Hutchins Center Working Paper #115

September 2026

How Local Homeowners Insurance Market Stress Impacts Urban and Suburban Housing Markets

Matthew E. Kahn

University of Southern California

kahnme@usc.edu

Wesley A. Miller

University of Southern California

millerwe@usc.edu

Abstract

The seemingly constant wave of natural disasters and surging insurance premiums have fueled concerns of a homeowners insurance crisis. We construct a Local Insurance Strain Index to measure market stress based on homeowners insurance availability, affordability, and insurer profitability. Across U.S. ZIP codes, insurance strain rises with environmental risk. Within metropolitan areas, however, we document a striking divergence between this type of risk and homeowners insurance strain based on proximity to the central business district (CBD). Environmental risk increases with distance from the CBD, but insurance strain is concentrated in the urban core and declines along the spatial gradient. Urban homeowners face the lowest environmental risk but the highest insurance strain, driven by relatively unaffordable and less available coverage. This urban insurance penalty is not the product of insurer rent extraction, as private market underwriting profitability is similar in urban and suburban areas. While government supported backstops absorb substantial catastrophe risk in many suburban ZIP codes (e.g., wildfires in California and coastal flooding in Texas and Florida), urban policyholders receive no analogous peril-specific carve-out for non-environmental loss exposures that correlate with density (e.g., theft and burglary). Not only does local insurance strain negatively capitalize into housing markets, but we estimate the magnitude of capitalization is larger in the urban core. Thus, the spatial distribution of homeowners insurance strain directly impacts the relative cost of urban living and ultimately the location decisions of households.

Hutchins Center on Fiscal & Monetary Policy at BROOKINGS

THIS PAPER IS ONLINE AT

https://www.brookings.edu/articles/how-local-homeowners-insurance-market-stress-impacts-urban-and-suburban-housing-markets

1 Introduction

Under the backdrop of catastrophic natural disasters, homeowners insurance markets in the United States are widely perceived to be under pressure (Keys and Mulder, 2024). We quantify this pressure by constructing a Local Insurance Strain Index (LISI) that measures homeowners insurance market stress based on affordability, availability, and insurer profitability across ZIP codes from 2018 through 2022. High levels of local insurance strain are concentrated in areas that tend to face elevated environmental risks.

We explore the risk-strain relationship through the lens of urban economics by focusing on metropolitan areas and reveal a spatial pattern of environmental risk that rises with distance from the central business district (CBD). The likelihood of a natural disaster tends to be lowest where housing and employment density are highest, probably because of a combination of city selection, survivorship, and anthropogenic effects.[1]

While the urban-environmental risk gradient is an important finding, our primary contribution juxtaposes this spatial pattern against that for homeowners insurance strain. We show that homeowners insurance strain is highest in the urban core where environmental risk is lowest. Insurance strain declines as you move further from the city center, particularly within a five mile radius, until conditions flatten through a 30-mile range.

Decomposing our measure of homeowners insurance strain reveals two sources of the urban disadvantage. First, homeowners insurance is less available in urban ZIP codes as measured by nonrenewal rates. Insurers allow existing policies to expire at a rate roughly 20% higher in the urban core than in the same metropolitan area's suburban ZIP codes.[2] Second, homeowners insurance is less affordable near the CBD, where average premiums are 39% higher than in suburban ZIP codes.[3]

Surprisingly, we fail to detect meaningful differences in insurer profitability in urban and suburban ZIP codes, as the ratio of insured losses to premiums in the private market is effectively flat across the spatial gradient. Premiums spike in the city center because insured losses are highest in those same ZIP codes, suggesting that the urban insurance penalty is not the product of insurer rent extraction. This stability of insurer profitability characterizes the residual private market after catastrophe risks have been carved out to government-backed programs. Government-supported backstops, including the federal National Flood Insurance Program, state windstorm pools, and state FAIR Plans, absorb substantial catastrophe risk that concentrates in suburban ZIP codes (e.g., wildfires in California, windstorms in Texas and Florida).[4]

Standard homeowners insurance policies also cover non-environmental perils, and several density-correlated exposures plausibly drive up insurance costs even where environmental risk is lowest. Burglary and theft rates are typically higher in urban environments (Kneebone and Raphael, 2011). City centers tend to have older housing stock, which elevate the frequency of claims from plumbing and appliance failures. Higher labor and materials costs in urban areas combined with code-compliance requirements for repairs to older structures may also amplify the magnitude of losses. While suburban ZIP codes enjoy reduced insurance strain from government backstops of major environmental perils, policyholders receive no peril-

[1] If society considers physical risk when determining where to invest resources, we might expect economic activity to locate in relatively safe places. Locations that prove environmentally costly may eventually be abandoned or scaled down in favor of safer alternatives. Today's urban landscapes reflect the cumulation of long-run locational reoptimization. Furthermore, urbanization itself modifies the local physical environment. Infrastructure may reroute floodplains and waterways, modify ecological patterns that impact wildfire intensity, and alter microclimate conditions (e.g., urban heat islands).

[2] The mean nonrenewal rate among ZIP codes within one mile of their CBSA's CBD is 1.15% versus 0.94% for ZIP codes more than five miles out after CBSA$\times$year demeaning.

[3] We measure insurance unaffordability as the ratio of average homeowners insurance premiums per policy to the median household income for homeowners in a ZIP code.

[4] Some state FAIR Plans also cover non-catastrophe perils (e.g., vandalism), but these state-specific provisions are components of residual-market policies rather than comparably broad, peril-specific carve-outs for urban loss exposures.

specific carve-outs for the non-environmental loss exposures that elevate insurance stress in the urban core.

This asymmetry in public risk-bearing connects our within-metropolitan findings to the broader literature on cross-subsidization and pricing frictions in the U.S. homeowners insurance markets (Oh et al., 2026; Keys and Mulder, 2024). State-level rate regulation constrains insurers' ability to set risk-based premiums, deterring them from writing new policies, renewing existing policies, and potentially prompting them to exit markets where premiums cannot align with expected losses. Government-backed catastrophe programs (e.g., NFIP, state windstorm pools, FAIR Plans) absorb some catastrophe risks the private market has retreated from. Conditional on this institutional environment, the stability of private-market loss ratios across urban and suburban ZIP codes is consistent with competitive risk-based pricing *within* metropolitan areas.[5] Our results reveal this loss-ratio stability and highlight how the spatial equilibrium also depends on insurers' selective nonrenewal of high-risk policies. In particular, elevated nonrenewal rates near the CBD are consistent with binding pricing frictions *within* metropolitan areas that operate on non-environmental loss exposures. Oh et al. (2026) find similar adjustments on the extensive margin for insurers responding to losses in high-regulation states as well as insurer cross-state cross-subsidization via premium increases in low-regulation states.

Finally, we illustrate how local insurance strain has important consequences for housing markets. Unaffordable or unavailable homeowners insurance raises the carrying cost of single-family ownership and should depress home values. If single-family homes and apartments are substitutes, displaced demand should push multifamily rents upward. On the other hand, homeowners insurance strain is a function of location-specific disamenities that generate losses (e.g., natural hazards, theft, burglary, etc.). Disamenities reduce the desirability of a location for all types of housing, pulling down both single-family prices and multifamily rents.

We estimate that local insurance strain simultaneously decreases housing prices and rents. A one-standard-deviation increase in LISI is associated with a 10.6% decline in single-family home prices and a 3.1% decline in multifamily rents. Viewed through an asset-pricing framework, the negative rent response is consistent with a disamenity channel outweighing tenure-substitution effects. The larger price response and resulting 7.5% decline in the price-to-rent ratio are consistent with an additional ownership-specific insurance channel.[6] We also test for heterogeneity in capitalization across the spatial gradient and find disproportionate effects in the urban core. At the CBD, a one-standard-deviation increase in LISI is associated with a 19.1% decline in single-family home prices and a 6.2% decline in multifamily rents. Both effects decay smoothly to approximately zero by 30 miles out. These results support the urban-suburban distribution of insurance strain as a meaningful margin of spatial equilibrium that shapes the relative cost of living, local housing markets, and location decisions of households.

[5] In a frictionless competitive market with full pricing flexibility, insurers adjust premiums to match expected losses, equilibrating loss ratios across ZIP codes.

[6] A permanent shock that is fully and immediately reflected in current and expected future rents produces approximately equal percentage declines in prices and rents, whereas a transitory shock generally produces a larger current-rent response. An ownership-specific insurance-market shock instead depresses single-family values and can support multifamily rents through tenure substitution. Transitory insurance strain generally attenuates both effects. While standard price-to-rent models typically compare the price and rent of the same asset, our measure combines the Zillow Home Value Index for single-family homes with RentHub rents for purpose-built multifamily apartments in the same ZIP code. We interpret our measure as an approximation to the local price-to-rent ratio and a proxy for relative adjustment across housing segments.

2 Data and Empirical Strategy

2.1 Data

We construct a panel of homeowners insurance market outcomes, environmental risk, and housing market conditions at the ZIP code×year level over 2018–2022 by merging data from five sources. Our primary data source is the FIO Homeowners Insurance Data Collection, which provides ZIP code-level aggregates of the private homeowners insurance market under the authority of the Dodd-Frank Act.[7] We focus on three variables: (1) the nonrenewal rate (share of policies not renewed by insurers), (2) the paid loss ratio (losses paid divided by premiums earned), and (3) the average annual premium per policy. Our analytical FIO sample (excluding Texas, whose insurers did not provide nonrenewal data, and a small number of ZIP codes with the same gap) spans 70% of U.S. ZIP Code Tabulation Areas (ZCTAs) from 2018–2022, covering approximately 87% of the U.S. population.[8]

Household income at the ZIP code level comes from the American Community Survey (ACS) 5-year estimates of median household income for owner-occupied units, matched to ZIP codes via the HUD-USPS crosswalk.[9] Homeowner household income enters the LISI through the Unaffordability component, which is intended to capture the premium burden to homeowners' current cash flow rather than an asset-normalized carrying cost or a premium rate per dollar of dwelling coverage. Homeowners insurance is a recurring payment that must be met from current household resources, whereas home value is an illiquid asset and does not directly measure the resources available to meet annual premium payments.[10]

We combine these data to construct the Local Insurance Strain Index (LISI) that measures homeowners insurance market stress at the ZIP code×year level. The index aggregates information across three conceptual pillars: (1) *Unaffordability* (average premiums per policy as a share of median homeowner household income), (2) *Unavailability* (the nonrenewal rate), and (3) *Unprofitability* (insurers' paid loss ratio).We construct LISI so that a value of zero corresponds to the 2018 national median. Positive values of LISI indicate worse-than-baseline insurance market conditions, while negative values indicate the opposite.[11]

[7] Most policies in the FIO data are HO-3 "special form" contracts, which insure against a broad range of perils on an open-perils basis for the dwelling and a named-perils basis for personal property. The standard HO-3 policy excludes flood, earthquake, nuclear hazards, war, intentional acts, and ordinary wear and tear. Flood damage is covered separately by the federal National Flood Insurance Program (NFIP), with a small share of private flood insurance accounting for the remainder of the market. Wind and hurricane coverage is sometimes carved out of HO-3 policies in high-hurricane-risk states and placed in state windstorm pools. The FIO data reflect private homeowners coverage net of these carve-outs.

[8] The FIO data are subject to a confidentiality-driven suppression rule, i.e., ZIP codes with fewer than 10 reporting insurers or fewer than 50 policies are excluded. As a result, the thin rural markets are the most likely to be excluded. FIO data coverage is less than half of policies in Florida because that market relies disproportionately on smaller in-state insurers and the state's residual insurer, both largely outside the FIO's collection scope (Federal Insurance Office, 2025; Gourevitch and Kousky, 2025). We provide a comprehensive examination of FIO coverage, the reporting-gap exclusions, and selection in Appendix A.

[9] The ACS reports income at the ZCTA level. We prefer the median because it is less sensitive to upper-tail incomes and provides a clearer benchmark for the typical homeowner.

[10] Zhu et al. (2026) distinguish insurance cost burden from an effective insurance rate, while the Insurance Research Council (2025) similarly measures homeowners insurance affordability relative to median household income. HUD's broader housing-cost-burden framework also evaluates recurring owner costs relative to household income (Eggers and Moumen, 2008). Appendix C.1 constructs an otherwise identical asset-normalized index using the Zillow Home Value Index and shows that the urban gradient is not driven by the income denominator.

[11] We standardize each pillar against its 2018 cross-sectional median and interquartile range, and then we average the three standardized pillars with equal weights. We then recenter the composite index so that its 2018 cross-sectional median equals zero. Appendix B provides the complete mathematical formulation and summary statistics.

Environmental risk measures come from FEMA's National Risk Index (NRI), which provides standardized scores across 18 natural hazards at the Census tract level. We aggregate these risk metrics to ZIP codes using population-weighted averages.[12] For each hazard type, we use the annual frequency measure (AFREQ), which represents the expected number of events per year based on historical occurrence rates.[13]

Housing values come from the Zillow Home Value Index (ZHVI), which captures the typical home value for single-family residential properties in the middle tier (33rd–67th percentile). For multifamily rents, we use listing-level data from RentHub, computing the median rent in each ZIP-month.[14] We aggregate both monthly ZHVI and the monthly ZIP-level rent medians to annual means for our panel.

We assign each ZIP code to the nearest Core Based Statistical Area (CBSA) and compute the distance from the ZIP centroid to the CBSA's central business district (CBD).[15]

2.2 Sample Construction

Sample construction proceeds in stages. After excluding ZIP codes with nonrenewal reporting gaps, the insurance analysis panel is balanced at 23,806 ZIP codes in each of the five years, or 119,030 ZIP code×year observations. Requiring a published ACS median income for owner-occupied units yields the LISI construction panel of 116,610 ZIP code×year observations across 23,709 unique ZIP codes (2,420 observations drop out because they lack the income measure required for the Unaffordability pillar). Our main analysis concerns the within-metropolitan urban–suburban gradient, prompting us to restrict all main regression and gradient analyses to ZIP codes assigned to a CBSA and located within 30 miles of its CBD. This 30-mile analysis window covers approximately 66% of the U.S. population and 72% of the population in the states represented in the sample.[16] Within this window, the exact estimation sample varies by outcome. Home-price analyses require non-missing ZHVI data, while rent and price-to-rent analyses also require non-missing RentHub core multifamily rent data. Within the window, ZHVI is available for approximately 96% of observations, and the ZHVI–RentHub common sample spans approximately 50% and is disproportionately

[12] The NRI contains risk metrics for 18 natural hazard types: avalanche, coastal flooding, cold wave, drought, earthquake, hail, heat wave, hurricane, ice storm, landslide, lightning, riverine flooding, strong wind, tornado, tsunami, volcanic activity, wildfire, and winter weather. We construct a composite hazard measure by standardizing each hazard to z-scores and averaging across all 18 types. See https://hazards.fema.gov/nri/.

[13] We prefer AFREQ to NRI's expected annual loss measure because AFREQ is exogenous to local capital stock, population, and economic activity, whereas expected annual loss conflates hazard with exposure.

[14] We require a minimum of three listings per ZIP-month to reduce noise from thin markets. We focus on "core multifamily" listings (i.e., purpose-built rental apartments) which correlate strongly with Zillow's Observed Rent Index at the ZIP-year level ($r = 0.866$, mean difference $= \$32$/month) while providing substantially broader geographic coverage (RentHub covers 47.0% of our analysis-panel ZIP-years compared to 10.6% for ZORI; see Appendix A).

[15] We use the 2020 Census Bureau CBSA definitions and classify the CBD for each CBSA using a hybrid construction: for CBSAs in the contiguous 48 states, we identify the highest-employment-density Census tract using LEHD LODES8 Workplace Area Characteristics data; for CBSAs in Alaska (which lack LODES coverage), we use the highest-population-density tract from the 2020 Census as a fallback. The CBD is then defined as the centroid of the identified tract. We adopt this hybrid construction in place of single-source alternatives because LEHD employment density better identifies functional employment cores in major metros (where mixed retail and residential use can obscure population-density-based CBDs), while population density provides a defensible fallback in jurisdictions where LEHD data are unavailable. Distance is computed as the great-circle distance between the ZIP centroid and the CBD coordinate, in miles. Appendix D documents validation against published CBD lists and presents robustness to alternative CBD constructions.

[16] This restriction defines a common absolute-radius urban-suburban window rather than the full spatial extent of each CBSA. ZIP codes beyond 30 miles are typically rural or far-exurban areas with thin housing markets that represent a different empirical setting than the urban-suburban gradient we study. Appendix E documents cross-CBSA support through 30 miles, verifies that the gradient is unchanged in a common-support sample, and reports housing-market estimates under alternative distance caps.

urban. Appendix A reports the full sample-construction ladder.

2.3 Empirical Strategy

Our empirical analysis exploits the panel structure of our ZIP×year data. The key identification challenge is that insurance market outcomes correlate with unobserved metropolitan-level characteristics (e.g., regulatory stringency, market concentration, demographic composition) that also correlate with hazard exposure and housing market conditions. To address this, our preferred specification uses CBSA×year fixed effects:

$$Y_{z,t} = \alpha + \beta X_{z,t} + \theta_{m,t} + \varepsilon_{z,t} \tag{1}$$

where $Y_{z,t}$ is an outcome (e.g., ZHVI) for ZIP code z in year t, $X_{z,t}$ is an explanatory variable (e.g., LISI), and $\theta_{m,t}$ represents CBSA×year fixed effects for ZIP code z's assigned CBSA m. This specification absorbs all CBSA-level time-varying characteristics such as economic shocks, housing market trends, state-level regulatory changes, and metropolitan disaster events. Consequently, β is identified from cross-ZIP variation *within* the same CBSA and year. Standard errors are clustered at the CBSA×year level.

3 The Urban-Suburban Insurance Gradient

Across U.S. ZIP codes, homeowners insurance strain is positively associated with environmental risk. Table 1 establishes this relationship within metropolitan areas. A one-unit increase in annual hazard frequency is associated with a 0.16 increase in LISI, conditional on CBSA×year fixed effects (column 1). The pillar decomposition (columns 2–4) reveals that this risk-strain relationship operates through the household-facing margins of insurance market stress. Both Unaffordability and Unavailability are strongly positively associated with natural hazard exposure. Unprofitability, in contrast, is *negatively* associated with hazard within metropolitan areas. In other words, ZIP codes with higher environmental risk tend to have lower loss ratios compared to less risky neighborhoods in the same CBSA. Recall that the FIO private-market data exclude catastrophe-risk components that have been carved out to government-backed programs. If these carve-outs concentrate in the areas with the highest environmental risk (such as in the metropolitan periphery), then the residual private market's loss ratios will be mechanically supressed in high-hazard ZIP codes.

