Priced Out of Safety? Insurance Reform, Residential Sorting, and Climate Gentrification [1]

Sherry Tang [2]

April 2026

Very preliminary: Please do not circulate

Abstract

Flood insurance premiums in the United States have historically been set at roughly half their actuarial cost, implicitly subsidizing the most flood-exposed properties and distorting housing markets. This paper studies whether FEMA's Risk Rating 2.0 (the largest reform to NFIP pricing in 55 years) has begun to correct this distortion and reshuffle residential sorting by income and race. Replacing community-averaged rates with property-specific actuarial premiums, RR2.0 generates plausibly exogenous variation in insurance cost shocks across properties within the same housing market. I exploit this variation in a continuous-treatment difference-in-differences design using FEMA's pre-release projected premium change files and 5.7 million CoreLogic housing transactions from 2018 to 2024. The main result shows that higher flood insurance cost exposure significantly reduces house prices: a one-standard-deviation increase in local treatment intensity lowers log sale prices by 1.6%, or approximately $3,900 at the median transaction. A direction-reversal placebo using the share of policies receiving premium decreases yields a positive effect, confirming that price dynamics reflect genuine insurance cost capitalization rather than confounding trends. Ongoing work extends this analysis to examine whether these price pressures reshape the income and racial composition of buyers in flood-exposed neighborhoods and will complement the reduced-form evidence with a dynamic residential sorting model to quantify welfare consequences across income and racial groups and evaluate counterfactual policy reform.

Keywords: Flood Insurance, Risk Rating 2.0, Housing Markets, Residential Sorting, Climate Gentrification

[1] I am very grateful to Charles Taylor for invaluable support and encouragement. I thank Koichiro Ito, Peter Tufano, Wolfram Schlenker, Seung Min Kim, David Jesus Martinez, Nathanaël Soulage, and participants at Harvard-MIT Joint Econ Predoc Seminar, Harvard Financial Economics of Climate and Sustainability Reading Group and Harvard Salata Environmental Economics Lab Meeting for helpful suggestions and comments. Any opinions and conclusions expressed herein are my own and do not necessarily represent the views of Harvard University or its staff. All results have been reviewed to ensure that no confidential information is disclosed. All errors are my own.

[2] Predoctoral Fellow, Harvard University, sherrytang@fas.harvard.edu

1 Introduction

Flood insurance premiums in the United States have long been set at a fraction of their actuarial cost. From the NFIP's founding in 1968 through 2021, premiums were determined by community-averaged flood zones and maps that were often decades out of date, resulting in implicit subsidies of roughly 50% on the average policy. This pricing arrangement understated the true cost of occupying high-risk housing and inflated flood-zone property values. Gourevitch et al. (2023) estimate aggregate residential overvaluation of $121–$237 billion attributable to the gap between market prices and actuarial flood risk, while Georgic and Klaiber (2022) estimate that NFIP subsidy eligibility alone is capitalized at $12,352 per property on average. In October 2021, FEMA launched Risk Rating 2.0 (RR2.0), the most consequential reform to NFIP pricing in the program's 55-year history. By replacing community-averaged rates with property-specific actuarial premiums derived from each building's individual flood exposure, RR2.0 began unwinding this cross-subsidy on a glide path governed by an 18%-per-year statutory cap on annual premium increases. The result is a nationwide quasi-experiment in the repricing of climate risk: two neighbors within the same flood zone may now face very different annual premiums, with the higher-elevation property seeing a decrease while the flood-exposed neighbor faces increases that, in some cases, will ultimately exceed 300%.

This paper asks whether this differential repricing generates residential sorting, that is, systematic differences in who moves into and out of neighborhoods with different flood exposure, and, if so, whether the sorting constitutes "climate gentrification": lower-income and minority households displaced from newly expensive safe areas as demand shifts away from high-premium, high-risk neighborhoods. I address this question using two complementary strategies. The first is a continuous-treatment difference-in-differences design. The treatment intensity measure Δl_z (the share of single-family NFIP policies in ZIP code z projected to face monthly premium increases exceeding $10) is drawn from FEMA's pre-implementation projected change files released in September 2021, strictly before Phase II renewals began. This pre-determination ensures Δl_z cannot be contaminated by anticipatory responses to post-reform housing market outcomes. I match this measure to CoreLogic deed records to estimate the causal effect of treatment intensity on housing prices, transaction volumes, and the income and racial composition of inmovers, using property fixed effects and county-by-quarter fixed effects to control for time-invariant property characteristics and local economic shocks. The event-study design spans sixteen quarters (2018–2025), with April 2022 Phase II implementation as the event date. The second strategy is a dynamic residential sorting model in the tradition of Bayer et al. (2016), extended with heterogeneous flood risk beliefs, which translates the reduced-form price and sorting responses into welfare consequences across household types and enables counterfactual policy simulations.

