.009	
	(0.005)		(0.008)	
LU-ML Quintile 3	0.015***		−0.007	
	(0.005)		(0.008)	
LU-ML Quintile 4	0.030***		0.009	
	(0.005)		(0.008)	
LU-ML Quintile 5	0.043***		0.005	
	(0.005)		(0.008)	
Income Growth	0.308***	0.020	0.631***	0.789***
(Baseline: Quintile 1)	(0.026)	(0.023)	(0.056)	(0.070)
Income Growth ×	0.063*	0.020	0.041	0.004
LU-ML Quintile 2	(0.037)	(0.031)	(0.077)	(0.105)
Income Growth ×	0.043	0.031	0.012	−0.036
LU-ML Quintile 3	(0.038)	(0.027)	(0.074)	(0.095)
Income Growth ×	0.120***	0.119***	−0.007	−0.069
LU-ML Quintile 4	(0.035)	(0.030)	(0.074)	(0.092)
Income Growth ×	0.155***	0.168***	−0.113*	−0.241***
LU-ML Quintile 5	(0.043)	(0.036)	(0.068)	(0.082)
Time Fixed Effects	✓	✓	✓	✓
CBSA Fixed Effects		✓		✓

Notes: See the notes to Appendix Table B6. Before binning the LU-ML Supply Index, we residualize it to local housing demand factors measured at the CBSA level. The variables are the Lutz (2025b) natural amenity demand index and the 1990 values of the following: college share, median income, housing density, employment share in manufacturing, and employment share in computer services. These local factors account for 7.9% of the variation in the LU-ML Supply Index (R^2). This robustness check ensures that the LU-ML Index is not simply capturing the effects of these local demand factors.

Table B8: Stability of the LU-ML Mitigation Effect (β_3) to Unobserved Spatial Demand

LHS Variable	Bound Scenario	$R^2_{D\sim Z\|X}$	$R^2_{Y\sim Z\|D,X}$	Adjusted Estimate	Robust SE	Adjusted t-stat
Housing Units	0x Unadjusted Estimate	–	–	0.075	(0.009)	8.19
	1x Natural Amenities (Gmaps) × Decade FE	0.022	0.022	0.068	(0.009)	7.46
	2x Natural Amenities (Gmaps) × Decade FE	0.045	0.044	0.061	(0.009)	6.72
	3x Natural Amenities (Gmaps) × Decade FE	0.067	0.066	0.054	(0.009)	5.95
	1x 1990 College Share × Decade FE	0.031	0.012	0.069	(0.009)	7.49
	2x 1990 College Share × Decade FE	0.062	0.023	0.063	(0.009)	6.78
	3x 1990 College Share × Decade FE	0.093	0.035	0.057	(0.009)	6.05
	1x 1990 Median Income × Decade FE	0.011	0.019	0.070	(0.009)	7.74
	2x 1990 Median Income × Decade FE	0.022	0.039	0.066	(0.009)	7.28
	3x 1990 Median Income × Decade FE	0.034	0.058	0.061	(0.009)	6.82
	1x 1990 Computer Services Emp Share × Decade FE	0.005	0.003	0.073	(0.009)	8.06
	2x 1990 Computer Services Emp Share × Decade FE	0.010	0.006	0.072	(0.009)	7.93
	3x 1990 Computer Services Emp Share × Decade FE	0.015	0.009	0.071	(0.009)	7.80
House Prices	0x Unadjusted Estimate	–	–	-0.113	(0.026)	-4.37
	1x Natural Amenities (Gmaps) × Decade FE	0.027	0.013	-0.097	(0.026)	-3.71
	2x Natural Amenities (Gmaps) × Decade FE	0.054	0.026	-0.080	(0.026)	-3.03
	3x Natural Amenities (Gmaps) × Decade FE	0.081	0.040	-0.062	(0.026)	-2.34
	1x 1990 College Share × Decade FE	0.044	0.018	-0.088	(0.026)	-3.35
	2x 1990 College Share × Decade FE	0.089	0.037	-0.061	(0.027)	-2.30
	3x 1990 College Share × Decade FE	0.133	0.055	-0.033	(0.027)	-1.23
	1x 1990 Median Income × Decade FE	0.006	0.006	-0.108	(0.026)	-4.18
	2x 1990 Median Income × Decade FE	0.012	0.011	-0.103	(0.026)	-3.99
	3x 1990 Median Income × Decade FE	0.018	0.017	-0.098	(0.026)	-3.79
	1x 1990 Computer Services Emp Share × Decade FE	0.007	0.006	-0.107	(0.026)	-4.14
	2x 1990 Computer Services Emp Share × Decade FE	0.015	0.013	-0.101	(0.026)	-3.91
	3x 1990 Computer Services Emp Share × Decade FE	0.022	0.019	-0.095	(0.026)	-3.68