Table 1: LISI and Pillars on Composite Hazard Risk (CBSA×Year FE)

	(1) LISI Composite	(2) Unaffordability (Pillar)	(3) Unavailability (Pillar)	(4) Unprofitability (Pillar)
Composite AFREQ (z)	0.1642***	0.4234***	0.3323***	−0.2425***
	(0.0376)	(0.0577)	(0.0910)	(0.0330)
CBSA×Year FE	Yes	Yes	Yes	Yes
Zip Codes	13,595	13,595	13,616	13,616
Zip×Year Observations	67,190	67,190	68,080	68,080
R-squared	0.3513	0.4761	0.4719	0.1260
R-squared (within)	0.0014	0.0041	0.0022	0.0008

Notes: OLS regressions of LISI and its three pillars on composite hazard annual frequency (z-score), with CBSA×year fixed effects. Pillars are the standardized inputs to the composite LISI: Unaffordability (premium per policy / homeowner income), Unavailability (non-renewal rate), and Unprofitability (loss ratio). Each component is robust-standardized to the 2018 cross-section using (value − 2018 median) / 2018 IQR. Sample restricted to ZIP codes within 30 miles of the nearest CBSA's CBD. Standard errors clustered at the CBSA×year level in parentheses. Significance: *** p<0.01, ** p<0.05, * p<0.10.

While we expect environmental risk to predict insurance strain, our primary interest pertains to how these factors change as you move from the central business district to the periphery of the metropolitan area. Figure 1 plots the LISI (panel a) and our composite environmental risk measure (panel b) across the urban-suburban gradient.[17] The two panels reveal strikingly different patterns. Composite environmental risk is lowest in the urban core and increases with distance from the CBD through the metropolitan periphery. Tornados and wildfires are the largest contributors to the composite gradients, but coastal and riverine flood exposure also systematically increase with distance to the urban core. In contrast, homeowners insurance strain is concentrated near the CBD, falls sharply within a five mile radius, and then flattens through the suburban ring out to 30 miles. The opposite nature of these patterns along the spatial gradient suggests environmental risk is not driving stress in the private insurance market.

Figure 1: Insurance Strain Versus Physical Risk Across the Urban-Suburban Gradient

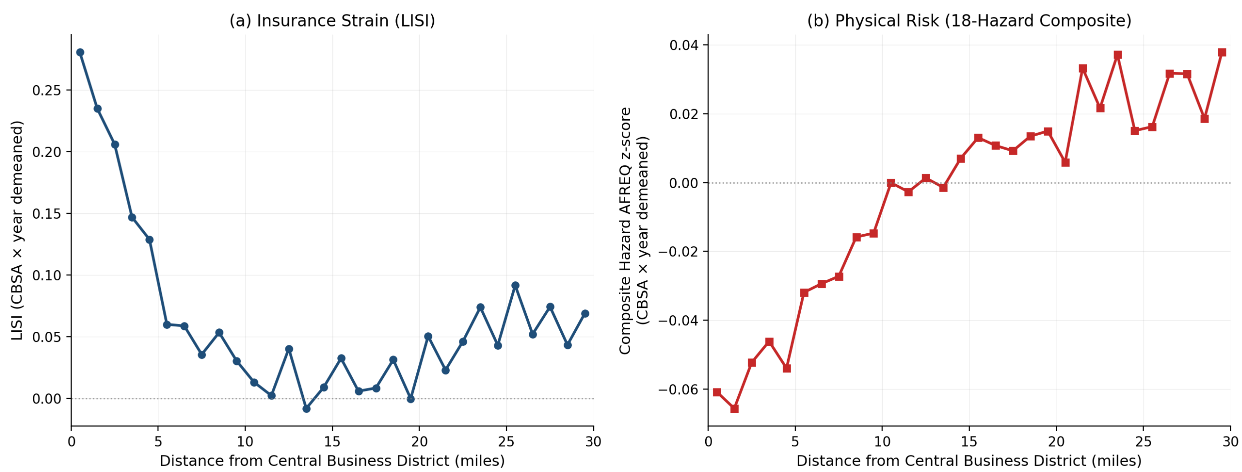

Notes: Panel (a) shows mean LISI by distance from CBD; panel (b) shows the composite physical risk z-score (standardized average of 18 NRI hazard annual frequency measures). Both variables are CBSA-demeaned: each ZIP code's value is adjusted as $\tilde{y}_z = (y_z - \bar{y}_m) + \bar{y}$, where $\bar{y}_m$ is the CBSA mean and $\bar{y}$ is the grand mean, so that gradients reflect within-CBSA spatial variation consistent with the CBSA × year fixed effects used in the regression analysis. Distance is binned in 1-mile intervals from the CBD; the x-axis is truncated at 30 miles. Sample restricted to CBSA×year cells with at least three ZIP codes in the 0–30 mile range. Based on 67,170 ZIP-year observations across 13,591 unique ZIP codes and 324 CBSAs, pooled over 2018–2022.

Figure 2 decomposes the composite LISI pattern into its three sub-indices and reveals two sources of the urban disadvantage. First, homeowners insurance is less available in urban ZIP codes. Insurers allow existing policies to expire at a rate roughly 20% higher in the urban core than in suburban ZIP codes of the same metropolitan area. Second, homeowners insurance is less affordable near the CBD, where average premiums are 39% higher than in suburban ZIP codes. In contrast, underwriting profitability is approximately flat across the urban-suburban gradient. The flat paid-loss ratio gradient is conditional on the set of policies that private insurers continue to write. Higher nonrenewal rates near the CBD can change the composition of the retained portfolio by removing risks for which regulatory-compliant premiums are insufficient. The observed loss-ratio stability, therefore, reflects both premium adjustment on retained policies and quantity adjustment through nonrenewal. The loss ratio pattern also masks substantial spatial co-movement of premiums and claim costs, since both claim frequency and average claim severity are likewise elevated at the urban core. In other words, average premiums per policy are higher in the urban core because that is where insurers pay more in losses.

The stability of private-market loss ratios across the urban-suburban gradient is consistent with risk-based pricing *within* metropolitan areas (subject to state regulation), suggesting that the elevated urban

[17] Appendix F.2 documents the spatial gradient of each individual hazard type. When aggregated, hazards that increase with distance dominate, producing the monotonic composite gradient shown in panel (b).

Figure 2: Insurance Strain Components Across the Urban-Suburban Gradient

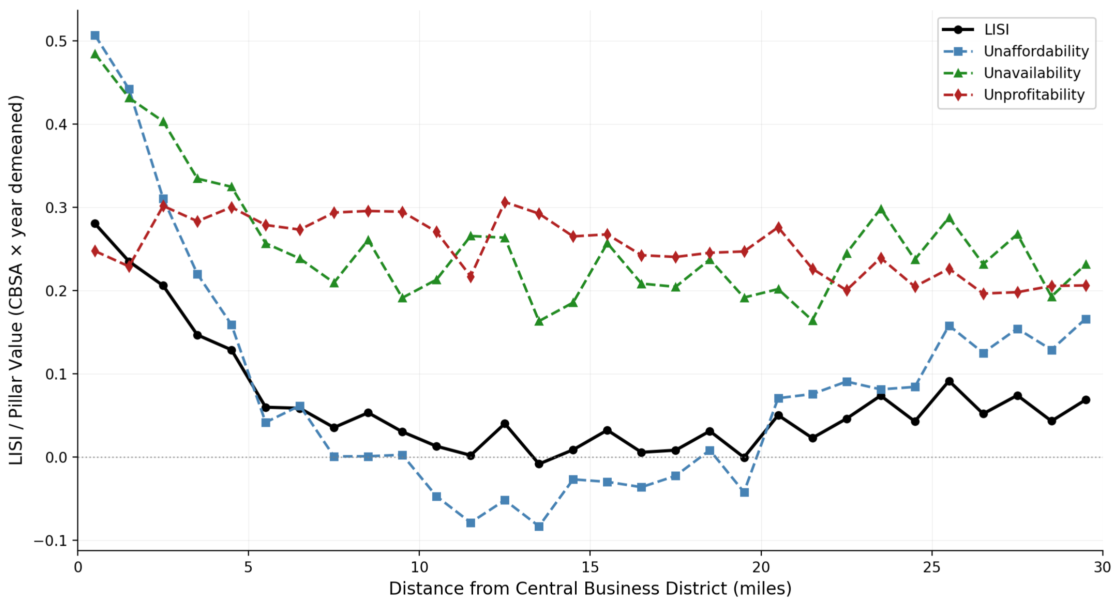

Notes: Mean LISI composite and three pillars (Unaffordability, Unavailability, Unprofitability) by distance from CBD. All four series are CBSA × year demeaned: each observation is adjusted as $\tilde{y}_{z,t} = (y_{z,t} - \bar{y}_{m,t}) + \bar{y}$, where $\bar{y}_{m,t}$ is the CBSA × year mean and $\bar{y}$ is the grand mean, so that gradients reflect within-CBSA spatial variation consistent with the CBSA × year fixed effects used in the regression analysis. Each pillar is robust-standardized to the 2018 cross-section using (value − 2018 median) / 2018 IQR; the LISI composite is the equal-weighted average of the three pillars and is recentered so the 2018 LISI median equals zero. Distance is binned in 1-mile intervals; the x-axis is truncated at 30 miles. Sample restricted to CBSA × year cells with at least three ZIP codes in the 0–30 mile range. Based on 67,170 ZIP-year observations across 13,591 unique ZIP codes and 324 CBSAs, pooled over 2018–2022.

strain is not the result of insurer rent extraction. These results characterize the private insurance market after catastrophe risks have been carved out to government-backed programs. The absence of these backstops would likely decrease affordability and availability in areas with the highest environmental risk, flattening the urban-to-suburban decline in LISI documented in Figure 2.

Government-supported insurance backstops, however, do not target the loss exposures that typically drive up premiums near the CBD. The fact remains that urban ZIP codes exhibit the highest insurance strain despite the lowest environmental risk. This pattern points to non-environmental loss exposures that correlate with density. Burglary and theft rates are typically higher in urban environments (Kneebone and Raphael, 2011). City centers tend to have older housing stock, which elevates the frequency of claims from plumbing and appliance failures. Higher labor and materials costs in urban areas, combined with code-compliance requirements for repairs to older structures, amplify the magnitude of losses. These density-correlated exposures likely explain the gap between urban homeowners insurance premiums and environmental risk. We turn next to whether the resulting spatial heterogeneity in insurance strain has measurable consequences for local housing markets.

4 Housing Market Capitalization

We examine three residential housing outcomes via regressions on LISI using our CBSA×year fixed effects specification (Equation 1). Table 2 presents the unconditional capitalization estimates.[18]

In the full sample (column 1), a one-unit increase in LISI is associated with a 7.6% decline in home prices, or a 5.5% decline per standard deviation of LISI.[19] A one-standard-deviation increase in LISI is roughly equivalent to moving a ZIP code from the median of the LISI distribution to approximately the 85th to 90th percentile in any given year, i.e., from a typical neighborhood to one experiencing meaningfully elevated insurance market stress. The resulting 5.5% decline in home prices is sizable in absolute terms, comparable in magnitude to capitalization estimates for other local disamenities documented in the urban economics literature (Linden and Rockoff, 2008). In the common sample of ZIP codes with both home-price and rent data (which is disproportionately urban) a one-unit increase in LISI is associated with a 15.4% decline in home prices (column 3) and a 4.5% decline in multifamily rents (column 2), or per-standard-deviation declines of 10.6% and 3.1%, respectively. The price-to-rent ratio declines by 10.9% per unit of LISI, or 7.5% per standard deviation (column 4). All effects are statistically significant at the 1% level.

Table 3 re-estimates Equation (1) on the LISI pillars themselves to explore the primary determinants of the LISI capitalization. Panel A enters each pillar on its own, so every row is a separate regression. The association is strongest for Unaffordability, but results are not solely driven by the income-normalized component. Unavailability and Unprofitability are negatively associated with all four outcomes. The smaller estimated effects for Unprofitability are expected given that the loss ratio is an insurer-side metric compared to the household-facing nature directly captured by Unaffordability and Unavailability.

Panel B enters all three pillars jointly, so each column is a single regression and each coefficient is a partial association holding the other two pillars fixed. Every coefficient retains its sign and remains statistically significant, and the joint estimates sit close to their separately estimated counterparts in Panel A.

The combination of results point toward two channels that link local insurance strain to housing markets. First, LISI may capture a homeownership-specific insurance-market wedge. At prevailing home prices,

[18] We use capitalization in a reduced-form equilibrium sense because local homeowner income, average premiums, and housing outcomes are jointly determined, so these regressions do not identify willingness to pay for an exogenous insurance-market shock.

[19] The within-sample standard deviation of LISI is 0.73 in the full sample and 0.69 in the common sample. Per-standard-deviation effects are computed by multiplying the regression coefficient on LISI by the corresponding within-sample standard deviation. For example, the column 1 coefficient of -0.076 implies a per-standard-deviation effect of $-0.076 \times 0.73 \approx -0.055$.

Table 2: Insurance Strain and Housing Market Outcomes

	(1) Log Home Prices	(2) Log Rents	(3) Log Home Prices	(4) Log Price-to-Rent
LISI	-0.076^{***}	-0.045^{***}	-0.154^{***}	-0.109^{***}
	(0.022)	(0.009)	(0.030)	(0.021)
CBSA×Year FE	Yes	Yes	Yes	Yes
Zip Codes	13,098	9,153	9,153	9,153
Zip×Year Observations	64,270	33,770	33,770	33,770
R-squared	0.6926	0.6614	0.6833	0.3920
R-squared (within)	0.0143	0.0063	0.0335	0.0229

Notes: OLS regressions of housing market outcomes on the Local Insurance Strain Index (LISI) with CBSA×year fixed effects. LISI is the equal-weighted average of three robust-standardized pillars: Unaffordability (premium / homeowner income), Unavailability (non-renewal rate), and Unprofitability (loss ratio). Column (1) uses the full sample of ZIP codes with non-missing LISI and Zillow Home Value Index (ZHVI). Columns (2)–(4) use a common sample further restricted to ZIP codes with non-missing RentHub core multifamily rents: Column (2) regresses log rents (RentHub core multifamily median), Column (3) regresses log home prices, and Column (4) regresses the log price-to-rent ratio (log ZHVI − log RentHub rent). Sample restricted to ZIP codes within 30 miles of the nearest CBSA's CBD; full-range robustness in Appendix E. Panel covers 2018–2022. Standard errors clustered at the CBSA×year level in parentheses. Significance: *** p<0.01, ** p<0.05, * p<0.10.

Table 3: Housing Market Outcomes by Insurance Strain Pillar

	(1) Log Home Prices (Full)	(2) Log Rents (Common)	(3) Log Home Prices (Common)	(4) Log Price-to-Rent (Common)
Panel A: Pillars entered separately				
Unaffordability (std.)	-0.042^{**}	-0.023^{***}	-0.086^{***}	-0.063^{***}
	(0.017)	(0.008)	(0.027)	(0.019)
Unavailability (std.)	-0.036^{***}	-0.012^{**}	-0.049^{***}	-0.036^{***}
	(0.007)	(0.005)	(0.015)	(0.010)
Unprofitability (std.)	-0.009^{**}	-0.013^{***}	-0.034^{***}	-0.020^{**}
	(0.005)	(0.004)	(0.012)	(0.008)
Panel B: Pillars entered jointly				
Unaffordability (std.)	-0.041^{**}	-0.023^{***}	-0.083^{***}	-0.060^{***}
	(0.017)	(0.008)	(0.027)	(0.019)
Unavailability (std.)	-0.034^{***}	-0.011^{**}	-0.043^{***}	-0.032^{***}
	(0.007)	(0.005)	(0.014)	(0.009)
Unprofitability (std.)	-0.010^{**}	-0.014^{***}	-0.035^{***}	-0.021^{***}
	(0.005)	(0.004)	(0.012)	(0.008)
R-squared (within)	0.0186	0.0069	0.0380	0.0267
CBSA×Year FE	Yes	Yes	Yes	Yes
Zip Codes	13,098	9,153	9,153	9,153
Zip×Year Observations	64,270	33,770	33,770	33,770

Notes: OLS regressions of housing market outcomes on the robust-standardized LISI pillars. In Panel A each pillar is entered on its own, so **every cell is a separate regression**. In Panel B all three pillars are entered together, so **each column is a single regression** and the coefficients are partial associations holding the other two pillars fixed. Panel B is reported in place of a composite index that omits one pillar: it identifies each pillar's independent contribution without requiring a re-normalized index. The specification, sample, and clustering match Table 2 exactly: Zip codes within 30 miles of the nearest CBSA's CBD, CBSA×year fixed effects, and standard errors clustered at the CBSA×year level. Column (1) uses the full sample of Zip codes with non-missing LISI and Zillow Home Value Index. Columns (2)–(4) use the common sample further restricted to Zip codes with non-missing RentHub core multifamily rents: column (2) regresses log rents, column (3) log home prices, and column (4) the log price-to-rent ratio (log ZHVI − log rent). Both panels are estimated on identical samples. Panel covers 2018–2022. Significance: *** p<0.01, ** p<0.05, * p<0.10.

higher premiums reduce households' willingness to pay for single-family homes. The resulting price decline is the capitalization response and can partially offset the higher recurring ownership cost for marginal buyers. Any displacement from ownership creates upward pressure on multifamily rental demand. Second, LISI may reflect location-specific disamenities that generate insured losses and reduce demand for both owner-occupied and rental housing.

The estimated negative multifamily-rent coefficient is consistent with the disamenity channel dominating upward substitution pressure from households displaced from ownership.[20] Under the standard benchmark, a permanent disamenity that is fully and immediately reflected in current and expected future housing-service flows would produce roughly equal percentage declines in prices and rents. A transitory disamenity shock would generally reduce current rents more than asset values. The decline in single-family prices, however, is larger than that for multifamily rents, consistent with a homeownership-specific insurance channel. We interpret the results as consistent with both mechanisms.[21]

We next examine whether capitalization varies along the urban-suburban gradient. In addition to adding controls for CBD distance and its square, Table 4 interacts CBD distance with LISI to allow the capitalization effect to vary along the spatial gradient. Results indicate that the capitalization of insurance strain is sharply concentrated near the urban core. At the CBD in the common sample, a one-unit increase in LISI is associated with a 27.7% decline in single-family home prices (column 3), an 8.9% decline in multifamily rents (column 2), and an 18.7% decline in the price-to-rent ratio (column 4).[22] All three effects decay smoothly with distance and approach zero by 30 miles out.[23] The spatial concentration of capitalization mirrors the spatial concentration of LISI itself. High-strain ZIP codes near the CBD experience both the largest LISI levels and the largest housing-market response per unit of LISI.

There are two primary implications drawn from these results. First, we confirm that local homeowners insurance strain is an important dimension of within-metropolitan housing prices and rents. Second, the spatial pattern of capitalization indicates that the welfare consequences of insurance strain are not uniformly distributed within metropolitan areas. They concentrate where strain is highest, compounding urban affordability burdens in an era of general housing unaffordability. The distribution of insurance strain is, therefore, a margin of urban spatial equilibrium that shapes the relative cost of urban living and, by extension, the location decisions of households.

[20] An additional consideration is that multifamily landlords face commercial property insurance rather than homeowners insurance. Commercial insurance costs may share local risk drivers with homeowners insurance, but they need not reprice one-for-one or pass fully into rents. Kim et al. (2026) find that commercial property insurance costs rose substantially faster than homeowners insurance costs across counties from 2019 to 2022, particularly where homeowners pricing faced greater regulatory frictions, and document incomplete rent pass-through.