Preliminary evidence from the insurance data establishes that the reform generates substantial and well-identified treatment variation. The median NFIP single-family policy carries a full-risk target premium of approximately $1,290 per year under RR2.0, compared to the current median of $786—a nationwide implicit subsidy of roughly $504 per policy per year. This subsidy varies dramatically by state: New Jersey (51%), New York

(49%), and Louisiana (47%) carry the largest proportional gaps. Across 2,786 counties, the share of NFIP policies projected to face monthly increases exceeding $10 averages 7% but has a pronounced right tail. The reform's aggregate cost impact is corroborated by the Keys and Mulder (2024) mortgage escrow data, which show a 30.5% nominal increase in weighted-mean property insurance premiums from 2020 to 2023, with the acceleration concentrated in the two years immediately following RR2.0 Phase I. The OpenFEMA policy data confirm adoption was complete by 2023.

My main result is that higher flood insurance cost exposure significantly reduces house prices: a one-standard-deviation increase in local treatment intensity (11 percentage points more NFIP policies facing premium increases exceeding $10/month) lowers log sale prices by 1.6%, or approximately $3,900 at the median transaction. Event-study estimates show flat pre-trends through 2021Q3 and price effects that deepen progressively over six to seven post-reform quarters, which is consistent with the program's 18%/year statutory glide path phasing out legacy subsidies. A direction-reversal placebo using the share of policies receiving premium decreases yields a positive and significant coefficient, confirming that price dynamics reflect genuine insurance cost capitalization rather than confounding trends. Ongoing work extends this analysis to examine whether these price pressures are already reshaping the income and racial composition of buyers in flood-exposed neighborhoods—a form of insurance-driven climate gentrification—and will complement the reduced-form evidence with a dynamic residential sorting model to quantify welfare consequences across income and racial groups and evaluate counterfactual policy reforms including means-tested premium caps and flood adaptation investments.

Related literature. This paper connects two literatures. The first is the growing body of evidence on flood insurance, climate risk, and housing markets. Several studies document that flood risk is capitalized into housing prices—Bernstein et al. (2019) find a 7% discount for sea-level-rise exposure, and Hino and Burke (2021) find a 2.1% price decline upon assignment to a SFHA—but that subsidized insurance premiums have partially masked this risk signal, contributing to overvaluation (Gourevitch et al., 2023). Keys and Mulder (2024) show directly that insurance cost increases reduce housing values and transaction volumes. The closest empirical example is Fabian (2024), who uses RR2.0 to study supply-side housing development, finding a 27% reduction in new permits in North Carolina, but does not study demand-side price and sorting effects. Ortega and Petkov (2024) and Mulder and Kousky (2023) document that RR2.0 reduced NFIP enrollment with lapses largest in the lowest-income ZIP codes, but do not estimate housing market or welfare consequences. Gruber and Solomon (2026) characterize the optimal flood insurance pricing rule in a second-best environment, providing the theoretical welfare benchmark against which the reform's distributional outcomes should be evaluated. To the best of my knowledge, no existing paper has used RR2.0 as a quasi-experiment to estimate causal housing price effects or the distributional sorting consequences.