Notes: This table reports the stability to unobserved spatial demand (Cinelli and Hazlett, 2020) of the change in the income elasticities of housing units and house prices associated with a 1 standard deviation loosening of supply constraints. Specifically, we evaluate the stability of the mitigation effect (β_3) from the following model: $\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_1 \Delta \log(Income_{it}) + \beta_2 LU\text{-}ML_{it} + \beta_3[\Delta \log(Income_{it}) \times LU\text{-}ML_{it}] + \mathbf{X}_i'\lambda_t + \epsilon_{it}$, where α_i represents CBSA fixed effects and δ_t represents decade fixed effects. We construct the LU-ML Supply Index using the decadal log change in housing units from 1990–2020 as the target variable. It is standardized over the pooled sample to have zero mean and unit variance; higher values indicate looser supply constraints. $Income_{it}$ is BEA total income. The vector $\mathbf{X}_i'\lambda_t$ represents baseline CBSA-level controls interacted with decade fixed effects. See Table 3 for the full regression results and list of controls. The unadjusted baseline reports β_3 assuming no unobserved confounding. The bounds simulate an unobserved spatial demand confounder that has 1, 2, or 3 times the explanatory power of the chosen benchmark covariates. Because the CBSA fixed effects absorb the base-decade (1990s) interaction for each benchmark, the benchmarks reflect the explanatory power of the time-varying slopes for the 2000s and 2010s. $R^2_{D\sim Z|X}$ and $R^2_{Y\sim Z|D,X}$ represent the hypothetical partial R^2 of the unobserved confounder with the treatment interaction term and the outcome, respectively. The housing unit and house price panels are estimated on their respective full samples; the house price sample is restricted to CBSAs with available FHFA house price data, leading to differing benchmark $R^2_{D\sim Z|X}$ values across the two panels. To account for spatial clustering by CBSAs, we scale the adjusted IID standard errors by the ratio of the baseline clustered standard error to the baseline IID standard error. This assumes that the intra-cluster correlation does not change with the addition of the omitted variable.

Table B9: Intra-City Income Elasticities Across Housing Supply Constraint Proxies