[21] This interpretation does not require LISI itself to be permanent. LISI is a contemporaneous composite of Unaffordability, Unavailability, and Unprofitability, each of which can fluctuate and can simultaneously reflect persistent local loss conditions and transitory insurance-market frictions.

[22] In per-standard-deviation terms, these effects are 19.1%, 6.2%, and 12.9%, respectively. Per-standard-deviation effects in the interaction model are computed by multiplying the implied marginal effect of LISI at distance d by the within-sample standard deviation of LISI (0.69 in the common sample and 0.73 in the full sample). The implied marginal effect at distance d miles is $\hat{a} + \hat{c} \cdot d$, where $\hat{a}$ is the LISI main-effect coefficient and $\hat{c}$ is the LISI$\times$CBD interaction coefficient (per mile of CBD distance). At the CBD ($d = 0$), the implied marginal effect collapses to the main-effect coefficient $\hat{a}$; for example, the column 3 coefficient $\hat{a} = -0.277$ implies a per-standard-deviation effect at the CBD of $-0.277 \times 0.69 \approx -0.191$.

[23] The standalone distance terms describe the within-metropolitan housing profile at LISI=0, which corresponds to the index's 2018 national-median normalization point. At this reference level of insurance strain, the positive linear and negative quadratic coefficients trace a concave gradient in which single-family values rise with distance from the CBD, peak near 16 miles, and decline thereafter. Because distance is also interacted with LISI, the total distance gradient varies with the level of insurance strain. We report the distance coefficients for transparency but do not treat them as primary objects of interest.

Table 4: Housing Market Capitalization of LISI Across the Urban-Suburban Gradient (CBSA×Year FE, 30-mile cap)

	(1) Log Home Prices (Full)	(2) Log Rents	(3) Log Home Prices (Common)	(4) Log Price-to-Rent
LISI	-0.1694^{***}	-0.0893^{***}	-0.2765^{***}	-0.1872^{***}
	(0.0395)	(0.0131)	(0.0396)	(0.0285)
LISI × CBD Distance (per mi)	0.00709^{***}	0.00329^{***}	0.01003^{***}	0.00675^{***}
	(0.00138)	(0.00073)	(0.00191)	(0.00134)
CBD Distance (per mi)	0.02611^{***}	-0.00057	0.01900^{***}	0.01957^{***}
	(0.00314)	(0.00160)	(0.00377)	(0.00261)
CBD Distance^2 (per mi^2)	-0.000827^{***}	-0.000016	-0.000568^{***}	-0.000552^{***}
	(0.000085)	(0.000056)	(0.000104)	(0.000068)
CBSA×Year FE	Yes	Yes	Yes	Yes
Zip Codes	13,098	9,153	9,153	9,153
Zip×Year Observations	64,270	33,770	33,770	33,770
R-squared	0.7016	0.6627	0.6926	0.4097
R-squared (within)	0.0434	0.0103	0.0621	0.0514

Notes: OLS regressions of housing market outcomes on LISI, its interaction with distance to the CBD (per mile), and a standalone CBD-distance main effect and its quadratic, with CBSA×year fixed effects. Conditioning on the distance profile is what allows the LISI main effect to be read as the capitalization rate at the CBD rather than as a term absorbing any within-CBSA spatial gradient in prices unrelated to insurance strain. The implied marginal effect of LISI at distance d miles is $\hat{a}+\hat{c}\cdot d$, where $\hat{a}$ is the LISI main-effect coefficient and $\hat{c}$ is the LISI×CBD coefficient; at the CBD ($d = 0$) it collapses to $\hat{a}$. Column (1) uses the full sample of Zip codes with non-missing LISI and ZHVI; Columns (2)–(4) use the common sample further restricted to Zip codes with non-missing RentHub core multifamily rents. Sample restricted to Zip codes within 30 miles of the nearest CBSA's CBD. Standard errors clustered at the CBSA×year level in parentheses. Significance: *** p<0.01, ** p<0.05, * p<0.10.

5 Discussion

The within-metro pattern of insurance strain documented above raises broader questions about how risk is shared across U.S. homeowners insurance markets, and what the urban concentration of strain implies for policy and future research.

5.1 Pricing, Availability, and Public-Risk Bearing

Our results contribute to our understanding of cross-subsidization in U.S. homeowners insurance markets. Cross-subsidization operates on at least three distinct margins. The first margin is *between states*. Oh et al. (2026) document that rate regulation in highly-regulated states forces insurers to under-price catastrophe risk, and insurers compensate by raising rates in less-regulated states. The cross-subsidy flows from policyholders in lightly-regulated states to policyholders in heavily-regulated, catastrophe-exposed states.

The second margin of cross-subsidization is *within states* through residual markets and government-backed programs. The federal National Flood Insurance Program covers a substantial share of flood-related catastrophe exposure outside the private homeowners insurance market entirely, funded through its own policyholder premiums and Treasury backstops. State windstorm pools (e.g., Florida Citizens, the Texas Windstorm Insurance Association) and state FAIR Plans cover high-risk wind, hurricane, and wildfire exposures and are typically funded through a layered structure of pool premiums, reserves, reinsurance, and assessments on private insurers. Where assessments are triggered, insurers recover the costs through surcharges or future rate filings, broadly raising private homeowners premiums across the state. The geographic incidence of these assessment costs depends on state regulatory constraints on territorial rating.

The FIO data, therefore, characterize the residual private market *after* these catastrophe carve-outs. The total burden of insuring catastrophe-prone suburban and exurban properties (including the implicit cross-subsidy through state-level assessments) is larger than the private-market data alone can reveal.

Our analysis is particularly informative on the third margin of homeowners insurance cross-subsidization, i.e., *within metropolitan areas* across the urban-suburban gradient. Conditional on the regulatory environment and the catastrophe carve-outs that operate at the state and federal levels, we find that private-market loss ratios are approximately flat across the urban-suburban gradient. This spatial stability indicates that insurers typically price ZIP codes in proportional measure to its private-market realized losses. We interpret this as an absence of evidence of private-market cross-subsidization within metropolitan areas.[24]

Insurers determine their portfolio on both an intensive and an extensive margin. Not only can they adjust premiums (conditional on their state's regulation), but they can decline to underwrite or renew policies. Price adjustment and nonrenewal are, therefore, alternative responses at the individual-policy level, and both contribute to the observed equilibrium. Selective nonrenewal can preserve a stable loss ratio by removing risks that cannot be repriced sufficiently. The combination of elevated urban nonrenewal rates and loss ratios comparable to those in the suburban levels is consistent with this mechanism. The relationship between loss ratios and nonrenewal also illustrates why Unavailability and Unprofitability enter the LISI as separate dimensions of market stress. For example, a market with stable loss ratios may still suffer from strain if insurers restrict coverage on the extensive margin.

Government policy affects both margins. Rate regulation shapes the extent and speed of premium adjustment. The NFIP, state windstorm pools, and FAIR Plans alter which risks remain on private balance sheets. Building codes and mitigation policies change expected losses, while zoning and construction regulation change the location and quantity of future exposure. Holding pricing constraints fixed, the absence of catastrophe backstops would tend to shift additional loss exposure toward households and private carriers in hazard-prone locations, increasing premiums, nonrenewals, or both. Conversely, greater pricing flexibility would tend to shift adjustment away from insurer-initiated nonrenewal and toward premiums.

This policy-mix perspective is related to Ostriker and Russo (2026), who study a floodplain regime in which insurance pricing and purchase requirements operate alongside floodplain designation and elevation standards. They show that non-price regulation affects both the location and flood resilience of new construction. The common lesson is that insurance prices should not be interpreted separately from coverage availability, public risk-bearing, and policies that alter the underlying stock and location of insured risk.

Most regulatory and policy attention in homeowners insurance has focused on catastrophe risk in environmentally exposed regions such as coastal hurricane zones, wildland-urban interfaces, and other geographically defined catastrophe corridors. Our results suggest this framing is incomplete. A substantial share of homeowners losses comes from non-environmental perils rather than from natural hazards, and these non-environmental loss exposures are systematically concentrated in urban cores. For policy purposes, this source of insurance strain is not addressable through the natural-hazard-pricing reforms that have dominated regulatory debate. Where pricing flexibility is constrained, insurers shift toward the availability margin and decline to renew policies they cannot reprice sufficiently. Elevated nonrenewal rates in the urban core are consistent with this mechanism operating on the dimension of non-environmental loss exposures. The homeowners insurance burden, however, is not symmetric across the urban-suburban gradient. Catastrophe-exposed suburbs are more likely to benefit from government-backed programs that absorb the highest-risk components of their hazard exposure. The urban core has no analogous peril-specific carve-out for its non-environmental loss exposures, leaving urban policyholders to absorb the resulting cost

[24] This does not establish actuarial adequacy for every policy, rule out cross-subsidization *within* ZIP codes, nor measure total economic profits.

more directly through higher premiums and elevated nonrenewal rates.

Consequently, efforts to improve urban housing affordability should account for homeowners insurance affordability *and* availability when evaluating the policy approaches typically applied to insurance and housing markets, (e.g., rate regulation, residual-market expansion, and zoning). If the elevated cost of insuring urban property is driven primarily by non-environmental loss exposures, then reductions in the underlying loss-generating factors (e.g., crime, structural fire risk, plumbing failures, liability exposure) would reduce the insurance burden at its source rather than redistributing it through cross-subsidies or rationing it through nonrenewal.

5.2 Future Research

The within-metro pattern of local homeowners insurance strain raises empirical questions that fall outside the scope of this analysis but are central to its policy implications. If non-environmental loss exposures drive the urban concentration of insurance strain, then improvements in urban conditions should meaningfully reduce insurance burden. Reductions in crime would reduce theft and vandalism claim activity; building-code modernization and infrastructure investment would reduce water-damage and fire-related claim activity; improvements to school quality, lighting, sidewalks, and other components of the urban environment could reduce liability exposure and, more broadly, the disamenities that depress housing demand in high-strain neighborhoods. These improvements are not typically discussed as insurance-market policies, but our results suggest they could be more effective than rate regulation at addressing the spatial concentration of insurance strain in urban cores.

Causal identification of these effects requires policy variation that we do not exploit here. Future research could leverage municipal interventions in crime reduction, building-code enforcement, school quality, or infrastructure investment to estimate the causal effect of urban improvements on homeowners insurance market outcomes. The results would have implications for both urban policy and insurance regulation. The geographic incidence of insurance burden is not a purely insurance-market phenomenon, and policies that improve the underlying urban environment may produce insurance-market benefits as a byproduct.

A related priority is to measure the spatial relationship between homeowners and commercial property insurance markets. Multifamily buildings face many of the same local hazards and repair-cost conditions as single-family homes, but commercial policies are generally less tightly constrained by state rate regulation, individually underwritten, and subject to different coverage and lender requirements. Kim et al. (2026) find that commercial multifamily insurance costs have risen substantially faster than homeowners premiums across counties, particularly in high-risk locations where homeowners pricing is more constrained, and that part of these commercial cost increases is reflected in rents. Their comparison nevertheless captures a reduced-form pricing wedge across different contracts rather than matched premium rates. We do not observe multifamily commercial premiums, coverage limits, or deductibles and therefore cannot determine whether commercial and homeowners insurance exhibit the same within-metropolitan urban gradient. Matched policy data across the two lines would allow future research to separate common local loss exposure from differences in regulation, underwriting, coverage, and cost pass-through.

6 Conclusion

Public discussion of the homeowners insurance market is dominated by concerns about climate risk and the affordability of coverage in catastrophe-prone regions. We construct the Local Insurance Strain Index (LISI) to measure homeowners insurance market stress at the ZIP code level over 2018–2022, combining three dimensions—unaffordability, unavailability, and unprofitability—into a single composite index. Applied across U.S. ZIP codes, LISI rises with environmental risk, consistent with the conventional climate-driven view of insurance market stress.

Within metropolitan areas, however, the relationship inverts. Environmental risk increases with distance from the central business district, but homeowners insurance strain is concentrated in the urban core and declines along the spatial gradient before flattening through the suburban ring. Urban homeowners face the lowest environmental risk but the highest insurance strain, driven by less affordable and less available coverage. This urban burden is not a product of insurer rent extraction. Paid loss ratios in the private market are approximately flat across the urban-suburban gradient, indicating that premiums move roughly in proportion to realized paid losses (conditional on insurers' renewal decisions and government-supported catastrophe carve-outs). The elevated cost of insuring urban properties likely reflects non-environmental loss exposures (e.g., theft, non-flood water damage, liability) that correlate with population density, the age and complexity of urban housing stock, and the cost of repair labor in dense markets.

These spatial differences in homeowners insurance burden capitalize into housing markets. A one-standard-deviation increase in LISI is associated with a 10.6% decline in single-family home prices and a 3.1% decline in multifamily rents, with the price-to-rent ratio declining by 7.5% per standard deviation. The capitalization is sharply concentrated in the urban core. At the CBD, the corresponding effects are 19.1% for home prices and 6.2% for multifamily rents, decaying to approximately zero by 30 miles out. The negative association of LISI with both prices and rents is consistent with a disamenity channel that dominates the upward pressure on multifamily rents from tenure substitution. At the same time, the substantially larger single-family price decline and the resulting fall in the price-to-rent ratio are consistent with an additional ownership-specific insurance channel.

Our findings reframe the homeowners insurance crisis as a within-metro phenomenon as much as a cross-state climate-risk phenomenon. Regulatory and policy attention has focused on catastrophe risk in environmentally exposed regions, but a substantial share of homeowners insurance losses comes from non-environmental perils that concentrate in urban cores. Efforts aimed at reducing urban housing affordability burdens should account for insurance affordability and availability when considering familiar policy approaches (e.g., rate regulation, residual-market expansion, and zoning). The cost of insuring urban property is shaped by local conditions (e.g., crime, building stock, infrastructure, liability environment) that often lie outside the conventional scope of insurance policy. The spatial distribution of insurance burden is an understudied margin of urban spatial equilibrium, with consequences for housing affordability and the location decisions that are distinct from the cross-state climate-risk dynamics that have dominated prior research.

References

Davis, Morris A., William D. Larson, Stephen D. Oliner, and Jessica Shui, "The Price of Residential Land for Counties, ZIP Codes, and Census Tracts in the United States," *Journal of Monetary Economics*, 2021, *118*, 413–431.

Eggers, Frederick J. and Fouad Moumen, "Trends in Housing Costs: 1985–2005 and the 30-Percent-of-Income Standard," Technical Report, U.S. Department of Housing and Urban Development, Office of Policy Development and Research, Washington, DC June 2008.

Federal Insurance Office, "Analyses of U.S. Homeowners Insurance Markets, 2018–2022: Climate-Related Risks and Other Factors," Technical Report, U.S. Department of the Treasury, Washington, DC January 2025. Released January 16, 2025.

Gourevitch, Jesse D. and Carolyn Kousky, "New Homeowners Insurance Data Reveals Insights into Market Trends and Suggests Future Research Needs," *Risk Management and Insurance Review*, 2025, *28* (2), 324–338.

Insurance Research Council, "Homeowners Insurance Affordability: Countrywide Trends and State Comparisons," Research Brief, Insurance Research Council June 2025. Published June 10, 2025.

Keys, Benjamin J. and Philip Mulder, "Property Insurance and Disaster Risk: New Evidence from Mortgage Escrow Data," Working Paper 32579, National Bureau of Economic Research 2024.

Kim, Minjoo, Prateek Mahajan, and Zirui Wang, "Who Bears Rising Commercial Property Insurance Costs?," Working Paper 110, Hutchins Center on Fiscal and Monetary Policy, Brookings Institution July 2026. Published July 2, 2026.

Kneebone, Elizabeth and Steven Raphael, *City and Suburban Crime Trends in Metropolitan America*, Washington, DC: Brookings Institution Metropolitan Policy Program, 2011.

Linden, Leigh and Jonah E. Rockoff, "Estimates of the Impact of Crime Risk on Property Values from Megan's Laws," *American Economic Review*, 2008, *98* (3), 1103–1127.

Oh, Sangmin S., Ishita Sen, and Ana-Maria Tenekedjieva, "Pricing of Climate Risk Insurance: Regulation and Cross-Subsidies," *Journal of Finance*, 2026, *81* (3), 1161–1215.

Ostriker, Abigail and Anna Russo, "The Effects of Floodplain Regulation on Housing Markets," Working Paper, Boston University and Harvard University March 2026.

Wessel, David, "Homeowners Insurance in an Era of Climate Change," Brookings Institution commentary June 2025. Published June 17, 2025.

Zhu, Jun, Janneke Ratcliffe, John Walsh, and Bryson Berry, "Property Insurance Affordability: How Rising Costs Burden Mortgage Borrowers," Research Report, Urban Institute, Washington, DC May 2026.

A Sample Selection and Coverage

Our analysis draws on multiple data sources, each with distinct geographic coverage. This appendix documents the sample construction pipeline stage by stage, characterizes the ZIP codes included and excluded at each stage, and assesses the direction and magnitude of potential selection bias. Table 5 summarizes the ladder; every count reported in this appendix is derived from the same panel and reproduced by a single script.

Table 5: Sample Construction, Stage by Stage (2018–2022)

Stage	Zip-years	Zip codes	Restriction
FIO release (all reporting Zip codes)	127,965	25,593	25,593 Zip codes × 5 years; balanced
− non-renewal reporting gaps	119,030	23,806	drops Texas and persistent-zero non-renewal Zips
+ ACS B25119 owner-occupied income	116,610	23,709	LISI panel; income merge is the only other binding restriction
+ ZHVI (home price analyses)	106,945	22,080	Zillow Home Value Index non-missing
+ RentHub core multifamily rent	43,581	12,736	adds the rent measure
+ assigned to a CBSA (estimation sample)	42,778	12,496	CBSA×year fixed effects require a CBSA

Notes: Each row applies the stated restriction cumulatively to the row above. The FIO release is balanced: every one of the 25,593 Zip codes appears in all five years, and every FIO variable we use (premium per policy, non-renewal rate, paid loss ratio, claim frequency, claim severity) is recorded for every Zip-year. The second row removes reporting gaps that the public file records as zeros rather than as missing: Texas insurers' submissions did not include non-renewal or cancellation information (1,510 Zip codes), and 277 non-Texas Zip codes report a zero non-renewal rate in every year, where genuine zeros cannot be distinguished from unreported fields. All index standardization anchors are computed on the retained sample. The reduction from the second row to the LISI panel is entirely attributable to the ACS B25119 owner-occupied median income merge: 2,420 Zip-years lack a published owner-occupied median income, affecting 1,104 Zip codes, of which 97 are absent in all five years and 1,007 in some but not all. The LISI panel is consequently unbalanced. The final row is the estimation sample for the housing-market regressions, which require a CBSA for the CBSA×year fixed effects. Home-price regressions that do not use rents run on the 106,945 Zip-year sample in the fourth row.

A.1 Stage 1: FIO Insurance Panel

The Federal Insurance Office (FIO) Homeowners Insurance Data Collection provides our source panel: 25,593 ZIP codes observed annually over 2018–2022, or 127,965 ZIP-year observations. This release is balanced. Every one of the 25,593 ZIP codes appears in all five years, so the ZIP-year count is exactly $25{,}593 \times 5$, and every variable we use (non-renewal rate, paid loss ratio, premium per policy, claim frequency, and claim severity) is recorded for every observation, so there is no literal item-level missingness within the release.