The second literature is residential sorting with environmental amenities. Following Tiebout (1956), households reveal preferences for local attributes through location choices; when RR2.0 raises the effective cost of high-risk neighborhoods, households sort toward

safety in ways that depend on their income and risk preferences. Bakkensen and Ma (2020) find that lower-income and minority households disproportionately occupy high-flood-risk neighborhoods, suggesting they are most exposed to large premium increases yet least able to absorb them—a precondition for climate gentrification. Bakkensen and Barrage (2022) show that belief heterogeneity already drives sorting into coastal flood zones, amplifying the potential price and demographic effects of actuarial repricing. In environmental economics, Banzhaf and Walsh (2008) demonstrate that neighborhoods experiencing reductions in environmental hazard attract higher-income households and see rising property values, potentially displacing original lower-income residents—the exact mechanism this paper tests in the flood risk context. Keenan et al. (2018) supply the first empirical evidence for climate gentrification specifically, documenting that higher-elevation Miami-Dade properties appreciated faster after climate concerns intensified. This paper examines whether RR2.0 activates the same insurance cost channel at national scale.

Contributions. My contribution is threefold. First, I provide the first quasi-experimental estimates of RR2.0's demand-side housing market effects, filling a gap left by the supply-side analysis of Fabian (2024) and the insurance demand analysis of Ortega and Petkov (2024). Second, I directly test the climate gentrification mechanism by examining whether premium shocks drive sorting by income and race, using buyer composition from CoreLogic deed records as the outcome. To my knowledge, very few papers use a causal quasi-experimental design to test whether climate-related cost shocks generate residential sorting by race and income at national scale. Third, by embedding the reduced-form estimates in a structural sorting model, I go beyond documenting effects to quantifying welfare consequences and evaluating which policy interventions most efficiently address the distributional costs of actuarially fair flood pricing.

The remainder of the paper is organized as follows. Section 2 describes the background. Section 3 describes the data. Section 4 presents the empirical strategy. Section 5 lays out the structural sorting model. Section 6 presents preliminary results. Section 7 describes the next steps.

2 Background

2.1 The National Flood Insurance Program

Flood insurance in the United States is provided almost exclusively by the federal government. The National Flood Insurance Program (NFIP), launched in 1968 and administered by FEMA, carries approximately 4.7 million policies in force covering over $1.28 trillion in residential building and contents value, representing more than 90% of all residential flood insurance in the country.[3] Despite this scale, NFIP participation is far from univer-

[3]The private flood market has grown in recent years, reaching roughly 350,000 policies in 2023, but remains small relative to the NFIP.

sal: around one in twenty U.S. households holds a flood insurance, though penetration reaches 25% in some high-exposure coastal markets.

The NFIP issues a single product, the Standard Flood Insurance Policy (SFIP), which covers up to $250,000 for structural damage and $100,000 for contents. Policies are written and claims administered by "Write-Your-Own" private carriers, but all contracts and risk are ultimately borne by the federal program. The NFIP carries $36.5 billion in cumulative borrowing from the U.S. Treasury since 2005, a direct consequence of the subsidized pricing described below.

Insurance purchase is mandatory for properties in Special Flood Hazard Areas (SFHAs) [4] that carry a federally backed mortgage. SFHAs encompass roughly 6–7% of all U.S. homes; inside these zones the insurance mandate creates a direct link between flood insurance premiums and effective housing costs for leveraged buyers. Outside SFHAs, insurance purchase is voluntary, yet 17.9% of the NFIP portfolio (based on the OpenFEMA policy sample) consists of pre-FIRM properties that have historically received the largest subsidies and face the steepest increases under the reform.

2.2 The Legacy Pricing Regime and Its Distortions

Prior to Risk Rating 2.0, NFIP premiums were set using a system of geographic flood zones and structural indicators established when the community's Flood Insurance Rate Map (FIRM) was originally drawn, often decades ago. All properties within the same flood zone paid rates determined by community-averaged risk rather than their individual exposure. As a result, two adjacent properties with identical structures but meaningfully different elevations paid the same premium. This community-averaging implicitly cross-subsidized high-risk properties at the expense of lower-risk ones within the same zone. By FEMA's own estimate, the average NFIP policy received a subsidy of approximately 50% relative to its actuarial cost, with the largest subsidies accruing to the highest-risk, lowest-elevation properties.