Geography	Constraint Proxy	Fixed Effects	Baseline Elasticity (β_1)	Secular Growth Premium (β_2)	Mitigation Effect (β_3)
Panel A: LHS = Housing Unit Growth					
County	LU-ML Supply Index	CBSA + Time FE	0.036** (0.016)	0.036*** (0.004)	0.059*** (0.011)
County	LU-ML Supply Index	County + Time FE	0.034** (0.017)	0.042*** (0.003)	0.061*** (0.010)
County	Baum-Snow and Han	CBSA + Time FE	0.075*** (0.022)	0.027*** (0.004)	0.061*** (0.020)
County	Baum-Snow and Han	County + Time FE	0.075*** (0.023)	–	0.059*** (0.018)
County	Saiz	CBSA + Time FE	0.063*** (0.019)	-0.015 (0.012)	-0.015 (0.013)
County	Saiz	County + Time FE	0.063*** (0.019)	–	-0.015 (0.012)
Zip5	LU-ML Supply Index	CBSA + Time FE	0.095* (0.052)	0.014*** (0.003)	0.078*** (0.018)
Zip5	LU-ML Supply Index	Zip5 + Time FE	0.094* (0.051)	0.014*** (0.003)	0.095*** (0.020)
Zip5	Baum-Snow and Han	CBSA + Time FE	0.108 (0.085)	0.046*** (0.005)	0.077*** (0.018)
Zip5	Baum-Snow and Han	Zip5 + Time FE	0.096 (0.081)	–	0.011 (0.030)
Panel B: LHS = House Price Growth					
County	LU-ML Supply Index	CBSA + Time FE	0.871*** (0.059)	0.009 (0.007)	-0.207*** (0.026)
County	LU-ML Supply Index	County + Time FE	0.872*** (0.059)	0.011 (0.008)	-0.211*** (0.027)
County	Baum-Snow and Han	CBSA + Time FE	0.911*** (0.080)	-0.009* (0.005)	0.010 (0.035)
County	Baum-Snow and Han	County + Time FE	0.914*** (0.079)	–	0.030 (0.039)
County	Saiz	CBSA + Time FE	0.954*** (0.070)	-0.011 (0.024)	0.175*** (0.048)
County	Saiz	County + Time FE	0.954*** (0.070)	–	0.175*** (0.048)
Zip5	LU-ML Supply Index	CBSA + Time FE	0.063** (0.030)	-0.005 (0.003)	-0.013 (0.013)
Zip5	LU-ML Supply Index	Zip5 + Time FE	0.065** (0.031)	-0.006 (0.004)	-0.026 (0.016)
Zip5	Baum-Snow and Han	CBSA + Time FE	0.035 (0.043)	-0.004* (0.002)	-0.020 (0.020)
Zip5	Baum-Snow and Han	Zip5 + Time FE	0.034 (0.045)	–	-0.021 (0.022)

Notes: This table reports the results of decadal panel regressions evaluating the impact of supply constraints on housing units (Panel A) and FHFA house prices (Panel B) within CBSAs across counties or zip codes. Estimates are from the following specification: $\Delta \log(Y_{i,c,t}) = \alpha_i + \gamma_c + \delta_t + \beta_1 \Delta \log(Income_{c,t}) + \beta_2 Constraint_{i,c,t} + \beta_3[\Delta \log(Income_{c,t}) \times Constraint_{i,c,t}] + \epsilon_{i,c,t}$. $Y_{i,c,t}$ is housing units or house prices in zip code or county i in CBSA c measured at the end of decade t. The specification includes either CBSA fixed effects (γ_c) or local fixed effects (α_i for county or zip code), along with decade fixed effects (δ_t), as indicated in the "Fixed Effects" column. $Income_{c,t}$ is BEA total income for CBSA c in decade t. We demean income growth across all CBSAs in the pooled BEA dataset so that β_2 measures the secular growth premium. $Constraint_{i,c,t}$, the constraint proxy for zip code or county i in CBSA c, is either the Baum-Snow and Han (2024, BSH) measure, the Saiz (2010) measure, or our time-varying LU-ML Supply Index. All constraint indices are standardized to zero mean and unit variance, and signed such that larger values indicate looser constraints. The county-level LU-ML Supply Index targets county-level housing units and employs a spatial cross-fitting design by training exclusively on counties outside of the target county's 2020 Commuting Zone to prevent contamination from local labor markets. The zip-code-level LU-ML Index similarly targets the decadal log change in zip-code housing units. Because the BSH and Saiz proxies are time-invariant, the county or zip code fixed effects (α_i) absorb their secular growth premiums (β_2) when applied. Robust standard errors clustered at the CBSA level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B10: Time-Varying Income Elasticities using the LU-ML Supply Index, 1990–2020