State-specific reporting gaps and the analysis sample. Recorded is not the same as collected, and one gap in the underlying collection is disguised rather than literal. Texas insurers' submissions to the Property and Casualty Market Intelligence data call did not include information on non-renewals or cancellations (Federal Insurance Office, 2025); FIO excludes Texas from its own national non-renewal calculations, and published maps of these data display Texas as unavailable for non-renewals (Wessel, 2025). The public data file, however, records these fields as zeros rather than as missing for all 1,510 Texas ZIP codes in every year. Left in place, those zeros would enter the Unavailability pillar as if Texas insurers renewed every policy. We therefore exclude Texas from the analysis sample. We also exclude the 277 non-Texas ZIP codes (1.2% of the remainder) whose non-renewal rate is zero in every observed year, for which genuine zeros cannot be distinguished from unreported fields; no other state exhibits a structural pattern (Texas aside, no state's share of zero non-renewal ZIP-years exceeds 7% in any variable). The resulting analysis panel is 23,806 ZIP codes and 119,030 ZIP-year observations, and every standardization anchor in the index (the 2018 medians and interquartile ranges of each pillar, and the composite recentering) is computed on this retained sample.

A second, milder gap concerns Florida, which participated only partially in the data call and whose market is served disproportionately by smaller in-state insurers and the state's residual insurer (Citizens), both largely outside the collection's scope. Florida is the one state where the collection covers less than half of HO-3 and HO-5 direct premiums (Federal Insurance Office, 2025; Gourevitch and Kousky, 2025), and its covered ZIP codes represent 56% of its ZCTAs and 76% of its population, the lowest coverage among the contiguous states. Florida observations reflect the reporting national carriers rather than the full state market; because Florida's non-renewal data are populated and vary across ZIP codes, we retain Florida in the sample and note the representativeness caveat here.

Within the analysis panel, selection operates entirely on whether a ZIP code is included at all, not on what is observed within an included ZIP-year: when a ZIP code is in the panel, it is in for every variable, every year. We therefore do not face differential reporting across variables that would complicate a multivariate analysis. This is why Table 7 reports the full 119,030 observations for the non-renewal rate and the loss ratio: those two pillars are observed wherever the analysis panel is observed.

Geographic universe: ZCTAs versus USPS ZIP codes. Coverage rates for ZIP-code-level data depend heavily on which ZIP-code universe is used as the denominator, and two universes are frequently conflated. ZIP Code Tabulation Areas (ZCTAs) are Census tabulation geographies built from census blocks; the 2020 vintage contains 33,774 of them, and they are the unit on which the FIO release and the ACS are both reported. USPS ZIP codes are a mail-delivery construct numbering roughly 42,000, a total that includes PO-box, single-building, and point ZIP codes that have no residential land area and are never tabulated by the Census. The two are not interchangeable and should not share a denominator.

On the tabulation-area basis, which is the appropriate one, all 25,593 FIO ZIP codes match to a ZCTA. They represent 75.8% of all 33,774 ZCTAs and 77.6% of the 32,970 ZCTAs with non-zero housing units in the ACS, leaving 7,393 residential ZCTAs uncovered. Against the broader USPS universe the same 25,593 ZIP codes work out to roughly 61%, but that figure understates residential coverage because its denominator contains tens of thousands of non-residential delivery codes; we report it only to reconcile with coverage rates computed on that basis elsewhere. Because the uncovered ZCTAs are small and thinly populated, coverage measured in people is far higher than coverage measured in ZIP codes: the FIO release accounts for 95.8% of the U.S. population, 96.1% of the population in occupied housing units, and 96.5% of the population in owner-occupied housing units. Our analysis sample, which further excludes Texas and the persistent-zero non-renewal ZIP codes described above, covers 70.5% of all ZCTAs and 86.9% of the U.S. population; measured against the non-Texas United States, which is the analysis sample's own footprint, it covers 95.2% of the population and 95.9% of the population in owner-occupied housing units.

ZIP codes outside the release are excluded by a confidentiality-driven suppression rule (ZIP codes with fewer than 10 reporting insurers or fewer than 50 policies are dropped from the public release). Coverage of the included ZIP codes is broadly proportional to state population. California (1,552 ZIPs), New York (1,536), and Pennsylvania (1,524) have the largest representation; all 50 states plus the District of Columbia appear in the release, and 49 states plus the District of Columbia appear in the analysis sample after the Texas exclusion. The median state has roughly four hundred ZIP codes covered, with a range from a handful (Alaska, Hawaii, District of Columbia) to over fifteen hundred in the largest states. Within each state, the suppression rule disproportionately excludes thin rural markets where fewer than 10 insurers actively write coverage; the suppression rate is highest in states with substantial rural land area and lowest in densely populated coastal and Great Lakes states.

The roughly 22% of residential ZCTAs outside the panel are predominantly rural and low-density areas where the private homeowners insurance market is thin. FIO collects data from the largest private insurers, and ZIP codes with very few policies may fall below reporting thresholds or lack sufficient insurer participation to generate reliable aggregates. This selection implies that our panel underrepresents very rural communities where housing stock is sparse and insurance markets are informal, areas dominated by non-standard insurance arrangements (e.g., surplus lines carriers, self-insurance), and small ZIP codes in states where one or two regional carriers dominate and may not participate in FIO reporting.

The direction of the resulting selection bias for our main findings is ambiguous. Rural ZIP codes excluded from the FIO release are likely to have higher insurance strain (fewer competitive insurers, higher non-renewal rates, greater weather exposure) but also lower housing market liquidity. The principal limitation imposed by FIO selection is interpretive rather than statistical. Our results characterize the insurance market conditions of the densely populated United States outside Texas: the analysis sample spans 70.5% of residential ZCTAs and 87% of the U.S. population, which is 95% of the population of the 49 states and the District of Columbia it covers. Conclusions about insurance market conditions in deep rural America, or in Texas, cannot be drawn directly from our analysis. Two features of our identification strategy further attenuate concerns about the coverage gap. First, the within-CBSA-year fixed effects we use compare ZIP codes against others in the same metropolitan area and year; the ZIP codes excluded by the suppression rule contribute little within-metro variation in any case. Second, our preferred specification restricts to ZIP codes within 30 miles of their CBSA's central business district (see Appendix E), where market thickness is generally well above the FIO suppression threshold.

A.2 Stage 2: ACS Owner-Occupied Income Merge

The LISI panel contains 23,709 unique ZIP codes and 116,610 ZIP-year observations. The reduction of 2,420 ZIP-years from the filtered analysis panel is entirely attributable to the ACS merge, not to anything missing in the FIO data. As noted above, premium per policy, the non-renewal rate, and the loss ratio are recorded throughout the analysis panel; what binds instead is the Unaffordability denominator, the ACS B25119 published median household income for owner-occupied units, which the Census does not publish for every ZCTA in every vintage. The pattern is consistent with a small-area suppression mechanism: ZIP-years without a published owner-occupied median have a median of 268 housing units, against 2,366 for those with one, and three-quarters of them fall below 471 housing units. In areas that small the ACS 5-year sample is too thin to support a published median.

The gap is spread across 1,104 ZIP codes. Only 97 of them lack a published owner-occupied median in all five years and drop out of the panel entirely; the remaining 1,007 are missing in some years but not others. The LISI panel is therefore unbalanced, unlike the analysis panel from which it is drawn, and $116{,}610 < 23{,}709 \times 5$. Regressions accommodate this directly, since all specifications use fixed effects estimated on an unbalanced panel.

Because the binding constraint is an ACS publication threshold rather than an insurance-market outcome, this stage is unlikely to select on insurance strain conditional on ZIP code size. It does select on size, in the same direction as the FIO suppression rule already documented, which reinforces the interpretive caveat above rather than introducing a new one.

A.3 Stage 3: ZHVI Coverage

The Zillow Home Value Index (ZHVI) is available for 108,216 of the 119,030 analysis-panel ZIP-years, or 90.9%, and for 106,945 of the 116,610 LISI panel ZIP-years, or 91.7%. The latter is the sample for home-price analyses. ZHVI requires sufficient transaction volume and housing stock for Zillow's repeat-sales methodology to produce reliable estimates; the roughly 9% of ZIP-years lacking ZHVI tend to be in areas with very low transaction volumes, including small rural communities, declining markets, and newly developed areas without sufficient sales history.

ZHVI availability and ACS income availability do not nest. Of the ZIP-years with ZHVI, 1,271 lack a published owner-occupied median income and so fall outside the LISI panel, which is why the ZHVI-on-LISI-panel count (106,945) is smaller than the ZHVI-on-analysis-panel count (108,216) by more than the difference in panel sizes alone. Appendix C.1, which compares the primary index against a home-value-denominated alternative, restricts to the 106,945 jointly observed ZIP-years for exactly this reason. This exclusion is unlikely to introduce meaningful bias, because the excluded ZIP codes are a modest fraction of the sample and their absence is driven by thin housing markets rather than by insurance market characteristics.

A.4 Stage 4: RentHub Common Sample

Our housing market analysis uses a "common sample" of ZIP codes with non-missing data for both ZHVI and RentHub core multifamily rents. Table 6 reports year-by-year coverage.

Table 6: Rent Data Coverage: RentHub vs. Zillow ZORI

Year	Total	RentHub	ZORI	RentHub %	ZORI %	Gain
2018	23,806	9,430	1,778	39.6	7.5	7,652
2019	23,806	12,104	1,937	50.8	8.1	10,167
2020	23,806	12,029	2,198	50.5	9.2	9,831
2021	23,806	9,136	2,896	38.4	12.2	6,240
2022	23,806	13,205	3,796	55.5	15.9	9,409
Total	119,030	55,904	12,605	47.0	10.6	43,299

Notes: Unit of observation is ZIP code × year. The Total column is the analysis panel, which excludes Texas and the persistent-zero non-renewal ZIP codes and is balanced at 23,806 ZIP codes in each of the five years; see Table 5. RentHub coverage here includes all property types (multifamily and single-family); the narrower core multifamily series used as the primary rent measure in the housing-market analysis covers fewer ZIP-years. ZORI covers single-family residences only. Gain = additional ZIP-year observations from RentHub relative to ZORI.

RentHub rents of any property type cover 47.0% of analysis-panel ZIP-years (55,904 observations across 15,003 ZIP codes that are also in the analysis panel), a substantial improvement over Zillow's ZORI, which covers only 10.6% (12,605 ZIP-years across 3,796 ZIP codes). The RentHub source data span a wider footprint than the analysis panel itself, 17,324 ZIP codes in total; the 15,003 figure is the intersection with the analysis panel and is the relevant one here. Our primary rent measure is the narrower core multifamily series, which covers 43,889 analysis-panel ZIP-years. Intersecting that with the LISI panel and ZHVI leaves 43,581 ZIP-year observations, and requiring an assigned CBSA for the CBSA×year fixed effects gives the estimation sample of 42,778 ZIP-year observations across 12,496 unique ZIP codes reported in the housing-market tables.

The RentHub common sample is not a random subset of the FIO panel. RentHub coverage requires active multifamily rental listings, which are more prevalent in urban and suburban areas with sufficient rental housing stock. This generates predictable selection:

- **Urban bias.** Common sample coverage is highest within 10 miles of CBDs and declines with distance. Figure 3 documents this gradient. The common sample represents approximately 50% of ZIP codes within 80 miles of metropolitan CBDs, but coverage rates fall from over 60% in urban cores to below 30% in exurban areas.

- **Higher-value areas.** ZIP codes with active multifamily rental markets tend to be in higher-demand locations such as dense urban neighborhoods, college towns, and amenity-rich suburbs, where both home values and rents are above the national median.
- **Higher insurance strain.** The common sample exhibits higher mean LISI than the full sample at every five-mile CBD distance bin, consistent with selection toward coastal and high-value locations that face elevated insurance costs.

Figure 3: Sample Coverage by Distance from CBD (2022)

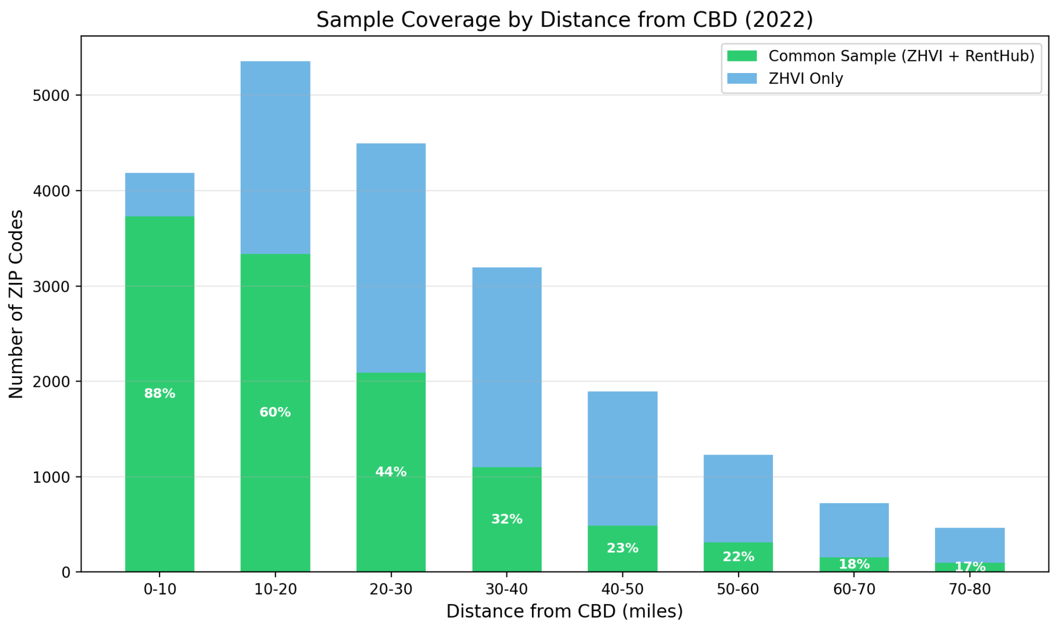

Notes: Stacked bar chart showing ZIP code counts by CBD distance for the common sample (ZIP codes with both ZHVI and RentHub core multifamily rents) and ZHVI-only ZIP codes. Percentages inside bars indicate common sample share within each distance bin. Based on 21,381 ZIP codes within 80 miles of CBSA centers, 2022 cross-section.

The implications for our estimates differ by outcome. For home prices, we report results in both the full sample and common sample; the full-sample estimate is less susceptible to selection bias. For rents and price-to-rent ratios, results are necessarily conditional on RentHub coverage and should be interpreted with caution, as they describe the relationship between insurance strain and rental markets in areas with active multifamily rental activity rather than in the broader population.

A.5 Implications for Key Findings

Our two main findings, the spatial concentration of insurance strain in the urban core and the cross-subsidization pattern across the urban-suburban gradient, are estimated on the full LISI panel using CBSA×year fixed effects and do not depend on rent data availability. The sample selection issues documented above therefore do not threaten the core identification. The CBD gradient analysis (Section 3) uses all LISI panel ZIP codes within 30 miles of CBSA centers regardless of rent coverage.

The housing market results (Section 4) are more vulnerable to selection. The negative composite LISI–home price relationship is observed in both the full sample and the common sample, so the qualitative result is robust to RentHub conditioning. The substantially larger common-sample coefficient may, however, reflect confounding from selection into urban, high-value locations.

B LISI Construction and Summary Statistics

This appendix provides the complete mathematical formulation of the Local Insurance Strain Index (LISI) summarized in Section 2, together with descriptive statistics for the index and its underlying components over the 2018–2022 panel.

B.1 Component Definitions

LISI aggregates three ZIP code–level inputs, each one component of a conceptual pillar:

Unaffordability.

$$\text{Unaffordability}_{z,t} = \frac{\text{Premium per Policy}_{z,t}}{\text{Median Owner-Occupied Income}_{z,t}}$$

Annual homeowners insurance premium per policy as a share of the median household income for owner-occupied units in ZIP code z in year t. Premiums come from the Federal Insurance Office Homeowners Insurance Data Collection; the owner-occupied income measure (ACS B25119) is drawn from the American Community Survey 5-year estimates and matched to ZIP codes via the HUD-USPS crosswalk. This definition treats Unaffordability as a cash-flow burden because annual insurance premiums are recurring obligations relative to current household resources. Appendix C.1 reports the distinct asset-normalized alternative based on home value.

Unavailability.

$$\text{Unavailability}_{z,t} = \text{NonRenewalRate}_{z,t}$$

The fraction of policies not renewed by insurers in ZIP code z in year t. Higher values indicate that insurers are exiting more aggressively from the local market.

Unprofitability.

$$\text{Unprofitability}_{z,t} = \frac{\text{Insured Losses Paid}_{z,t}}{\text{Premiums Earned}_{z,t}}$$

The paid loss ratio: insured losses paid divided by premiums earned. Higher values indicate that, conditional on the prevailing premium schedule, insurers are paying out a larger share of premiums as claims.

B.2 Robust Standardization to the 2018 Baseline

Because the three components are reported in different units and exhibit very different dispersion, we standardize each before aggregation. For each component $y_{z,t}$ we compute:

$$y^*_{z,t} = \frac{y_{z,t} - \text{Median}_{2018}(y)}{\text{IQR}_{2018}(y)}, \tag{2}$$

where $\text{Median}_{2018}(y)$ and $\text{IQR}_{2018}(y)$ are the median and interquartile range of y in the 2018 cross-section across all ZIP codes. We use the median and IQR rather than the mean and standard deviation to limit the influence of right-tail outliers, which are common in insurance-loss data. Anchoring the standardization to

Table 7: LISI Component Summary Statistics (2018–2022 Pooled)

Statistic	Unaffordability	Unavailability	Unprofitability	Composite
	Prem / HO Income (Ratio)	Non-Renewal (Rate)	Loss Ratio (Loss / Prem)	LISI (Std.)
N	116,610	119,030	119,030	116,610
Mean	0.0234	0.95%	0.3771	0.079
Std Dev	0.0154	1.13%	0.3931	0.701
P10	0.0122	0.00%	0.0587	−0.524
P25	0.0151	0.37%	0.1705	−0.317
Median	0.0199	0.72%	0.3253	−0.036
P75	0.0274	1.20%	0.5172	0.337
P90	0.0379	1.91%	0.7592	0.779

Notes: Summary statistics for the three raw components of the Local Insurance Strain Index (LISI) and the resulting composite, pooled across ZIP-year observations 2018–2022. **Unaffordability**: average annual premium per policy divided by the ACS B25119 owner-occupied median household income. **Unavailability**: fraction of policies not renewed by insurers. **Unprofitability**: loss ratio (losses paid divided by premiums earned). The three components are robust-standardized by subtracting the 2018 cross-sectional median and dividing by the 2018 interquartile range, then averaged with equal pillar weights and recentered so the 2018 LISI median equals zero. The composite N reflects ZIP-years with all three components observed.

the 2018 baseline rather than recomputing it each year allows post-2018 levels of strain to shift in absolute terms: a positive $y^*_{z,t}$ in year t indicates a ZIP code with worse conditions than the 2018 national median, regardless of how the median has moved since.