These pricing distortions had direct implications for residential location choices. To the extent that insurance costs are capitalized into housing prices, community-averaged premiums understated the true cost of occupying high-risk neighborhoods and overstated the cost of safe ones. Georgic and Klaiber (2022) estimate that NFIP subsidy eligibility is capitalized at approximately $12,352 per property on average, consistent with the view that subsidized premiums inflated demand for flood-exposed housing. Gourevitch et al. (2023) find aggregate residential overvaluation of $121–$237 billion attributable to the gap between market prices and actuarial flood risk. The reform is thus not merely an insurance pricing correction but a potential shock to the implicit relative price of safe versus risky housing, which motivates the empirical analysis in this paper.

[4] Zones defined by FEMA as having at least a 1% annual flood probability.

2.3 Risk Rating 2.0: Reform Design and Identifying Variation

On October 1, 2021, FEMA replaced the zone-based legacy system with Risk Rating 2.0, which assigns each property a premium based on its individual physical flood risk. The new methodology incorporates the likelihood of multiple flood perils (inland flooding, storm surge, coastal erosion), building characteristics (foundation type, first-floor elevation, replacement cost value), and updated flood loss statistics. The result is that two properties in the same FIRM flood zone can now face very different premiums depending on elevation above the base flood elevation, distance from water sources, and structural features.

The reform rolled out in two phases. Phase I (October 1, 2021) applied RR2.0 to all new policies and renewals qualifying for immediate decreases. Phase II (April 1, 2022) extended RR2.0 to all remaining renewals, completing the transition by April 1, 2023. FEMA's analysis projected that 23% of policyholders would see immediate decreases (averaging \$86/month) and 66% would see increases of \$0–\$10/month; the remaining 11% face increases of more than \$10/month, with a 4% tail facing increases exceeding \$20/month. My own calculations from the FEMA county-level projected change files confirm this distribution across 2,786 counties, with substantial right-tail variation: at the county level, the mean share of policies facing >\$10/month increases is 7.0%, but this figure reaches 53% in Charlotte County, FL.

To protect policyholders from large immediate increases, the Homeowner Flood Insurance Affordability Act of 2014 (HFIAA) caps annual premium increases at 18% for primary residences and 25% for other properties. This cap creates a multi-year "glide path" toward full actuarial pricing. As of December 2022, the GAO reports that the median risk-based target premium was \$1,288 while the median premium actually paid was \$689—roughly half the actuarial rate. An estimated 9% of policyholders will eventually require increases exceeding 300%, and the GAO projects it will take until approximately 2037 for 95% of policyholders to reach full-risk pricing.

Three features of the reform are critical for identification. First, FEMA published the ZIP-code-level projected premium change distributions in September 2021, before Phase II renewals began. This pre-announcement means the treatment intensity measure Δl_z (the share of policies in ZIP z projected to face >\$10/month increases) is strictly predetermined with respect to post-reform housing market outcomes. Second, the reform changed only prices, not insurance access, SFHA boundaries, or mortgage regulations, which isolates the insurance cost channel from confounding institutional changes. Third, because the within-ZIP variation in premium changes is driven by property-level physical characteristics (elevation, foundation type) rather than by housing market demand, conditional on ZIP fixed effects it is plausibly orthogonal to latent trends in neighborhood desirability.

3 Data

3.1 Treatment Variable: FEMA RR2.0 Projected Premium Changes

The treatment intensity measure Δl_z is constructed from two publicly released FEMA files. The *county-level projected premium change profiles* (released September 2021) report, for each of 2,786 U.S. counties, the share of single-family NFIP policies projected to fall in each of 22 monthly change bins, ranging from $<-\$100$ to $>\$100$ per month. The analogous *ZIP-code-level files* (released April 2025, reflecting the current policy portfolio) cover 17,342 ZIP codes with at least five active single-family policies. Both files are constructed by FEMA from the same actuarial model used to implement RR2.0 and are strictly pre-determined relative to post-reform housing market outcomes.

From these files I construct two treatment intensity measures at the ZIP level: (i) the share of NFIP policies projected to face a monthly increase exceeding \$10 (the primary continuous treatment variable Δl_z in equation (1)), and (ii) the share projected to receive any decrease, which serves as a falsification treatment for a direction-reversal test. The 2021 county-level files are used for the main DiD specification, since they were published before Phase II began; the 2025 ZIP files are used to study the current steady-state distribution of premium burdens and to construct treatment intensity for the structural model.