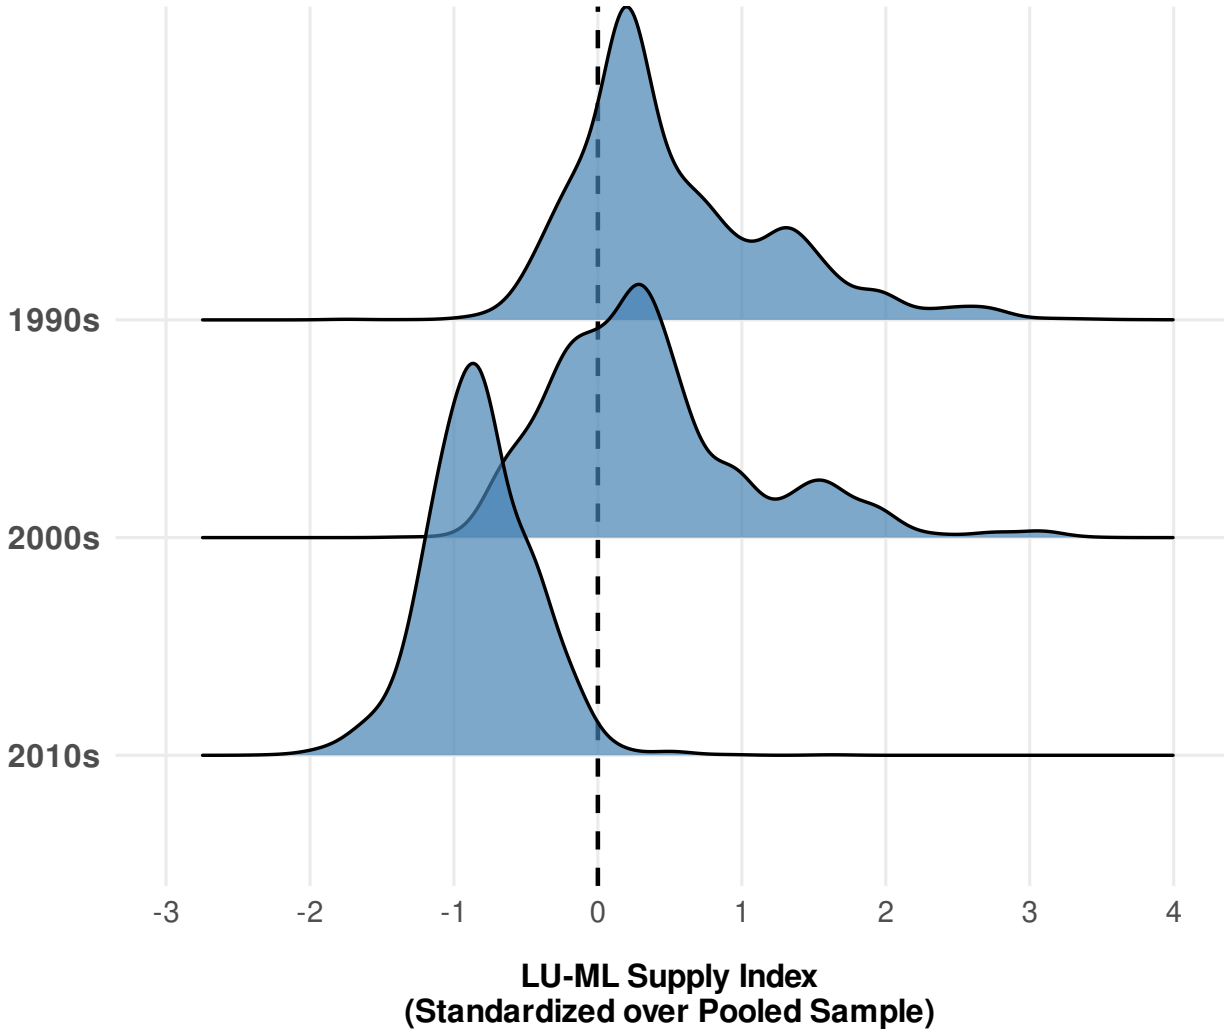
	Dependent variable:			
	$\Delta \log(\text{HU})$	$\Delta \log(\text{HU})$	$\Delta \log(\text{HP})$	$\Delta \log(\text{HP})$
	(1)	(2)	(3)	(4)
$\Delta \log(\text{Income})$	0.483^{***}	0.148^{***}	0.398^{***}	0.561^{***}
	(0.024)	(0.028)	(0.055)	(0.063)
LU-ML Supply Index	0.022^{***}	0.029^{***}	-0.017^{**}	-0.022^{**}
	(0.003)	(0.003)	(0.008)	(0.010)
$\Delta \log(\text{Income}) \times$ LU-ML	0.029	0.039^{**}	0.026	0.017
	(0.022)	(0.018)	(0.041)	(0.045)
$\Delta \log(\text{Income}) \times$ 2000s	-0.234^{***}	-0.086^{**}	0.303^{***}	0.240^{***}
	(0.034)	(0.043)	(0.066)	(0.078)
$\Delta \log(\text{Income}) \times$ 2010s	-0.273^{***}	-0.190^{***}	0.479^{***}	0.395^{**}
	(0.051)	(0.064)	(0.118)	(0.163)
LU-ML $\times$ 2000s	0.016^{***}	0.007^{*}	0.033^{***}	0.034^{***}
	(0.004)	(0.004)	(0.010)	(0.011)
LU-ML $\times$ 2010s	-0.001	0.006	0.056^{***}	0.057^{***}
	(0.007)	(0.010)	(0.015)	(0.022)
$\Delta \log(\text{Income}) \times$ LU-ML $\times$ 2000s	0.099^{***}	0.056^{*}	-0.324^{***}	-0.365^{***}
	(0.037)	(0.029)	(0.060)	(0.072)
$\Delta \log(\text{Income}) \times$ LU-ML $\times$ 2010s	-0.128^{**}	-0.057	0.016	0.004
	(0.054)	(0.067)	(0.108)	(0.156)
Time FE	✓	✓	✓	✓
CBSA FE		✓		✓