B.3 Composite LISI and Recentering

The composite index is the equal-weighted average of the three standardized pillars:

$$\widetilde{\text{LISI}}_{z,t} = \frac{1}{3}\left(\text{Unaffordability}^*_{z,t} + \text{Unavailability}^*_{z,t} + \text{Unprofitability}^*_{z,t}\right).$$

Equal weighting across pillars ensures that each conceptual dimension receives one-third of the total weight, rather than being implicitly weighted by the dispersion of its underlying input. Because the median of averages does not generally equal the average of medians when components have imperfect rank correlation, the 2018 cross-sectional median of $\widetilde{\text{LISI}}$ is not exactly zero. We therefore recenter the composite by subtracting its 2018 cross-sectional median:

$$\text{LISI}_{z,t} = \widetilde{\text{LISI}}_{z,t} - \text{Median}_{2018}\left(\widetilde{\text{LISI}}\right).$$

With this two-stage standardization, $\text{LISI}_{z,t} = 0$ corresponds exactly to the 2018 national median: positive values indicate worse-than-baseline insurance market conditions in ZIP code z in year t, negative values indicate the opposite.

B.4 Summary Statistics

Table 7 reports pooled summary statistics for the three raw LISI inputs and the resulting composite index, computed on the analysis panel described in Appendix A. The typical ZIP code has an annual homeowners insurance premium equal to roughly 2.0% of owner-occupied median income, a non-renewal rate just above 0.7%, and a paid loss ratio of about 0.33. The right tail is substantial in all three pillars: the 90th-percentile ZIP code pays nearly 4% of homeowner income in premiums, has a non-renewal rate close to 2%, and pays out roughly 76 cents in claims for every dollar of premium earned. This dispersion across ZIP codes, rather than the typical level, is what gives LISI its identifying variation.

Table 8 traces the LISI distribution year by year. By construction the 2018 median is zero. Median LISI is slightly negative in every subsequent year (reaching −0.01 in 2022), so the typical ZIP code experienced marginally better-than-baseline conditions through the end of the panel. At the same time, the standard deviation rises from 0.62 in 2018 to 0.75 in 2022, and the 99th percentile remains between 1.9 and 2.3 throughout. Aggregate strain has not deteriorated for the typical ZIP code, but the distribution has widened: a small share of ZIP codes face progressively more acute conditions while most do not.

Table 8: LISI Distribution by Year (2018–2022)

Year	N	Mean	Std Dev	P10	P25	Median	P75	P90	P95	P99
2018	23,388	0.115	0.622	−0.498	−0.292	−0.000	0.391	0.820	1.177	2.271
2019	23,373	0.086	0.652	−0.515	−0.308	−0.042	0.323	0.783	1.178	2.317
2020	23,271	0.022	0.809	−0.558	−0.353	−0.081	0.280	0.700	1.015	2.019
2021	23,301	0.061	0.653	−0.527	−0.321	−0.042	0.316	0.740	1.032	1.900
2022	23,277	0.109	0.750	−0.525	−0.309	−0.013	0.373	0.845	1.217	2.304

Notes: Distribution statistics for the Local Insurance Strain Index (LISI) by year. LISI is normalized so that the 2018 national median equals zero; positive values indicate above-median strain, negative values indicate below-median strain. The index is the equal-weighted average of three robust-standardized pillars: Unaffordability (premium per policy as a share of homeowner income), Unavailability (non-renewal rate), and Unprofitability (loss ratio).

C Alternative LISI Specifications

This appendix tests the robustness of the Local Insurance Strain Index along two construction dimensions: the denominator of the Unaffordability pillar and the structure of the pillars themselves.

The first two subsections hold the three-pillar structure fixed and vary only the Unaffordability denominator, so that any change in the results is attributable to that single choice. Subsection C.1 presents **LISI-Asset**, which replaces household income with the Zillow Home Value Index (ZHVI), asking whether the spatial pattern we document depends on measuring the premium burden against income rather than against home value. Subsection C.2 presents **LISI-Mean**, which retains an income denominator but replaces the median with the mean owner-occupied household income computed from the same ACS source. Subsection C.3 then turns to **LISI Extended**, a five-component construction that alters the pillar structure itself; because it changes the burden and cost-pressure pillars simultaneously, it does not isolate the denominator question, and we summarize it only briefly here.

All three alternatives preserve the median/IQR robust standardization, the equal-weighted pillar aggregation, and the 2018-baseline recentering used in the primary specification.

C.1 Home Value Denominator: LISI-Asset

The primary specification measures Unaffordability as a recurring premium obligation relative to household income. This cash-flow denominator is our preferred measure because homeowners insurance premiums must be met from current household resources, whereas home value is an illiquid asset and does not directly measure the capacity to meet recurring payments. An asset-based normalization nevertheless captures a distinct and useful concept, i.e., the annual insurance carrying cost relative to the market value of the home. We therefore reconstruct the index by replacing income with the Zillow Home Value Index:

$$\text{Unaffordability}^{\text{Asset}}_{z,t} = \frac{\text{Premium per Policy}_{z,t}}{\text{Home Value (ZHVI)}_{z,t}}.$$

Although the underlying FIO data call requested aggregate Coverage A limits, Treasury's public ZIP-level release used in this paper does not include that field. We therefore cannot construct premiums per dollar of dwelling coverage. ZHVI therefore provides an available but imperfect asset-based benchmark because market value includes land and location value and is not equivalent to the insured replacement value of the structure. The LISI-Asset is otherwise identical to the primary LISI, i.e., the Unavailability and Unprofitability pillars are unchanged, all three pillars are robust-standardized against their own 2018 cross-sectional median and IQR, the three are averaged with equal weight, and the composite is recentered so its 2018 cross-sectional median equals zero. The question the subsection answers is whether the urban concentration of insurance strain documented in Section 3 is an artifact of measuring burden against income rather than against home value.

Comparison with the primary specification. Table 9 reports pooled summary statistics for the two composite indices and the two Unaffordability ratios that distinguish them. Because the subsection's claim is that the denominator is the sole construction difference, the table is computed on the common sample of 106,945 ZIP-years on which both indices are observed, so that no part of the comparison is attributable to differences in geographic coverage. The two supports do not nest: the primary index is available for 116,610 ZIP-years and the asset-denominated index for 108,216, since ZHVI reaches fewer ZIP codes than ACS income overall while a small number of ZIP codes carry a ZHVI value without a published owner-occupied income. The value-denominated ratio is roughly one-third as large in raw units—a pooled median of 0.0068, or $6.84 of annual premium per $1,000 of home value, compared with 0.0196 for the income-denominated version—because home values are roughly three times homeowner household income in the pooled data. More substantively, the two ratios are far from interchangeable: their cross-sectional correlation is $\rho = 0.539$, well below the $\rho = 0.955$ between the median- and mean-income ratios in Section C.2. Income and home values are distinct bases, and a ZIP code that looks burdened relative to one need not look burdened relative to the other. The composite indices nonetheless correlate at $\rho = 0.868$ pooled across 2018–2022 (0.80 to 0.93 within years), and their pooled distributions are close, with medians of -0.0322 and -0.0529 and interquartile ranges of 0.641 and 0.623.

The burden gradient in natural units and after standardization. Figure 4 presents the value-denominated burden twice: panel (a) in natural units, as dollars of annual premium per $1,000 of home value in 2018 and 2022, and panel (b) after robust standardization, as the pillar actually enters the index and alongside the primary income-denominated pillar.

Panel (a) makes the magnitude concrete and speaks to a distinct concern, namely whether the unusual housing-market dynamics of 2018–2022 reshaped the spatial pattern. The CBSA $\times$ year fixed effects used throughout absorb shocks common to all ZIP codes within a metropolitan area and year, but they do not eliminate *differential* within-metro shifts between CBD and suburban ZIP codes. Asset burden makes those shifts directly observable, since it falls when home values rise faster than premiums. Two features stand out. First, the level fell substantially across the entire gradient, from a mean of $8.89 per $1,000 in 2018 to $6.36 in 2022, a 28% decline reflecting broad-based home-value appreciation outpacing premium growth. Second, and more importantly, the shape is preserved: burden peaks at the CBD, declines sharply through the first ten miles, troughs in the mid-teens, and rises modestly through the outer bins in both years. The CBD-to-trough gap narrows in dollar terms ($4.93 to $3.65) but is essentially unchanged in proportional terms, with the CBD bin standing 65% above the trough in 2018 and 66% above it in 2022. The dollar compression is therefore a mechanical consequence of the broad level decline rather than evidence of within-metro reallocation, and the urban asset-burden gradient documented here survives the period intact.

Panel (b) shows that the value-denominated pillar reproduces the primary gradient in shape ($\rho = 0.936$

Table 9: Primary LISI versus LISI-Asset: Composite Index and Unaffordability Ratio (2018–2022 Pooled, Common Sample)

	Composite Index		Unaffordability Ratio	
	LISI (primary)	LISI-Asset	LISI (Prem/Income)	LISI-Asset (Prem/Home Value)
N	106,945	106,945	106,945	106,945
Mean	0.0824	0.0613	0.0231	0.0084
Std Dev	0.6953	0.6505	0.0155	0.0061
P10	-0.5101	-0.5088	0.0121	0.0030
P25	-0.3079	-0.3193	0.0149	0.0045
Median	-0.0322	-0.0529	0.0196	0.0068
P75	0.3328	0.3034	0.0270	0.0105
P90	0.7679	0.7426	0.0373	0.0151

Notes: Pooled summary statistics, 2018–2022, on the **common sample** of 106,945 Zip code×year observations (22,080 unique Zip codes) for which both composite indices are observed. All four columns are computed on identical Zip-years, so the differences reported here reflect the Unaffordability denominator alone and not any difference in geographic coverage. Unrestricted, the primary index is available for 116,610 Zip-years and LISI-Asset for 108,216; neither support nests the other, because ZHVI reaches fewer Zip codes than ACS B25119 income overall while a small number of Zip codes carry a ZHVI value without a published owner-occupied income. **The two column groups are reported in different units and are not comparable across the group divide.** Columns (1)–(2) report the composite index, which is unitless: each of the three pillars is robust-standardized against its own 2018 cross-sectional median and IQR, the three standardized pillars are averaged with equal weight, and the composite is recentered so that its 2018 cross-sectional median equals zero. Both indices are constructed exactly as defined, on their own full support, and only then restricted to the common sample for this table; the index values are therefore identical to those used elsewhere in the paper, and the recentering constant is the one implied by each index's own 2018 cross-section, so neither pooled median is mechanically zero here. Columns (3)–(4) report the underlying Unaffordability pillar in natural units. The LISI (primary) column divides average premium per policy by ACS B25119 median owner-occupied household income; the LISI-Asset column divides by the Zillow Home Value Index, so the pooled median of 0.0068 corresponds to a premium of $6.84 per $1,000 of home value. The Unavailability (non-renewal rate) and Unprofitability (loss ratio) pillars are omitted because they are identical by construction across the two specifications; the Unaffordability denominator is the sole difference between the two indices.

across distance bins) and steepens it in magnitude: the CBD-to-outer-ring gap widens from $+0.33$ under the income denominator to $+0.45$ under the home-value denominator. Premiums therefore rise more steeply toward the urban core than home values do, so switching from income to home value as the burden base *widens* the urban penalty rather than explaining it away. The concentration of insurance strain in the urban core is not an artifact of the income denominator.

Figure 4: Insurance Burden Against Home Value Across the Urban Gradient

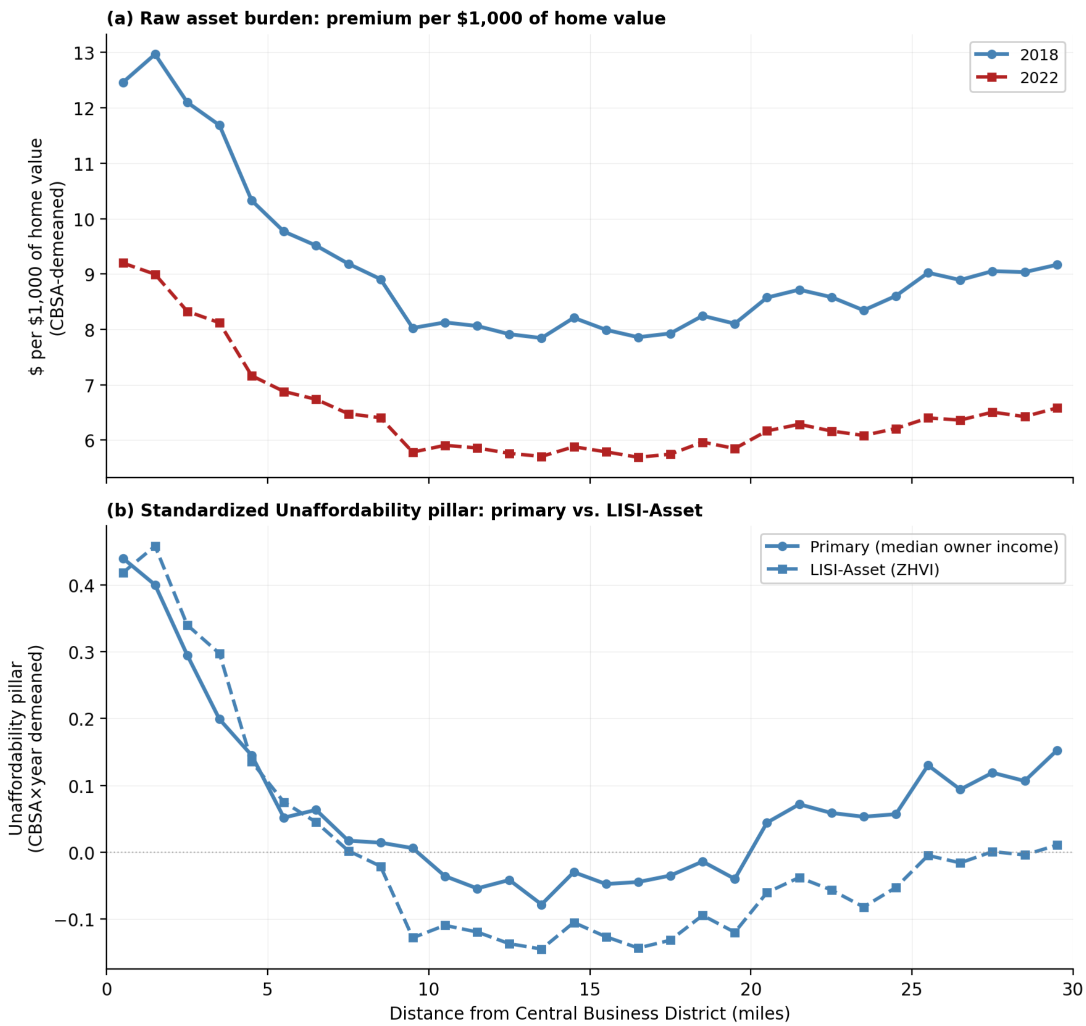

Notes: Panel (a) plots mean asset burden (premium per $1,000 of home value) by distance from the CBD, separately for 2018 and 2022, on a balanced panel of ZIP codes observed in both years. Each year is CBSA-demeaned within itself: $\tilde{y}_z = (y_z - \bar{y}_m) + \bar{y}$, where $\bar{y}_m$ is the CBSA mean and $\bar{y}$ the grand mean for that year. Based on 12,768 ZIP codes across 324 CBSAs observed in both years. Panel (b) plots the mean standardized Unaffordability pillar by distance from the CBD, comparing the primary specification (solid line, premium per policy divided by ACS B25119 median owner-occupied income) to LISI-Asset (dashed line, premium per policy divided by ZHVI), pooled over 2018–2022 and CBSA×year demeaned; based on 64,250 ZIP-year observations across 324 CBSAs. Both panels use 1-mile distance bins from 0 to 30 miles and restrict to CBSA cells with at least three ZIP codes.

The home-value normalization is, if anything, conservative as a proxy for insurance-price intensity. In the standard monocentric-city model, and on average in U.S. appraisal data, land accounts for a larger share of property value near the CBD (Davis et al., 2021). To the extent that this implies a smaller structure-replacement-cost share near the urban core, holding premiums per dollar of dwelling coverage fixed would produce a lower premium-to-home-value ratio near the CBD. We find the opposite. Conditional on this structure-share gradient, the observed pattern would imply that premiums per dollar of insurable structure rise more steeply toward the CBD than the premium-to-home-value gradient alone indicates. Because the public FIO data do not report Coverage A and homeowners premiums finance coverages beyond the dwelling structure, we interpret this as a conditional implication rather than a direct estimate of insurance-price intensity.

Why we do not report housing-market capitalization under LISI-Asset. We deliberately omit the housing-capitalization regressions from this subsection. The LISI-Asset Unaffordability pillar is constructed by dividing premiums by the Zillow Home Value Index, which is the dependent variable in those regressions. A coefficient estimated by regressing log home prices on an index built from log home prices would be mechanically negative regardless of any economic relationship, and its magnitude would reflect the construction rather than the capitalization of insurance strain. This is the same consideration that leads the primary specification to exclude the asset-burden term carried by LISI Extended, discussed in Section C.3. LISI-Asset is therefore informative about the *spatial distribution* of insurance strain, which Figure 4 establishes, but it is not a valid right-hand-side variable for the housing-market analysis of Section 4, and we make no claim from it on that margin.

C.2 Mean Income Denominator: LISI-Mean

The primary LISI Unaffordability pillar measures the affordability burden as

$$\text{Unaffordability}_{z,t} = \frac{\text{Premium per Policy}_{z,t}}{\text{Median Owner-Occupied Income}_{z,t}},$$

using the ACS B25119 published median household income for owner-occupied units. We defend the median against the mean in the Section 2 footnote on the grounds that the mean is more sensitive to outliers and could understate the affordability burden faced by the typical homeowner. This subsection tests that argument empirically by reconstructing LISI under the mean rather than the median.

Mean owner-occupied income from ACS B25120 and B25003. The ACS does not publish a direct mean owner-occupied household income at the ZIP code tabulation area level, but it does publish the components from which a mean can be computed in a manner methodologically symmetric to the B25119 published median: ACS table B25120 (*Aggregate Household Income in the Past 12 Months by Tenure and Mortgage Status*) publishes the aggregate income of owner-occupied households, and ACS table B25003 (*Tenure*) publishes the count of owner-occupied households. Both are derived by the Census Bureau from the underlying microdata. We compute

$$\overline{\text{Income}}_{z,t} = \frac{\text{B25120}_{002,z,t}}{\text{B25003}_{002,z,t}}.$$

Across the 112,135 ZIP-years jointly observed with the LISI panel, this published-derived mean exceeds the B25119 published median in 97.8% of observations, with a median mean-to-median ratio of 1.219 and a mean ratio of 1.250. This pattern is consistent with greater upper-tail income dispersion in central-city ZIP codes.