3.2 NFIP Policy Transactions

The FIMA NFIP Redacted Policies dataset (v2, OpenFEMA) contains approximately 72.6 million individual policy transactions going back to 2009, updated monthly. The key variable for validating the treatment timing is `rateMethod`, which takes the value `"RatingEngine"` for all policies rated under RR2.0; I confirm that this share rose from zero in 2020 to 99.5% by 2023, consistent with the Phase I/II rollout. Additional variables used in the analysis include the total insurance premium, subsidized rate type (pre-FIRM subsidy indicator), flood zone, primary residence indicator, occupancy type, and building replacement cost value. Geographic identifiers at the ZIP, county, flood-region, and census tract levels enable merging with housing transaction and demographic data. This dataset is also used to document the insurance demand response (policy lapses and new enrollment) as a complement to the housing market analysis.

3.3 Housing Transactions

The primary outcome data are CoreLogic deed records, which cover approximately 95% of U.S. residential property transactions. The dataset provides the sale price, transaction date, parcel identifier, property address, and buyer and seller names for each arm's-length transaction. Parcel identifiers allow construction of a repeat-sales panel: for each

property that transacts multiple times, I observe the price at each sale, enabling the property fixed effects μ_i in equation (1) to absorb all time-invariant unobservables at the property and immediate neighborhood level. The transaction date is used to define the pre/post indicator Post_t at the quarterly frequency.

The raw CoreLogic file contains 364 million residential records. After restricting to the study period (2018Q1–2024Q1), retaining only arm's-length transactions (sale prices between \$5,000 and \$50 million), dropping records with missing property identifiers or ZIP codes, and inner-joining to the NFIP ZIP treatment intensity file to restrict to ZIP codes with active NFIP policies, the panel contains 8,563,362 transactions from 6,771,758 unique residential properties across 16,617 ZIP codes and 49 states. Non-disclosure states (where sale prices are legally suppressed) account for 9.2% of transactions and are excluded from the main regressions. Corporate buyers, identified by CoreLogic's corporate indicator flag, represent 19.5% of transactions and are also excluded from the main estimation sample.

Buyer income and race for the distributional sorting analysis are imputed from CoreLogic mortgage origination records linked to Home Mortgage Disclosure Act (HMDA) filings, which report borrower income and race/ethnicity for federally regulated mortgage applications. For cash transactions and cases without HMDA matches, I supplement with ZIP-quarterly ACS estimates of in-mover demographics as an alternative outcome.

3.4 Neighborhood Demographics

Neighborhood demographic composition, such as income distribution, racial composition, homeownership rate, share of housing units in SFHA, is drawn from the American Community Survey 5-year estimates at the ZIP Code Tabulation Area (ZCTA) level. Pre-reform (2016–2020) ACS estimates serve as baseline neighborhood controls $\mathbf{X}_{it}$ in equation (1) and as the initial conditions for the structural sorting model. Post-reform (2021–2023) estimates are used as outcomes in the ZIP-quarterly regressions studying the demographic composition of in-movers.

3.5 Flood Risk

Objective flood risk at the property level is measured using First Street Foundation Property Level Flood Risk Statistics (v2.0), which provides estimates of annual expected flood damage as a share of property value for each of approximately 142 million U.S. properties. This measure serves two roles: (i) as the flood risk variable ρ_j in the structural utility equation (3), and (ii) as an alternative treatment instrument to validate the FEMA-based Δl_z measure, under the assumption that larger RR2.0 increases should accrue to properties with higher objective first-party flood risk. FEMA Flood Insurance Rate Map (FIRM) flood zone designations (SFHA vs. non-SFHA) are used to stratify the main estimation by mandatory- versus voluntary-purchase status.

4 Empirical Strategy

The core identifying variation is the property-level flood insurance premium shock generated by Risk Rating 2.0. Because FEMA's new pricing algorithm assigns each property a risk-based premium based on physical flood exposure, two otherwise similar properties within the same neighborhood can receive very different premium changes depending on elevation, distance to water, and foundation type. Crucially, this within-neighborhood variation in Δl is orthogonal to the local amenity and demographic trends that drive housing demand, conditional on neighborhood-level fixed effects. The reform also introduced no simultaneous changes to flood zone boundaries, insurance access, or mortgage regulations, isolating the insurance cost channel.