Notes: This table reports the results of CBSA decadal panel regressions evaluating the time-varying impact of the LU-ML Supply Index on the income elasticities of housing units (Columns 1 and 2) and FHFA house prices (Columns 3 and 4). Estimates are from the following specification:

$$\Delta \log(Y_{it}) = \alpha_i + \delta_t + \beta_{1t} \Delta \log(Income_{it}) + \beta_{2t} LU\text{-}ML_{it} + \beta_{3t} [\Delta \log(Income_{it}) \times LU\text{-}ML_{it}] + \epsilon_{it}$$

where α_i and δ_t represent CBSA and decade fixed effects, respectively. The time-varying parameters (β_{kt}) are estimated via interactions with decade indicator variables, where the baseline coefficients correspond to the omitted 1990s decade. The LU-ML Supply Index is standardized over the pooled sample (zero mean, unit variance); higher values indicate looser supply constraints. Because $Income_{it}$ is demeaned across all U.S. CBSAs, β_{2t} isolates the secular growth premium: the baseline increase in housing units or house price growth associated with looser constraints at average income growth. Robust standard errors clustered at the CBSA level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure B3: Densities of the Time-Varying LU-ML Supply Index by Decade, Weighted by 1990 CBSA Housing Units



Notes: See the notes to Figure 16.

C Appendix: Housing Collateral and Entrepreneurship

Table C1: The Housing Collateral Channel For Entrepreneurship (2SLS Estimates, 1994–2022)

	Dependent variable:	
	$\Delta \log(\text{Firms}^{1-4})$	
	(1)	(2)
$\Delta \log(HP)$	0.113^{***}	0.090^{**}
	(0.017)	(0.035)
First-stage F-stat	164.73	80.16
Number of Zip Codes	18,870	18,870
Number of Observations	460,539	460,539
Zip Code FE	✓	✓
Division-Year FE	✓	
State-Year FE		✓

Notes: This table shows the impact of house price growth on entrepreneurship, proxied by the growth in firms with 1–4 employees. Estimates are derived from the following two-stage least squares (2SLS) specification:

$$\Delta \log(\text{Firms}^{1-4}_{i,d,t}) = \alpha_i + \delta_{d,t} + \beta \cdot \Delta \log(\text{HP}_{i,d,t}) + \epsilon_{i,d,t}$$

where $\Delta \log(\text{Firms}^{1-4}_{i,d,t})$ is the year-over-year (YoY) log first-difference in the number of firms with 1–4 employees in zip code i and Census division d at year t. α_i represents zip code fixed effects, $\delta_{d,t}$ signifies Census division-time fixed effects, and $\Delta \log(\text{HP}_{i,d,t})$ is the YoY log difference in zip code FHFA house prices. In Column (2), we replace $\delta_{d,t}$ with state-time fixed effects $\delta_{s,t}$. We instrument house price growth with the zip-code-level LU-ML Price Pressure Index. The reported first-stage F-statistics are Kleibergen-Paap statistics. Robust standard errors clustered at the commuting zone level are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