LISI-Mean construction. LISI-Mean is identical to the primary LISI in every respect except for the income denominator: the Unaffordability pillar uses $\overline{\text{Income}}_{z,t}$ in place of the ACS B25119 median. All three pillars are robust-standardized using their own 2018 cross-sectional median and IQR, the composite is the equal-weighted average of the three standardized pillars, and the composite is recentered so its 2018 cross-sectional median equals zero—the same two-stage construction described in Appendix B.

Comparison with the primary specification. The mean-income-based Unaffordability ratio is systematically lower than the median-based version in raw units, because dividing by a larger denominator shrinks the ratio: the pooled medians are 0.0162 and 0.0199 respectively. After robust standardization the two pillars are in comparable units and correlate at $\rho = 0.956$ in the cross-section. The two composite indices are more tightly related still, at $\rho = 0.987$ pooled across 2018–2022 and ranging from 0.98 to 0.99 within years, and their pooled distributions are nearly indistinguishable, with medians of -0.0365 and -0.0355 and interquartile ranges of 0.645 and 0.654. The Unavailability and Unprofitability pillars are identical by construction across the two specifications, since the income statistic is the sole point of difference. The median-versus-mean choice does not alter the sign, spatial shape, nor qualitative conclusions of the analysis.

The Unaffordability pillar across the urban gradient. Figure 5 traces both Unaffordability pillars across distance from the CBD. The two series are nearly coincident beyond roughly ten miles ($\rho = 0.966$ across distance bins) but separate sharply inside the urban core: at the innermost bin the primary pillar stands at 0.50 against 0.26 for LISI-Mean, and the CBD-to-outer-ring gap falls from $+0.37$ to $+0.16$. The mean denominator thus compresses the urban affordability penalty by more than half. This is the expected consequence of right-skewed income distributions being more right-skewed in central cities: high earners in the urban core pull the mean further above the median than they do in suburban ZIP codes, so dividing by the mean shrinks the measured burden disproportionately where incomes are most dispersed. The compression is a reason to prefer the median, not evidence against the urban penalty—the gradient retains its shape, its sign, and its sharp decline through the first ten miles under either denominator.

Figure 5: Unaffordability Pillar Across the Urban Gradient: Primary LISI vs. LISI-Mean

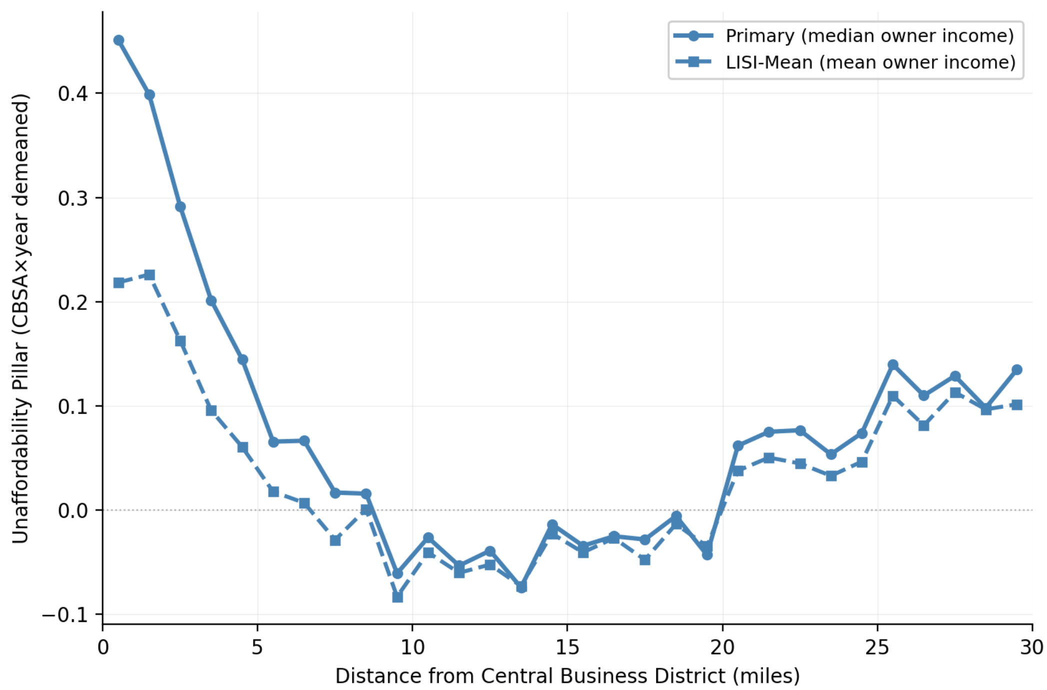

Notes: Mean standardized Unaffordability pillar by distance from the central business district, comparing the primary specification (solid line, premium per policy divided by ACS B25119 median owner-occupied income) to LISI-Mean (dashed line, premium per policy divided by mean owner-occupied income). Both series are CBSA×year demeaned (within each metropolitan area and year, the cell mean is subtracted and the grand mean added back), matching the specification used in the housing-capitalization regressions. Distance bins are 1-mile intervals from 0 to 30 miles, pooled over 2018–2022. Sample restricted to CBSA×year cells with at least three ZIP codes. Based on 65,145 ZIP-year observations across 324 CBSAs.

C.3 Five-Component LISI Extended

An earlier construction of the index, which we label "LISI Extended," builds the same three pillars from five underlying components rather than three.

Household Burden. Two components capture the household-level cost of insurance, weighted equally:

$$\text{CashFlowBurden}_{z,t} = \frac{\text{Premium per Policy}_{z,t}}{\text{Median Owner-Occupied Income}_{z,t}},$$
$$\text{AssetBurden}_{z,t} = \frac{\text{Premium per Policy}_{z,t}}{\text{Home Value}_{z,t}}.$$

The cash-flow burden uses the same ACS B25119 median owner-occupied household income as the primary specification's Unaffordability ratio. The asset-based burden expresses premiums as a share of the Zillow Home Value Index.

Availability. The fraction of policies not renewed by insurers, $\text{NonRenewalRate}_{z,t}$, identical to the Unavailability pillar in the primary specification.

Cost Pressure. Two components decompose realized loss experience:

$$\text{Frequency}_{z,t} = \frac{\text{Number of Claims}_{z,t}}{\text{Number of Policies}_{z,t}},$$
$$\text{Severity}_{z,t} = \frac{\text{Total Insured Losses}_{z,t}}{\text{Number of Claims}_{z,t}}.$$

Together these recover total losses per policy (Frequency $\times$ Severity), which the primary specification instead captures through the paid loss ratio.

Each of the five components is robust-standardized to the 2018 cross-section using

$$y^*_{z,t} = (y_{z,t} - \text{Median}_{2018}(y))/\text{IQR}_{2018}(y).$$

Within Household Burden and Cost Pressure the two standardized components are averaged with equal weight; Availability is its single standardized component. The three pillar sub-indices are then averaged with equal weight, and the composite is recentered so that its 2018 cross-sectional median equals zero.

We report LISI Extended for completeness but do not rely on it, for three reasons. First, its Household Burden pillar contains premium-to-home-value, which is constructed directly from the Zillow Home Value Index and would mechanically contaminate the housing-capitalization estimates of Section 4; this is the same consideration discussed in Subsection C.1. Second, the frequency-severity decomposition is unstable in ZIP codes with few or zero claims, whereas the paid loss ratio remains well-defined there and is the measure most directly relevant to insurer underwriting and exit decisions. Third, with two pillars containing two components and one containing a single component, the implicit weight on an individual input is no longer symmetric across inputs. Because LISI Extended changes the burden pillar and the cost-pressure pillar simultaneously, it also cannot isolate the denominator question that Subsection C.1 addresses, which is why the asset-denominated index rather than LISI Extended carries that argument.

Despite these differences, the broad spatial pattern survives. The pooled Pearson correlation between LISI Extended and the primary LISI is $\rho = 0.659$, so the two are substantively different indices in the cross-section, yet their CBD gradients track closely (Figure 6): both show a pronounced disadvantage at the CBD that declines sharply through the first five to ten miles and then stabilizes. The composites track because the amplification in the burden pillar is offset by noise in the cost-pressure pillar, and the availability pillar is identical by construction. Panel (d) makes the last point concrete: the two cost-pressure measures are in fact negatively correlated across distance bins, because claim frequency and severity swing sharply in ZIP codes with few claims while the loss ratio remains well-defined there. That instability is the clearest single reason we build the primary index on the loss ratio.

Figure 6: Primary LISI vs. LISI Extended Across the Urban Gradient

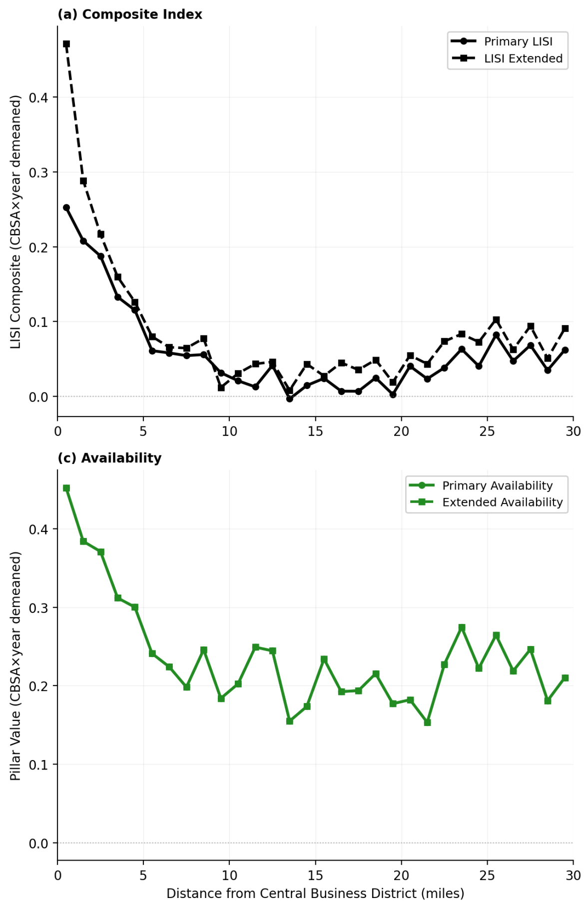

Notes: Four-panel comparison of primary LISI (solid lines) and LISI Extended (dashed lines) across distance from the central business district. All values are CBSA×year demeaned (within each metropolitan area and year, the cell mean is subtracted and the grand mean is added back), matching the specification used in housing-capitalization regressions. Panels: (a) Composite index ($\rho = 0.976$ across distance bins); (b) Affordability / Household Burden pillar ($\rho = 0.974$); (c) Availability pillar ($\rho = 1.000$, identical by construction); (d) Cost Pressure / Unprofitability pillar ($\rho = -0.081$). Distance bins are 1-mile intervals, 0–30 miles, pooled over 2018–2022. Sample restricted to CBSA×year cells with at least three ZIP codes, matching Figure 2: 67,170 ZIP-year observations across 13,591 ZIP codes and 324 CBSAs.

D Hybrid CBD Construction and Validation

For each Core Based Statistical Area (CBSA), we define the central business district (CBD) as the centroid of the highest-density Census tract within the CBSA. We use a hybrid construction that prioritizes employment density (which better identifies functional employment cores in large metros) and falls back to population density in jurisdictions where employment-density data are unavailable.

D.1 Primary source: LEHD LODES8 Workplace Area Characteristics

For CBSAs in the contiguous 48 states plus Hawaii, we identify the CBD by ranking each Census tract within the CBSA by total jobs per square mile, using the U.S. Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) Origin-Destination Employment Statistics (LODES) Workplace Area Characteristics (WAC) file for 2022. LEHD WAC reports total primary jobs at the Census block level, which we aggregate to tracts using the 2020 block-to-tract crosswalk. The tract with the highest jobs-per-square-mile within each CBSA is designated the CBD; its 2020 Census centroid (latitude, longitude) becomes the CBD coordinate for that CBSA. Population-density alternatives—such as identifying the tract with the highest residential population per square mile—tend to misidentify the CBD in metropolitan areas where residential density and employment density diverge (e.g., New York, where Midtown is the employment core but residential density peaks in Manhattan's Upper East Side and Upper West Side).

D.2 Fallback: 2020 Census population density

LODES does not cover Alaska or Puerto Rico, and Connecticut transitioned from county-based to Planning Region-based geographies in 2022, creating a temporary mismatch between LODES tract codes and our reference geography. For CBSAs in these jurisdictions we fall back to identifying the highest-population-density tract using the 2020 Decennial Census. The fallback affects three CBSAs in Connecticut, one in Alaska (Anchorage), and one in Puerto Rico (San Juan); for all other CBSAs the primary employment-density construction is used.

D.3 Validation

We validate the hybrid CBD assignment in two ways. First, we manually inspect the top 50 CBSAs by population and confirm that the identified CBD tract corresponds to the recognized central business district in each case (e.g., the Loop in Chicago, Downtown LA, Midtown Manhattan, Downtown Houston). Second, in robustness exercises we re-construct CBDs using alternative criteria—highest population density, highest total jobs (level rather than density), centroid of the principal city—and confirm that the within-metro distance gradients we report are qualitatively unchanged. The choice between employment-density and population-density CBD identification matters most in a small number of large, polycentric metros (New York, Los Angeles, the San Francisco Bay Area), where the choice shifts a handful of ZIP codes between distance bins but does not alter the overall gradient.

E The 30-Mile CBD Window and Alternative Geographic Restrictions

The 30-mile restriction defines a common absolute-radius urban-suburban window rather than an estimate of each CBSA's full spatial extent. We measure location in miles because the mechanisms studied in the paper—physical proximity to the employment core, density, accessibility, and density-correlated loss exposures—are naturally expressed in physical distance from the employment core. A percentile rank of observed ZIP-code distances is not a neutral rescaling: it depends on the number, size, and spatial arrangement of ZCTAs within each CBSA, and it mechanically distributes each CBSA's ZIP codes across the relative-distance scale even when many ZIP codes are tightly concentrated near the CBD. Nor do we normalize distance by a CBSA's maximum radial extent. CBSA boundaries are irregular, county-based, and not radially symmetric, so a maximum-distance denominator may be determined by a narrow geographic protrusion or a single distant ZIP code rather than by the economically relevant extent of the metropolitan area.

The analytical data also provide broad common support over the baseline window. Of the 324 CBSAs in the gradient sample, 311 (96%) have observed ZIP-code support reaching at least 30 miles from the CBD, whereas support declines progressively at greater distances: 85% of CBSAs reach 40 miles, 67% reach 50 miles, 53% reach 60 miles, and 30% reach 80 miles. Support is likewise broad throughout the interior of the window: 304 CBSAs (94%) contribute observations to the 25–30-mile band, and no five-mile band between the CBD and 30 miles draws on fewer than 297 of the 324 CBSAs. The main LISI gradient is also qualitatively unchanged when the sample is restricted to the 304 CBSAs with observed support in the outermost band: the restricted subsample retains 97.5% of the gradient observations, and its thirty one-mile bin means are nearly identical to the full-sample gradient (correlation 0.998, with the innermost-mile mean at $+0.26$ versus $+0.28$). These results indicate that the sharp urban-core gradient is not generated by a changing set of metropolitan areas within the baseline range. The baseline window also contains approximately 66% of the U.S. population, which is 72% of the population of the states the analysis sample spans, and Table 10 shows that the housing-market associations retain their signs and qualitative interpretation under alternative distance caps.

Table 10 reports the per-standard-deviation LISI effect on log home prices, log rents, and the log price-to-rent ratio for the common sample at each cap. Three patterns are visible. First, every coefficient retains its sign across all caps. Second, the magnitude of each effect grows monotonically as the cap tightens—from -8.9% on home prices at no cap to -20.2% at the 10-mile cap. This pattern is consistent with the spatial heterogeneity in capitalization documented in Section 4: the LISI effect is sharply concentrated near the CBD, so restricting the sample toward the urban core mechanically increases the average effect. Third, the 30-mile cap chosen in the main text sits comfortably within the range of qualitatively similar estimates; it is restrictive enough to focus on the urban-suburban gradient (excluding rural and far-exurban ZIPs that are outside the paper's scope) without being so tight that the sample loses ZIP codes essential for identifying the within-metro gradient.

Table 10: LISI Capitalization Across Alternative CBD Distance Caps (Per 1-SD, Common Sample)

CBD Cap	ZIP-Years	Home Prices	Rents	Price-to-Rent
No cap	40,156	−8.9%	−2.4%	−6.5%
80 miles	39,502	−9.1%	−2.6%	−6.6%
50 miles	38,130	−9.3%	−2.6%	−6.7%
30 miles (main text)	**33,770**	**−10.6%**	**−3.1%**	**−7.5%**
20 miles	26,963	−13.9%	−4.5%	−9.4%
15 miles	21,904	−17.7%	−5.3%	−12.4%
10 miles	15,192	−20.2%	−5.6%	−14.6%

Notes: OLS regressions of log home prices, log rents, and log price-to-rent ratio on LISI with CBSA×year fixed effects. Each row corresponds to a different geographic restriction on the analytic sample, defined by the maximum allowable distance from each ZIP code's CBSA's CBD. The common sample includes ZIP codes with non-missing data on LISI, the Zillow Home Value Index, and RentHub core multifamily rents. Coefficients are reported per one standard deviation of LISI in the analytic sample. Standard errors clustered at the CBSA×year level (not shown); all reported coefficients are statistically significant at the 1% level.

The growth of the reported per-standard-deviation association as the cap tightens is consistent with the spatial heterogeneity that we document in the interaction specification (Table 4): the LISI–price relationship is largest near the urban core and decays smoothly to approximately zero by 30 miles out. Tightening the cap toward 10 miles effectively concentrates the analytic sample at the high-effect tail of the spatial distribution, mechanically increasing the average LISI coefficient. The 30-mile cap captures the bulk of the spatial heterogeneity (extending to where the LISI effect has decayed to near-zero) while preserving sample size for power.

The robustness exercise confirms that the qualitative findings in the main text—negative capitalization of LISI into home prices and rents, together with a decline in the price-to-rent ratio—do not depend on the specific 30-mile cap. They are present across all geographic restrictions, and grow in magnitude as the cap tightens around the urban core. The 30-mile cap balances the paper's substantive focus on the urban-suburban gradient against the need to retain sufficient sample for precise inference.

F Natural Hazard Exposure and Insurance Strain

This appendix collects the two hazard-related analyses that support the risk measures used in the main text. Subsection F.1 establishes which of the 18 NRI hazard types predict insurance strain. Subsection F.2 documents the spatial gradient of each individual hazard type, which underlies the composite gradient shown in the main text.

F.1 Hazard Predictors of Insurance Strain

This subsection presents the full analysis of which natural hazards predict insurance strain, summarized in the main text in Section 3.

F.1.1 Which Natural Hazards Predict Insurance Strain?

We examine which of the 18 NRI natural hazard risk types most strongly predict insurance market strain as measured by LISI. We regress LISI on all 18 hazard-specific annual frequency measures using three specifications: (1) 2022 levels without fixed effects, (2) 2022 levels with CBSA fixed effects, and (3) the 2018–2022 change in LISI with CBSA fixed effects. The fixed-effects columns identify relationships using only within-CBSA variation, comparing ZIP codes against others in the same metropolitan area.