4.1 Main Regression

The baseline reduced-form specification is a continuous-treatment difference-in-differences at the property level. Let Y_{izt} denote the outcome for property i in ZIP z in quarter t (transaction price or transaction volume). The estimating equation is:

$$Y_{izt} = \alpha + \beta \cdot \Delta l_z \times \text{Post}_t + \mu_i + \delta_{ct} + \mathbf{X}'_{it}\lambda + \varepsilon_{izt} \tag{1}$$

where $\text{Post}_t = \mathbf{1}[t \geq 2022\text{Q2}]$ marks the Phase II implementation quarter (April 1, 2022), μ_i are property fixed effects that absorb all time-invariant property and neighborhood characteristics, δ_{ct} are county $\times$ quarter fixed effects that absorb local economic shocks common to all properties in the same county-quarter, and $\mathbf{X}_{it}$ includes time-varying controls (mortgage rate environment, flood events). Standard errors are clustered at the ZIP-code level. The coefficient β identifies the average causal effect of a unit increase in treatment intensity on the outcome, relative to low-intensity ZIPs, in the post-reform period.

I construct the treatment intensity measure Δl_z at the ZIP-code level as the share of single-family NFIP policies in ZIP z projected to face a monthly premium increase exceeding \$10, using FEMA's pre-implementation projected premium change files. This measure is pre-determined (published October 2021 before Phase II renewals began) and does not depend on post-reform housing market outcomes, which solves the concern about the anticipation effect.

4.2 Event Study

To test the parallel trends assumption and trace the dynamic response, I estimate the event-study version of equation (1):

$$Y_{izt} = \alpha + \sum_{\tau=-8}^{8} \beta_\tau \cdot \Delta l_z \times \mathbf{1}[t = \tau] + \mu_i + \delta_{ct} + \mathbf{X}'_{it}\lambda + \varepsilon_{izt} \tag{2}$$

where τ indexes quarters relative to Phase II (2022Q2). The reference period is $\tau = -3$ (2021Q3), the last quarter before the Phase I transition window; $\tau = -2$ (2021Q4) and $\tau = -1$ (2022Q1) are excluded as Phase I observations. The pre-treatment coefficients $\{\beta_\tau\}_{\tau<0}$ serve as pre-trend tests: under parallel trends, they should be indistinguishable from zero. The post-reform coefficients $\{\beta_\tau\}_{\tau\geq 0}$ trace how prices respond as more policies roll onto RR2.0.

4.3 Handling Continuous Treatment Heterogeneity

The 18%/year statutory cap creates staggered, household-specific treatment timing: different properties reach their actuarial target in different years. I follow Fabian (2024) in applying the continuous treatment estimator of Callaway et al. (2024) to account for this treatment effect heterogeneity, using the pre-reform subsidy level ($\bar{l}^{\text{pre}} - l^{\text{RR2.0}}$) as the treatment dose and exploiting the differential timing of glide-path completion across properties.

4.4 Distributional Outcomes

To study residential sorting, the outcome variables are augmented to include: (i) the income percentile of the buyer (from CoreLogic deed records linked to mortgage origination data), (ii) the racial composition of in-movers at the ZIP-quarter level (from ACS), and (iii) the change in demographic composition D_{zt} of recent movers. These outcomes are estimated using the ZIP-quarter collapsed version of equation (1), with δ_{st} state $\times$ quarter fixed effects to allow comparisons within state-year cells.

5 Structural Model

The reduced-form estimates capture the causal price and sorting effects of RR2.0, but do not permit welfare analysis or counterfactual policy simulation. For this, I build a dynamic residential sorting model in the tradition of Bayer et al. (2016), extended with heterogeneous flood risk beliefs and insurance pricing.

There are J neighborhoods (ZIP codes), indexed by j, and a continuum of households indexed by type $\theta = (y, \pi, \text{race})$, where y is income and $\pi \in [0, 1]$ is the household's subjective belief about the probability of a flood-damage event. Households are forward-looking and can relocate each period at a moving cost $\kappa(\theta)$.