D Appendix: Replications using the LU-ML Price Pressure Index

Table D1: First Stage House Price Regression as in Chaney et al. (2012)

	Dependent variable:		
	MSA HP Residential Prices		
	(1)	(2)	(3)
Saiz Elasticity × Natl Mortgage Rate	0.029***		
	(0.005)		
Baum-Snow & Han × Natl Mortgage Rate		0.023	
		(0.017)	
LU-ML Price Pressure			0.767***
			(0.039)
First Stage F-Stat	40.23	1.79	387.87
First Stage Partial R^2	0.23	0.01	0.33
Number of MSAs	97	322	368
Observations	1,358	4,499	5,132
MSA Fixed Effects	✓	✓	✓
Time Fixed Effects	✓	✓	✓

Notes: Column (1) replicates the first stage regression in Table 3, column (1) of Chaney et al. (2012) that uses Saiz Elasticity as an instrument for house prices using their equation (2): $P_t^l = \alpha^l + \delta_t + \gamma \cdot \mathit{Elasticity}^l \times IR_t + u_t^l$. P_t^l is the normalized residential house price index (in levels) for MSA l in year t, α^l and δ_t are MSA and time fixed effects, $\mathit{Elasticity}^l \times IR_t$ is the Saiz Elasticity proxy for each MSA multiplied by the national real mortgage rate, and u_t^l is the error term. Column (2) uses our LU-ML instrument for 2007 MSAs instead of $\mathit{Elasticity}^l \times IR_t$. The sample period ranges from 1993 to 2007. Robust standard errors clustered at the MSA level are in parentheses. One, two, or three asterisks represent statistical significance at the 10, 5, and 1 percent levels, respectively.

Table D2: Replication of Mian and Sufi (2014) – Non-Tradable Employment and the Housing Net Worth

	Non-Tradable Emp. Growth, 2007-09							
	Rest. & Retail	Geog. Concen.	Rest. & Retail	Geog. Concen.	Rest. & Retail	Geog. Concen.	Rest. & Retail	Geog. Concen.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Delta Housing Net Worth, 2006-09	0.174***	0.166***	0.374***	0.208**	0.137	−0.109	0.207***	0.173***
	(0.022)	(0.017)	(0.132)	(0.086)	(0.174)	(0.182)	(0.077)	(0.053)
First Stage F-Stat			11.07	11.07	3.21	3.21	107.48	107.48
First Stage Partial R^2			0.12	0.12	0.03	0.03	0.25	0.25
Specification	OLS	OLS	IV	IV	IV	IV	IV	IV
Instrument			Saiz Elasticity	Saiz Elasticity	Baum-Snow & Han	Baum-Snow & Han	LU-ML	LU-ML
Number of Counties	944	944	540	540	580	580	936	936

Notes: Columns 1 to 4 replicate Mian and Sufi (2014) using the equation $\Delta \log E_i^{NT} = \alpha + \eta \cdot \Delta HNW_i + \varepsilon_i$. $\Delta \log E_i^{NT}$ is the log change in non-tradable employment for county i and HNW_i is the change in housing net worth for county i. Mian and Sufi proxy non-tradable employment via the restaurant and retail sector (*Rest. and Retail*) or through industries that have low geographic concentration (*Geog. Concen.*). Columns (5) to (6) use the elasticity proxy from Baum-Snow and Han (2024) as an instrument, while columns (7) and (8) employ the LU-ML IV. Controls include 23 two-digit 2006 employment shares. Robust standard errors are clustered by state. One, two, or three asterisks represent statistical significance at the 10, 5, and 1 percent levels, respectively.