Our analysis reveals that hurricane risk and wildfire risk dominate LISI predictions, consistent with these hazards generating the largest insured losses and driving insurer behavioral responses. Other hazards contribute minimally to explaining insurance market strain patterns.

Table 11 presents OLS regressions of LISI on all 18 NRI hazard-specific annual frequency measures. Column (1) shows baseline results without fixed effects using the 2022 cross-section; Column (2) includes CBSA fixed effects; Column (3) uses the change in LISI from 2018 to 2022 as the dependent variable with CBSA fixed effects.

In the baseline specification, wildfire exhibits by far the strongest association with insurance market strain (coefficient = 31.1***, indicating a 1-unit increase in annual wildfire frequency is associated with 31 units higher LISI). Hurricane (1.45***) and tornado (0.32***) also show positive and significant effects. Several hazards show significant negative associations (riverine flooding -0.020***, strong wind -0.038***, cold wave -0.023**), likely reflecting confounding: areas with high exposure to these hazards may systematically differ in ways that offset the direct effect of physical risk on insurance outcomes.

Adding CBSA fixed effects fundamentally changes the interpretation: we now identify hazard effects using only within-CBSA variation across ZIP codes. This absorbs all metropolitan-level characteristics—regulatory regimes, market structure, housing stock quality, demographics—that might confound cross-sectional comparisons.

The results reveal three critical patterns. First, wildfire and hurricane effects persist with CBSA FE. Wildfire remains the dominant predictor (27.2***, only modestly reduced from 31.1), and hurricane remains significant (0.80**). These patterns indicate that within metropolitan areas exposed to these hazards, meaningful within-metro risk gradients exist: ZIP codes facing higher hurricane or wildfire frequency within the same regulatory and market environment experience substantially higher insurance strain. Second, many

Table 11: Natural Hazard Predictors of Insurance Market Strain

	(1) $LISI_{2022}$	(2) $LISI_{2022}$	(3) $\Delta LISI_{2018-2022}$
Avalanche	0.316	0.167	0.222
	(0.295)	(0.305)	(0.300)
Coastal Flooding	−0.0046	0.0133	0.0150*
	(0.0074)	(0.0109)	(0.0091)
Cold Wave	−0.023**	−0.019	−0.012
	(0.011)	(0.023)	(0.022)
Drought	0.0059***	0.0019	0.0007
	(0.0004)	(0.0013)	(0.0012)
Earthquake	6.259	−6.418	−2.975
	(5.120)	(7.106)	(4.473)
Hail	0.057***	0.042***	0.037**
	(0.005)	(0.014)	(0.015)
Hurricane	1.447***	0.796**	−0.379
	(0.163)	(0.397)	(0.240)
Heat Wave	0.0228***	0.0004	−0.0175**
	(0.0026)	(0.0070)	(0.0077)
Ice Storm	−0.0042	−0.0729***	−0.0258
	(0.0098)	(0.0278)	(0.0239)
Landslide	−0.673	−0.871	−0.917
	(0.909)	(0.970)	(0.993)
Lightning	0.0030***	0.0004	0.0014**
	(0.0006)	(0.0006)	(0.0006)
Riverine Flooding	−0.0196***	−0.0000	0.0094*
	(0.0032)	(0.0079)	(0.0056)
Strong Wind	−0.038***	−0.030***	−0.018**
	(0.003)	(0.010)	(0.009)
Tornado	0.324***	0.130	0.021
	(0.058)	(0.089)	(0.095)
Tsunami	0.058	0.285	−0.733
	(1.684)	(1.526)	(1.144)
Volcanic Activity	7.721	21.318	17.229
	(10.583)	(17.900)	(14.886)
Wildfire	31.135***	27.245***	20.705***
	(3.895)	(6.060)	(6.023)
Winter Weather	−0.0052	0.0223***	0.0083
	(0.0032)	(0.0078)	(0.0055)
CBSA FE	No	Yes	Yes
R-squared	0.162	0.319	0.219
Observations	22,982	22,982	22,787

Notes: OLS regressions of insurance market strain on all 18 NRI hazard-specific annual frequency measures. Columns (1)–(2) use 2022 LISI levels as the dependent variable; column (3) uses the change in LISI from 2018 to 2022 (ΔLISI = $LISI_{2022}$ - $LISI_{2018}$). Column (1) presents baseline specification without fixed effects (robust HC3 standard errors). Columns (2)–(3) include CBSA fixed effects, identifying effects from within-CBSA variation across ZIP codes; standard errors clustered at the CBSA level. LISI is normalized such that its 2018 national median equals zero; positive values indicate worse insurance market conditions. The long-difference specification in column (3) tests whether hazard exposure predicts *changes* in insurance strain over the sample period. Significance: *** p<0.01, ** p<0.05, * p<0.10.

hazard coefficients lose significance or attenuate sharply with CBSA FE: tornado drops from 0.32*** to 0.13 (insignificant), hail attenuates from 0.057*** to 0.042***, and earthquake flips sign from $+6.3$ to -6.4, with neither estimate precise. These sign flips and attenuations indicate confounding by metropolitan-level characteristics. Third, R-squared increases substantially with CBSA FE: R^2 rises from 0.162 to 0.319, indicating that metro-level factors explain substantial variation in LISI beyond what hazard exposure alone captures.

Column (3) tests whether hazard exposure predicts *changes* in insurance strain from 2018 to 2022. Wildfire remains the dominant predictor of ΔLISI (20.7***), although hurricane is no longer significant (-0.38). The wildfire result indicates that high-wildfire-risk areas saw differential deterioration over the sample period relative to low-risk areas within the same CBSA, while hurricane-strain dynamics were essentially stable.

The cross-sectional and long-difference results together indicate that hazard exposure is closely tied to the level of insurance strain within metropolitan areas, and that wildfire risk in particular continued to shape the evolution of strain during 2018–2022. This is consistent with our finding in Table 12 that Hazard×Year interactions are insignificant with CBSA×year fixed effects, indicating limited steepening of the cross-sectional gradient even as wildfire-exposed locations saw differential level shifts.

F.1.2 Risk Gradient and Temporal Evolution

We next examine how the relationship between disaster risk and insurance market strain evolved over our 2018–2022 panel. Table 12 presents regressions of LISI on hazard annual frequency measures using our preferred CBSA×year fixed effects specification, with each hazard frequency standardized to z-scores.

Columns (1)–(2) use a composite measure constructed by standardizing each of the 18 NRI hazard annual frequencies to z-scores and averaging across hazards. This composite provides an objective summary of overall physical hazard exposure independent of socioeconomic factors. A one-standard-deviation increase in composite hazard frequency is associated with a 0.189-unit increase in LISI ($p<0.01$), indicating that ZIP codes with greater aggregate physical risk experience higher insurance market strain even when comparing locations within the same CBSA and year. Crucially, the Hazard×Year interaction term is not statistically significant (coefficient = -0.009, $p>0.10$), providing no evidence that the risk-strain gradient steepened during 2018–2022.

Columns (3)–(4) isolate the standardized wildfire annual frequency. A one-standard-deviation increase in wildfire annual frequency is associated with a 0.068-unit increase in LISI ($p<0.01$), reflecting the catastrophic loss potential and difficult insurability of wildfire risk. The Hazard×Year interaction (0.004) is statistically insignificant, indicating that the wildfire-strain relationship remained stable over the sample period despite major wildfire seasons in California during 2018–2020.

Columns (5)–(6) examine standardized hurricane annual frequency, which also significantly predicts LISI (coefficient = 0.066, $p<0.01$). The Hazard×Year interaction (-0.001) is again not statistically significant, indicating no steepening of the hurricane-strain gradient.[25]

The absence of significant Hazard×Year interactions across all three specifications—composite, wildfire, and hurricane—is a central finding. It indicates that the cross-sectional relationship between physical disaster risk and insurance market strain was *stable* during 2018–2022. High-risk areas entered the sample period with elevated strain relative to low-risk areas within the same CBSA, and this differential persisted

[25] In contrast, specifications using the NRI composite *risk score*—which incorporates property values, social vulnerability, and community resilience alongside physical hazard exposure—show a significant Risk×Year interaction with CBSA fixed effects. This divergence suggests that the apparent "steepening gradient" observed with the composite risk score reflects time-varying institutional responses (e.g., changes in insurer underwriting practices in high-value coastal areas, regulatory adjustments, or market structure shifts) rather than changes in the relationship between purely physical hazard exposure and insurance outcomes.

without widening. This pattern is inconsistent with narratives suggesting that climate-driven physical risks are generating accelerating divergence in insurance market outcomes. Instead, the localization of strain in high-risk areas appears to reflect a long-standing equilibrium relationship rather than an emerging crisis.

Table 12: LISI and Hazard-Specific Annual Frequency (CBSA×Year FE)

Dependent variable: Local Insurance Strain Index (LISI)						
	(1)	(2)	(3)	(4)	(5)	(6)
Hazard regressor:	Composite (18 hazards)		Wildfire only		Hurricane only	
Hazard regressor (z-score)	0.189***	0.207***	0.068***	0.060***	0.066***	0.069***
	(0.037)	(0.052)	(0.009)	(0.016)	(0.014)	(0.020)
× Year (centered at 2018)		−0.009		0.004		−0.001
		(0.030)		(0.005)		(0.009)
CBSA×Year FE	Yes	Yes	Yes	Yes	Yes	Yes
ZIP codes	23,404	23,404	23,404	23,404	23,404	23,404
ZIP code-year observations	115,233	115,233	115,233	115,233	115,233	115,233
R-squared	0.3025	0.3025	0.3055	0.3055	0.3019	0.3019
R-squared (within)	0.0030	0.0030	0.0072	0.0073	0.0021	0.0021

Notes: OLS regressions of LISI on hazard-specific annual frequency measures with CBSA×year fixed effects. Each hazard regressor is standardized to a z-score over the pooled 2018–2022 panel; the composite is the mean of all 18 hazard z-scores. Columns (2), (4), and (6) add an interaction with a linear year trend centered at 2018; the interaction coefficient is the per-year change in the hazard slope. Standard errors clustered at the CBSA×year level in parentheses. R-squared is reported both unconditionally (treating the absorbed CBSA×year fixed effects as part of the model) and within (after partialling out CBSA×year fixed effects). Significance: *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

F.2 Spatial Patterns of Individual Hazard Types

This subsection documents the heterogeneous spatial patterns of the 18 NRI natural hazard types as a function of distance from central business districts (CBDs), over the same 30-mile window and analysis sample used in the main text. Understanding these individual hazard gradients is essential for interpreting the composite physical-risk gradient documented in Section 3 and shown in Figure 1, panel (b), where the hazard types that rise with distance dominate the aggregate.

For each individual hazard, we condition on ZIP codes having non-zero exposure to that specific hazard. This is methodologically necessary because we seek to characterize the *intensity gradient*—how hazard frequency varies with distance *among locations exposed to that hazard*—rather than conflating the absence of hazard exposure with low intensity. Many ZIP codes have zero exposure to specific hazards (e.g., avalanche, coastal flooding) due to geographic constraints. Including these zeros would distort the gradient among exposed locations. In contrast, our composite measures (composite AFREQ and LISI) do not require this conditioning because every ZIP code has meaningful exposure to some subset of the 18 hazards, and the composite averages across all hazard types.

We organize the hazard types into three categories based on the economic mechanisms driving their spatial patterns.

F.2.1 City Location Selection Effects

Several hazards exhibit spatial gradients that reflect where cities historically chose to locate (Figure 7). Cities selected locations based on terrain, elevation, and proximity to navigable water—choices that systematically correlate with natural hazard exposure. Avalanche and landslide risk increase steadily with distance from CBDs, reflecting that cities avoided steep, unstable terrain, although the avalanche gradient is noisy at this resolution because only 431 exposed ZIP codes fall within the window. Wildfire risk rises steadily toward the periphery, where lower-density development meets flammable vegetation. Tornado frequency increases with distance, possibly because urban built environments disrupt formation. Hurricane risk, in contrast, is approximately flat across the 30-mile window: hurricane exposure is regional in scale, so coastal cities and their inland suburbs face similar frequencies within this range.

Figure 7: City Location Selection Effects on Hazard Exposure

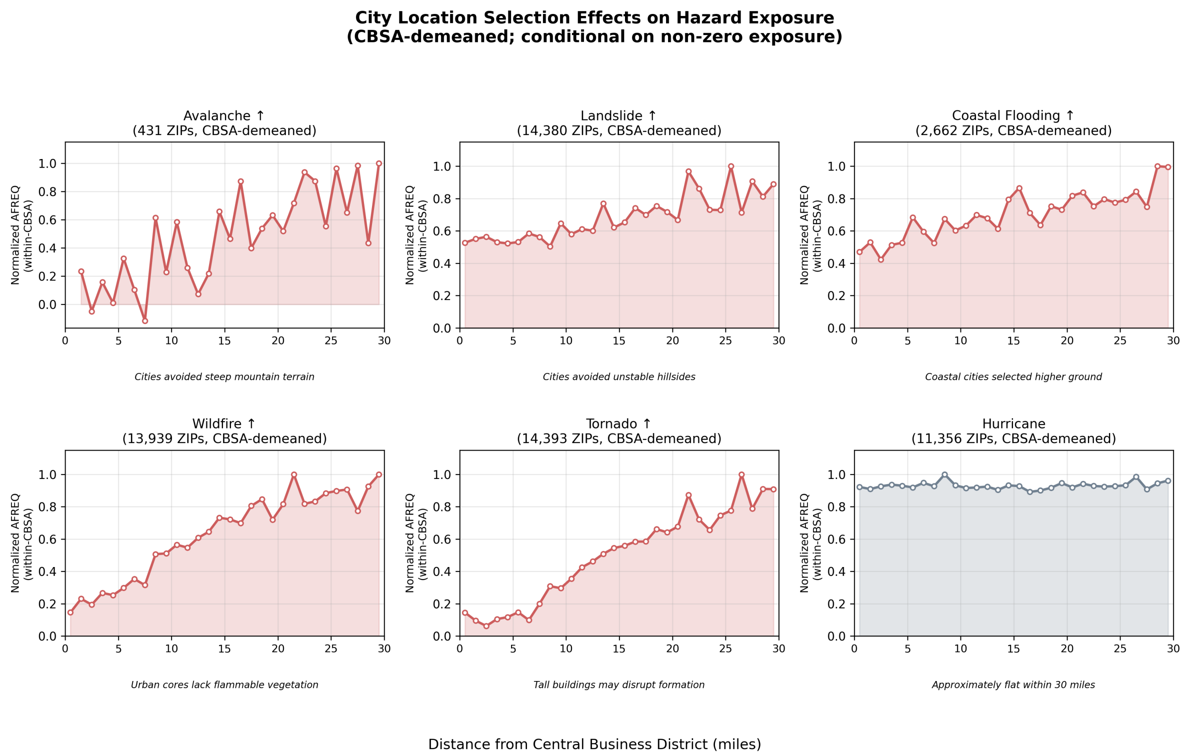

Notes: CBSA-demeaned annual hazard frequency (AFREQ) by distance from central business district, conditional on ZIP code having non-zero exposure to each hazard. Each hazard is demeaned within CBSA and rescaled by adding the grand mean ($\tilde{y}_z = (y_z - \bar{y}_m) + \bar{y}$), so that gradients reflect within-CBSA spatial variation consistent with the CBSA×year fixed effects used in the regression analysis. Values are then normalized to a 0–1 scale within each hazard type for comparability. Distance is binned in 1-mile intervals. Restricted to CBSAs with at least three exposed ZIP codes. Based on ZIP codes within 30 miles of CBSA centers, 2022 cross-section.

F.2.2 Urban Microclimate Effects

Temperature-related hazards exhibit gradients driven by the urban heat island effect (Figure 8). Dense urban areas generate and retain heat from buildings, vehicles, and reduced vegetation, modifying local climate. Within the 30-mile window these gradients are modest but consistently signed: cold wave and winter weather hazards rise gently with distance from CBDs because urban heat reduces the frequency of cold extremes, while heat wave frequency is higher in the inner half of the window and declines toward 30 miles because the urban heat island *intensifies* heat extremes.

F.2.3 Rural Transition Effects

The hazards in this group are tied to the transition from built to rural land use, which for most metropolitan areas lies at or beyond the edge of the 30-mile window, so their within-window gradients are correspondingly muted (Figure 9). Drought rises gently toward the periphery, reflecting the approach of agricultural land

Figure 8: Urban Microclimate Effects on Hazard Exposure

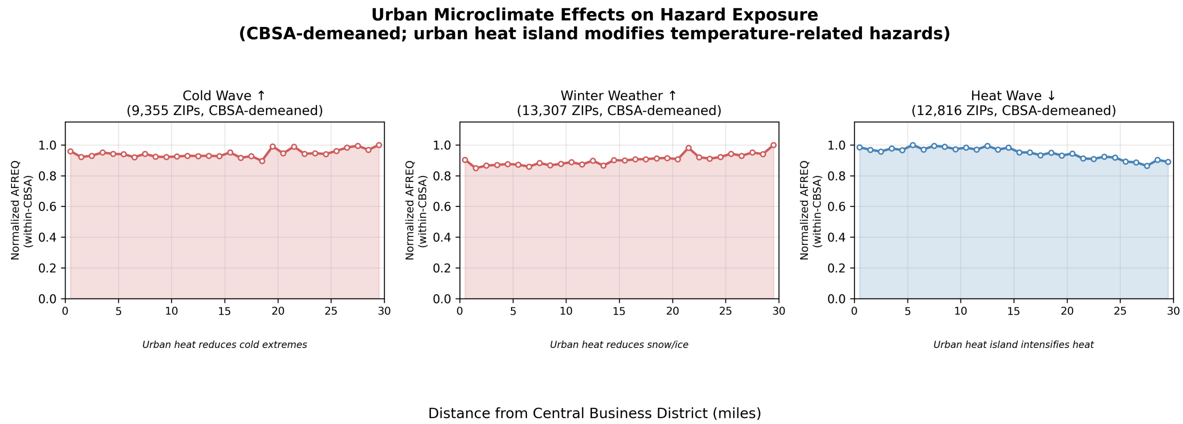

Notes: CBSA-demeaned annual hazard frequency (AFREQ) by distance from central business district, conditional on ZIP code having non-zero exposure to each hazard. Each hazard is demeaned within CBSA and rescaled by adding the grand mean ($\tilde{y}_z = (y_z - \bar{y}_m) + \bar{y}$), then normalized to a 0–1 scale within each hazard type. Urban heat island effects reduce cold-related hazards but intensify heat-related hazards. Distance is binned in 1-mile intervals. Restricted to CBSAs with at least three exposed ZIP codes. Based on ZIP codes within 30 miles of CBSA centers, 2022 cross-section.

use. Hail, ice storm, and strong wind are approximately flat, and lightning declines mildly with distance. Riverine flooding is the exception within the window, displaying a non-monotonic pattern that peaks in the 10–15-mile band as impervious urban surfaces give way to floodplains and then declines. The sharper rural-transition effects visible in wider windows fall outside the urban-suburban range this paper studies.