Table D3: First Stage House Price Regression as in Stroebel and Vavra (2019)

	House Price Growth					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 2001–06						
Saiz Elasticity	−0.069***	−0.058***				
	(0.012)	(0.011)				
Baum-Snow & Han			0.242***	0.152*		
			(0.084)	(0.083)		
LU-ML Price Pressure					1.027***	0.881***
					(0.064)	(0.066)
Number of MSAs	250	247	280	276	291	287
First Stage F-Stat	34.81	27.29	8.24	3.39	259.63	177.85
First Stage Partial R^2	0.23	0.20	0.03	0.02	0.44	0.35
Controls		✓		✓		✓
Panel B: 2007–11						
Saiz Elasticity	0.069***	0.038***				
	(0.012)	(0.007)				
Baum-Snow & Han			−0.080	−0.038		
			(0.104)	(0.076)		
LU-ML Price Pressure					1.007***	0.660***
					(0.093)	(0.094)
Number of MSAs	250	248	280	277	291	288
First Stage F-Stat	33.87	30.38	0.59	0.26	116.60	49.67
First Stage Partial R^2	0.19	0.08	0.00	0.00	0.34	0.21
Controls		✓		✓		✓

Notes: Columns (1) and (2) replicate the first stage regression in Table A2, columns (1) and (5) of Stroebel and Vavra (2019) using the equation $\Delta \log(HP_m) = \alpha + \beta Elasticity_m + \gamma X_m + \varepsilon_m$. $\Delta \log(HP_m)$ is the log difference in house prices for MSA m between 2001–06 (panel A) or 2007–11 (panel B). Stroebel and Vavra (2019) use proprietary house prices from CoreLogic. This replication uses publicly available Freddie Mac House Price Indices. $Elasticity_m$ is the Saiz elasticity proxy for MSA m. X_m is a vector of controls for MSA m and includes the change in the share of grocery retail employment, the change in the share of nontradable employment, the change in the share of construction employment, the change in the unemployment rate, and the change in the wage. See Table 1 in Stroebel and Vavra (2019) for more information on these controls. The number of MSAs in columns (1) and (2) is higher than in Table A2 of Stroebel and Vavra (2019) because their replication code includes an undisclosed outcome variable that limits the number of MSAs. Without this outcome variable, we could not determine which MSAs were used in their analysis. Here we report regressions results using all available MSAs in the Stroebel and Vavra (2019) data for each specification. Columns (3) and (4) use the elasticity proxy from Baum-Snow and Han (2024) and columns (5) and (6) use our LU-ML instrument. Robust standard errors are in parentheses. One, two, or three asterisks represent statistical significance at the 10, 5, and 1 percent levels, respectively.

Table D4: Replication of Guren et al. (2021), Table 1

	Dependent variable:		
	YoY Log Diff in Retail Emp Per Capita		
	(1)	(2)	(3)
Panel A: 1978–2017			
YoY Log Diff in HP Growth	0.083***	0.058***	0.084*
	(0.007)	(0.017)	(0.048)
First Stage F-Stat		249.08	19.67
First Stage Partial R^2		0.16	0.03
Observations	59,999	59,999	42,710
Panel B: 1990–2017			
YoY Log Diff in HP Growth	0.081***	0.072***	0.141***
	(0.008)	(0.015)	(0.037)
First Stage F-Stat		440.81	20.41
First Stage Partial R^2		0.27	0.04
Observations	41,985	41,985	29,867
Panel C: 2000–2017			
YoY Log Diff in HP Growth	0.068***	0.055***	0.134***
	(0.008)	(0.014)	(0.035)
First Stage F-Stat		351.11	21.12
First Stage Partial R^2		0.31	0.05
Observations	26,884	26,884	19,116
Specification	OLS	IV	IV
Instrument		Sensitivity	Saiz Elast
Num. CBSAs	380	380	270
Yr-Qtr FE			✓
Region, Yr-Qtr FE	✓	✓	
CBSA FE	✓	✓	✓

Notes: Replication of Guren et al. (2021), table 1 using the equation $\Delta y_{i,r,t} = \psi_i + \xi_{r,t} + \beta \Delta p_{i,r,t} + \Gamma X_{i,r,t} + \epsilon_{i,r,t}$. $\Delta y_{i,r,t}$ is the log annual change in quarterly retail employment per capita (a consumption proxy in year-over-year first-difference form) for CBSA i in census region r at time t. $\Delta p_{i,r,t}$ is the log annual change in quarterly house prices for CBSA i. ψ_i, $\xi_{r,t}$, and $X_{i,r,t}$ represent CBSA fixed effects, census region × time fixed effects, and other controls, such as industry shares, respectively. See Guren et al. (2021) for a full list of controls. Note that to ensure exact replication of Guren et al. (2021) Table 1, we execute the specifications precisely as implemented in their provided Stata code. Consequently, the Sensitivity instrument employs a sequential residualization to partial out CBSA fixed effects prior to estimating the region-time effects, while the Saiz elasticity specification is estimated using solely region-time fixed effects, mirroring the original authors' implementation. Robust standard errors clustered by time and CBSA are in parentheses. One, two, or three asterisks represent statistical significance at the 10, 5, and 1 percent levels, respectively.

Table D5: 2SLS Housing Wealth Elasticity Estimates as in Guren et al. (2021)

	YoY Log Diff in Retail Emp Per Capita				
	(1)	(2)	(3)	(4)	(5)
Panel A: 1978–2017					
YoY Log Diff in House Prices	0.085***	0.082	0.057***	0.289	0.074***
	(0.007)	(0.054)	(0.017)	(0.296)	(0.016)
First Stage F-Stat		14.37	235.16	1.12	286.13
First Stage Partial R^2		0.02	0.17	0.00	0.16
Panel B: 1990–2017					
YoY Log Diff in House Prices	0.081***	0.136***	0.072***	0.583	0.077***
	(0.008)	(0.040)	(0.015)	(0.703)	(0.016)
First Stage F-Stat		16.34	434.18	0.67	230.11
First Stage Partial R^2		0.04	0.28	0.00	0.21
Panel C: 2000–2017					
YoY Log Diff in House Prices	0.067***	0.121***	0.055***	0.427	0.065***
	(0.008)	(0.037)	(0.014)	(0.761)	(0.018)
First Stage F-Stat		16.26	347.48	0.32	166.93
First Stage Partial R^2		0.05	0.31	0.00	0.20
Specification	OLS	IV	IV	IV	IV
Instrument		Saiz	Sensitivity	Baum-Snow & Han	LU-ML
Num. CBSAs	376	270	376	311	376
CBSA FE	✓	✓	✓	✓	✓
Region × Date FE	✓	✓	✓	✓	✓

Notes: Columns (1) to (3) replicate Guren et al. (2021) using the equation $\Delta y_{i,r,t} = \psi_i + \xi_{r,t} + \beta \Delta p_{i,r,t} + \Gamma X_{i,r,t} + \epsilon_{i,r,t}$. $\Delta y_{i,r,t}$ is the log annual change in quarterly retail employment per capita (a consumption proxy in year-over-year first-difference form) for CBSA i in census region r at time t. $\Delta p_{i,r,t}$ is the log annual change in quarterly house prices for CBSA i. ψ_i, $\xi_{r,t}$, and $X_{i,r,t}$ represent CBSA fixed effects, census region × time fixed effects, and other controls, such as industry shares, respectively. See Guren et al. (2021) for a full list of controls. Column (1) employs OLS, while columns (2) and (3) use the Saiz Elasticity and Sensitivity instruments, respectively. Column (4) uses the elasticity proxy from Baum-Snow and Han (2024) and column (5) employs the LU-ML IV. Note that to ensure exact replication as in (Guren et al., 2021) Table 1, we execute the specifications precisely as implemented in their provided Stata code. Consequently, the Sensitivity instrument employs a sequential residualization to partial out CBSA fixed effects prior to estimating the region-time effects, while the Saiz elasticity specification (and our subsequent benchmarking using the BSH and LU-ML instruments) is estimated using solely region-time fixed effects, mirroring the original authors' implementation. Robust standard errors clustered by time and CBSA are in parentheses. One, two, or three asterisks represent statistical significance at the 10, 5, and 1 percent levels, respectively.