Figure 9: Rural Transition Effects on Hazard Exposure

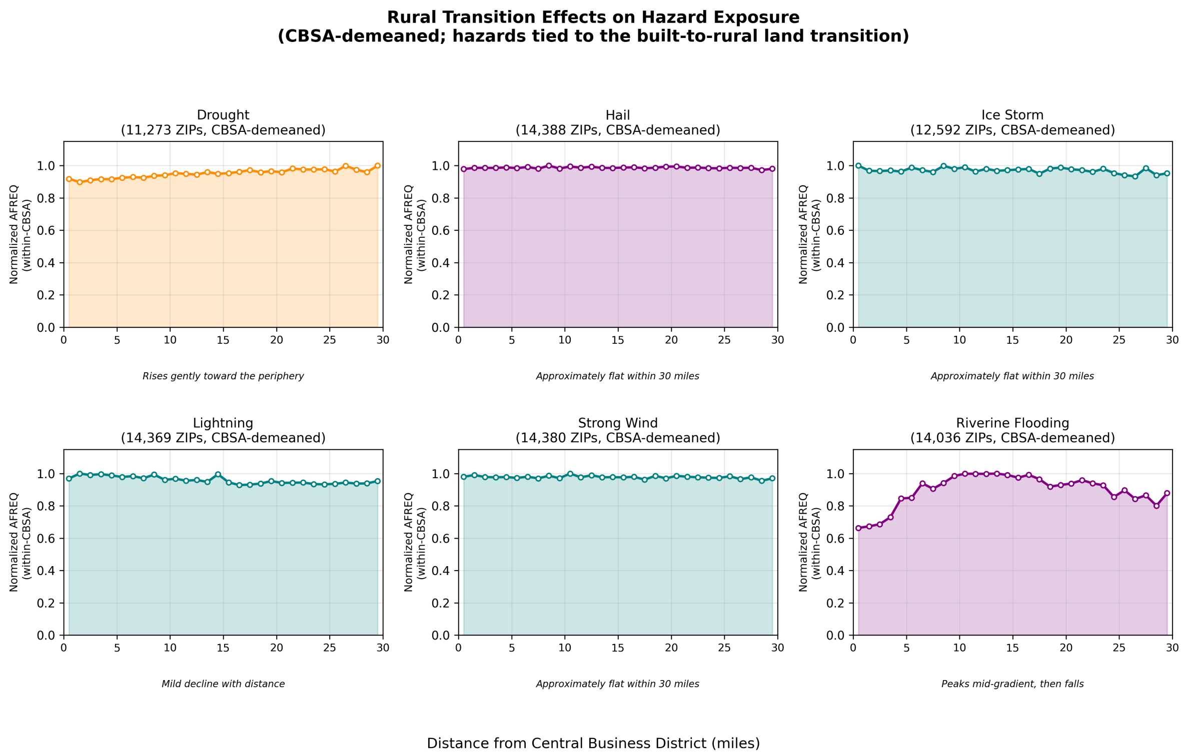

Notes: CBSA-demeaned annual hazard frequency (AFREQ) by distance from central business district, conditional on ZIP code having non-zero exposure to each hazard. Each hazard is demeaned within CBSA and rescaled by adding the grand mean ($\tilde{y}_z = (y_z - \bar{y}_m) + \bar{y}$), then normalized to a 0–1 scale within each hazard type. Distance is binned in 1-mile intervals. Restricted to CBSAs with at least three exposed ZIP codes. Based on ZIP codes within 30 miles of CBSA centers, 2022 cross-section.

G Population-Weighted Estimates

Our main analysis uses unweighted regressions that treat each ZIP code equally. This approach may be appropriate for understanding geographic patterns of insurance strain but may not accurately characterize the experience of the typical American household. As a robustness check, we re-estimate our central specifications using ZIP code population as analytic weights. We report the two regressions and the one descriptive gradient that bear most directly on the paper's claims—the hazard-to-strain relationship, the housing-market capitalization result, and the strain-versus-risk spatial gradient—rather than reproducing every unweighted exercise under weights.

G.1 LISI and Hazard Frequency (Population-Weighted)

Table 13 presents population-weighted panel regressions of LISI on hazard frequency measures with CBSA×year fixed effects, matching the sample and specification of Table 12 apart from the weights.

Table 13: LISI and Hazard Frequency: Population-Weighted

	Composite (18 Hazards)		Wildfire		Hurricane	
	(1) Baseline	(2) + Trend	(3) Baseline	(4) + Trend	(5) Baseline	(6) + Trend
Hazard AFREQ (z)	0.217***	0.220***	0.067***	0.057***	0.082***	0.084***
	(0.039)	(0.054)	(0.010)	(0.017)	(0.015)	(0.022)
Hazard × Year		−0.001		0.005		−0.001
		(0.030)		(0.006)		(0.010)
CBSA×Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Population Weights	Yes	Yes	Yes	Yes	Yes	Yes
Within R-squared	0.0041	0.0041	0.0083	0.0084	0.0032	0.0032
Observations	115,233	115,233	115,233	115,233	115,233	115,233

Notes: Population-weighted OLS regressions of LISI on NRI hazard annual frequency measures with CBSA×year fixed effects. Weights are ZIP code population. Hazard AFREQ measures are standardized to z-scores. Composite AFREQ is the average of 18 hazard z-scores. The Hazard×Year interaction tests whether the risk-strain gradient steepens over time (year normalized to 0 in 2018). Standard errors clustered at the CBSA×year level in parentheses. Significance: *** p<0.01, ** p<0.05, * p<0.10.

The weighted estimates closely track the unweighted ones: the composite hazard coefficient is 0.217 against 0.189 unweighted, wildfire 0.067 against 0.068, and hurricane 0.082 against 0.066, all significant at the 1% level. As in the unweighted specification, none of the hazard-by-year interactions is statistically significant, so there is no evidence that the risk-to-strain gradient steepened over the panel when ZIP codes are weighted by population.

G.2 Housing Market Outcomes (Population-Weighted)

Table 14 presents population-weighted housing outcome regressions.

The specification is identical to Table 2 in every respect except the weights: the 30-mile cap and sample restrictions match exactly, both tables run on the same 64,270 and 33,770 observations, and the comparison therefore isolates the effect of weighting alone. All four outcomes load negatively at the 1% level: log home prices in the full sample (-0.076 weighted versus -0.076 unweighted), log rents (-0.040 versus -0.045), log home prices in the common sample (-0.120 versus -0.154), and the log price-to-rent ratio (-0.080 versus -0.109). Weighting leaves the full-sample home-price estimate essentially unchanged and moderately attenuates the common-sample estimates, while every outcome retains its sign and 1% significance. The

Table 14: Insurance Strain and Housing Market Outcomes: Population-Weighted

	(1) Log Home Prices	(2) Log Rents	(3) Log Prices	(4) Log Price-to-Rent
LISI	−0.076***	−0.040***	−0.120***	−0.080***
	(0.017)	(0.010)	(0.028)	(0.019)
CBSA×Year FE	Yes	Yes	Yes	Yes
Population Weights	Yes	Yes	Yes	Yes
ZIP Codes	13,098	9,153	9,153	9,153
Observations	64,270	33,770	33,770	33,770
R-squared (within)	0.0119	0.0050	0.0219	0.0132

Notes: Population-weighted OLS regressions of housing market outcomes on the new primary LISI with CBSA×year fixed effects. Weights are ZIP code population. LISI is the equal-weighted average of three robust-standardized pillars (Unaffordability, Unavailability, Unprofitability). Column (1) uses the full sample; Columns (2)–(4) use a common sample restricted to ZIP codes with non-missing ZHVI and RentHub core multifamily rent data. Standard errors clustered at the CBSA×year level. Significance: *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

capitalization result is not an artifact of small or unrepresentative ZIP codes carrying the same weight as populous ones.

G.3 Insurance Strain and Physical Risk Gradients (Population-Weighted)

Figure 10 presents the population-weighted analog to Figure 1 in the main text. Population weighting is the *only* difference between the two figures: both are drawn by the same routine on the same 67,170 ZIP-year observations (13,591 ZIP codes across 324 CBSAs), with the same CBSA×year demeaning, the same 1-mile bins over 0–30 miles, and the same requirement of at least three ZIP codes per CBSA×year cell. Every ZIP code in that sample has a non-missing, strictly positive population, so weighting drops no observations and the two figures are estimated on an identical sample. Weights enter both the demeaning and the bin means, matching the within-transformation implied by weighted least squares.

Weighting leaves both gradients essentially unchanged. The LISI gradient is, if anything, marginally steeper: the CBD-to-outer-ring gap is $+0.216$ under population weighting against $+0.212$ unweighted, with the innermost bin falling slightly from $+0.281$ to $+0.267$ and the outermost from $+0.069$ to $+0.051$. The physical-risk gradient likewise retains its sign and magnitude, with a CBD-to-outer gap of -0.102 weighted against -0.099 unweighted. The divergence between insurance strain and physical risk documented in Section 3 is therefore not an artifact of treating small and large ZIP codes equally.

Figure 10: Insurance Strain (LISI) vs. Physical Risk by CBD Distance (Population-Weighted)

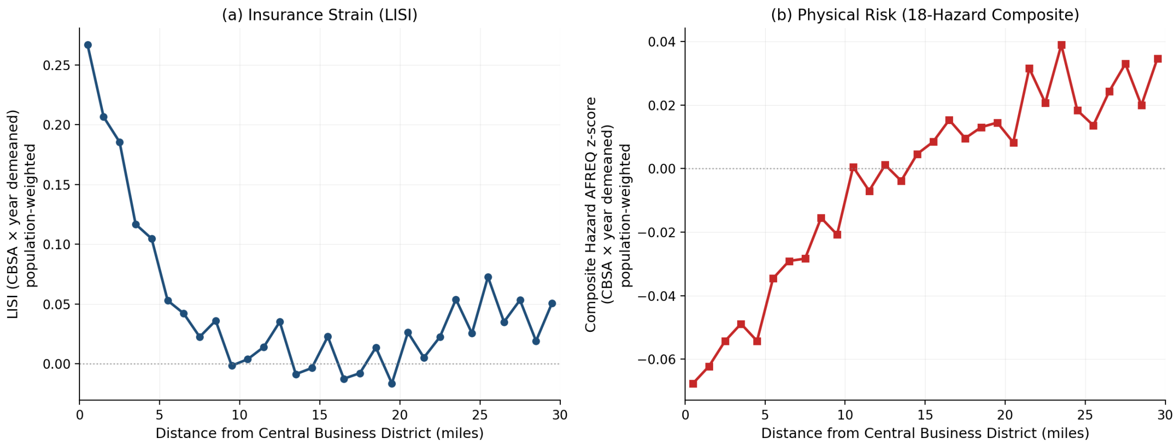

Notes: Population-weighted LISI composite (Panel a) and composite physical hazard (Panel b) across 1-mile distance bins from 0 to 30 miles from the nearest metropolitan CBD, pooled over 2018–2022. This figure is the population-weighted counterpart to Figure 1 and is produced by the same routine on the same sample; weighting is the only difference. Values are CBSA×year demeaned, $\tilde{y}_{z,t} = (y_{z,t} - \bar{y}_{m,t}) + \bar{y}$, with both the cell means and the grand mean population-weighted, matching the within-transformation implied by weighted least squares; bin means are population-weighted likewise. The composite hazard measure standardizes each of the 18 NRI annual frequency measures to z-scores and averages across them. Restricted to CBSA×year cells with at least three ZIP codes. Based on 67,170 ZIP-year observations across 13,591 ZIP codes and 324 CBSAs, identical to Figure 1: every ZIP code in that sample has strictly positive population, so weighting drops no observations.

H Renter Household Composition Across the Urban–Exurban Gradient

Figure 11 documents the national spatial distribution of renter households across the urban–exurban gradient. Panel (a) plots total renter-occupied housing units by distance from the central business district for all ZIP codes within 80 miles of a metropolitan CBD. Renter households are overwhelmingly concentrated near metropolitan cores: approximately 19.9 million renter-occupied units lie within 10 miles of a CBD, declining steeply through the suburban ring and to roughly 264,000 in the 70–80 mile bin. Panel (b) plots the county-demeaned renter household rate—the share of occupied housing units that are renter-occupied, net of county-level composition. The renter rate is highest in urban cores (approximately 32%) and falls sharply to a trough of roughly 22% in inner suburban areas (10–20 miles), consistent with the concentration of owner-occupied single-family housing in inner-ring suburbs. Beyond 20 miles, the rate is essentially flat at approximately 24–26%, indicating that the urban core–suburban distinction is the dominant spatial margin for tenure composition.

Figure 12 replicates the same analysis restricted to the RentHub common sample used for housing outcomes analysis (ZIP codes with both ZHVI and RentHub core multifamily rent data). Panel (a) shows a similar count gradient, with 19.5 million renter-occupied units within 10 miles. Panel (b) shows the county-demeaned renter rate in the common sample is higher overall (grand mean 33% vs. 26% nationally), reflecting the common sample's selection toward ZIP codes with sufficient multifamily rental activity to appear in RentHub. The urban core premium is preserved (38% at 0–10 miles), though estimates beyond 40 miles are based on substantially fewer ZIP codes and should be interpreted with caution.

Figure 11: Renter Household Composition by Distance from Central Business District: National

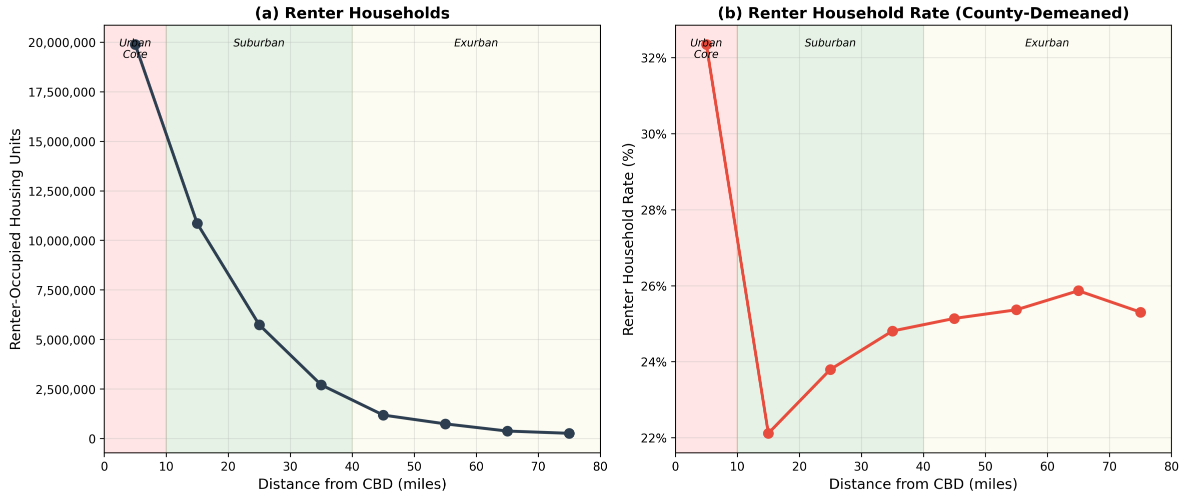

Notes: Panel (a) shows the total number of renter-occupied housing units summed across ZIP codes within each 10-mile distance bin from the nearest metropolitan central business district. Panel (b) shows the mean renter household rate (share of occupied housing units that are renter-occupied), county-demeaned to remove cross-county composition effects: $y_{\text{plot}} = (y_i - \bar{y}_c) + \bar{y}$, where $\bar{y}_c$ is the county mean and $\bar{y}$ is the grand mean. The sample includes all ZIP codes within 80 miles of a metropolitan CBD with available ACS tenure data (23,971 ZIP codes in 2022). Panel (b) further restricts to counties with at least three ZIP codes. Renter household rates follow the Census Bureau methodology: renter-occupied units divided by total occupied housing units (vacant units excluded). Shaded regions indicate urban core (0–10 miles), suburban (10–40 miles), and exurban (40–80 miles) zones. Source: American Community Survey 5-Year Estimates (2023), Table B25003.

Figure 12: Renter Household Composition by Distance from Central Business District: Common Sample

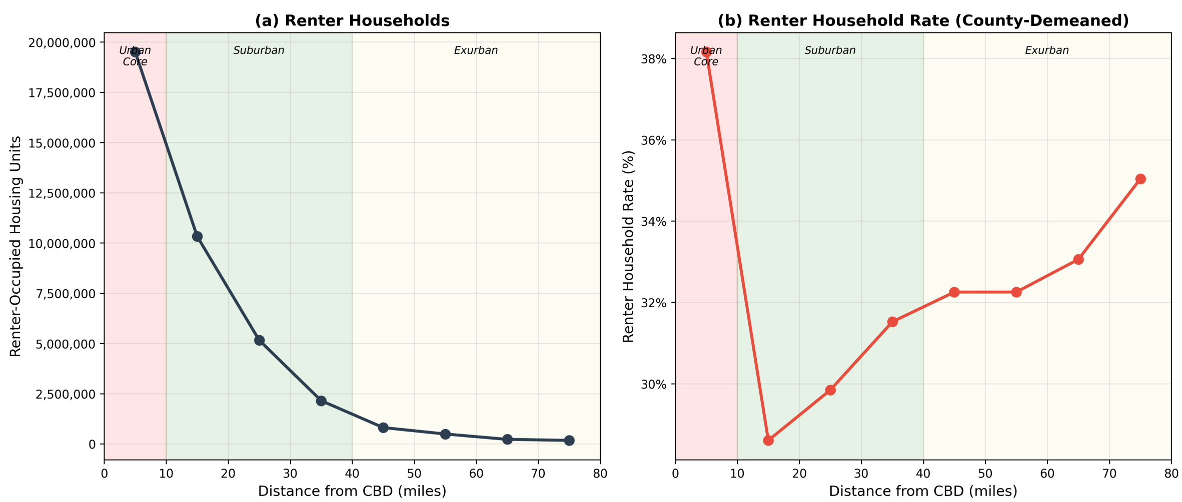

Notes: Panel (a) shows the total number of renter-occupied housing units summed across ZIP codes within each 10-mile distance bin from the nearest metropolitan central business district. Panel (b) shows the mean renter household rate (share of occupied housing units that are renter-occupied), county-demeaned to remove cross-county composition effects: $y_{\text{plot}} = (y_i - \bar{y}_c) + \bar{y}$, where $\bar{y}_c$ is the county mean and $\bar{y}$ is the grand mean. Both panels restrict the sample to ZIP codes in the RentHub common sample (those with non-missing Zillow Home Value Index and RentHub core multifamily rent data in 2022; 11,954 ZIP codes). Panel (b) further restricts to counties with at least three ZIP codes in the common sample (10,622 ZIP codes). Renter household rates follow the Census Bureau methodology: renter-occupied units divided by total occupied housing units (vacant units excluded). Shaded regions indicate urban core (0–10 miles), suburban (10–40 miles), and exurban (40–80 miles) zones. Source: American Community Survey 5-Year Estimates (2023), Table B25003.

Hutchins Center
on Fiscal & Monetary Policy
at BROOKINGS

The mission of the Hutchins Center on Fiscal and Monetary Policy is to improve the quality and efficacy of fiscal and monetary policies and public understanding of them.

Questions about the research? Email ESmedia@brookings.edu.
Be sure to include the title of this paper in your inquiry.