Each neighborhood j is characterized by: flood risk ρ_j (objective expected annual damage rate), insurance premium $l_j(t)$ (time-varying; changes under RR2.0), housing price $p_j(t)$ (endogenous), non-marketed amenity bundle A_j, and endogenous demographic composition $D_j(t)$.

The flow utility of type-θ household residing in neighborhood j at time t is:

$$u_{ij}(\theta, t) = \beta_1 \ln(y - p_j(t) - l_j(t)) - \beta_2 \, \pi \cdot \rho_j + \beta_3 A_j + \beta_4(\text{race}) \cdot D_j(t) + \varepsilon_{ij} \tag{3}$$

The first term captures the budget cost of housing and insurance; the second is the expected disutility of flood risk, where π weights objective risk ρ_j by subjective beliefs (following Bakkensen and Barrage 2022); the third is a non-marketed amenity value; and the fourth captures preferences for neighborhood composition, which may differ by race. The shock ε_{ij} is i.i.d. Type-1 Extreme Value.

The household's dynamic problem is:

$$V_j(\theta, t) = \max_{j'} \{u_{ij'}(\theta, t) - \kappa(\theta) \cdot \mathbf{1}[j' \neq j] + \beta \, \mathbb{E}_t[V_{j'}(\theta, t+1)]\} \tag{4}$$

Housing supply in neighborhood j responds to prices with elasticity η_j following Saiz (2010):

$$\ln H_j(t) = \eta_j \ln p_j(t) + \xi_j \tag{5}$$

where ξ_j absorbs geography and regulation. Market clearing requires $H_j(t) = \int s_j(\theta, t) \, dF(\theta)$, where $s_j(\theta, t)$ is the equilibrium share of type-θ households choosing neighborhood j.

Before RR2.0, the NFIP community-averaged premium $\bar{l}_j^{\text{pre}}$ involves a cross-subsidy: low-risk properties in high-risk communities subsidize high-risk properties. When RR2.0 moves premiums to $l_j^{\text{RR2.0}} = f(\rho_j, \text{RCV}_j)$, the implicit subsidy to high-risk neighborhoods disappears. The cost of residing in neighborhood j rises by $\Delta l_j = l_j^{\text{RR2.0}} - \bar{l}_j^{\text{pre}} > 0$ for most high-risk properties, reducing demand and housing prices there. Simultaneously, demand shifts toward low-risk neighborhoods, bidding up p_j for low-ρ_j locations. If low-income households have lower willingness to pay for flood safety (as Bakkensen and Ma 2020 document), they are displaced from newly desirable safe neighborhoods—the climate gentrification mechanism.

Estimation proceeds in two steps:

- **Step 1:** use pre-reform (2018–2021) data on housing transactions, NFIP premiums, neighborhood characteristics, and mover demographics to estimate the preference parameters $(\beta_1, \beta_2, \beta_3, \beta_4, \kappa)$ and supply elasticities η_j via BLP-style moment conditions, following the multi-step estimator of Bayer et al. (2016).
- **Step 2:** impose the observed RR2.0 premium changes Δl_j as a policy shock and solve for the new sorting equilibrium. The predicted price changes and demographic shifts are compared to the post-reform reduced-form estimates as an out-of-sample validation exercise. Successful model fit validates the preference estimates and enables counterfactual simulations: means-tested premium subsidies for low-income policyholders, public adaptation investments that reduce ρ_j, and managed retreat subsidies.

6 Preliminary Results

6.1 The National Premium Distribution Under RR2.0

Figure 1 shows the distribution of risk-based annual premiums for the 3.15 million single-family NFIP policies included in FEMA's national Risk Rating 2.0 exhibit. The distribution is right-skewed: 37.1% of policies have a full-risk target premium below $1,000 per year, and 69.6% fall below $2,000 per year. At the other extreme, only 0.5% face risk-based premiums exceeding $10,000 per year. The median full-risk premium nationally is approximately $1,290. It is important to note that these are target premiums under full actuarial pricing, not what policyholders currently pay. FEMA's statutory caps (18%/year for primary residences) mean policyholders are being phased toward these levels along a multi-year glide path. The median current premium is $786 nationally, implying a nationwide implicit subsidy of roughly $504 per policy per year that RR2.0 is designed to eliminate.

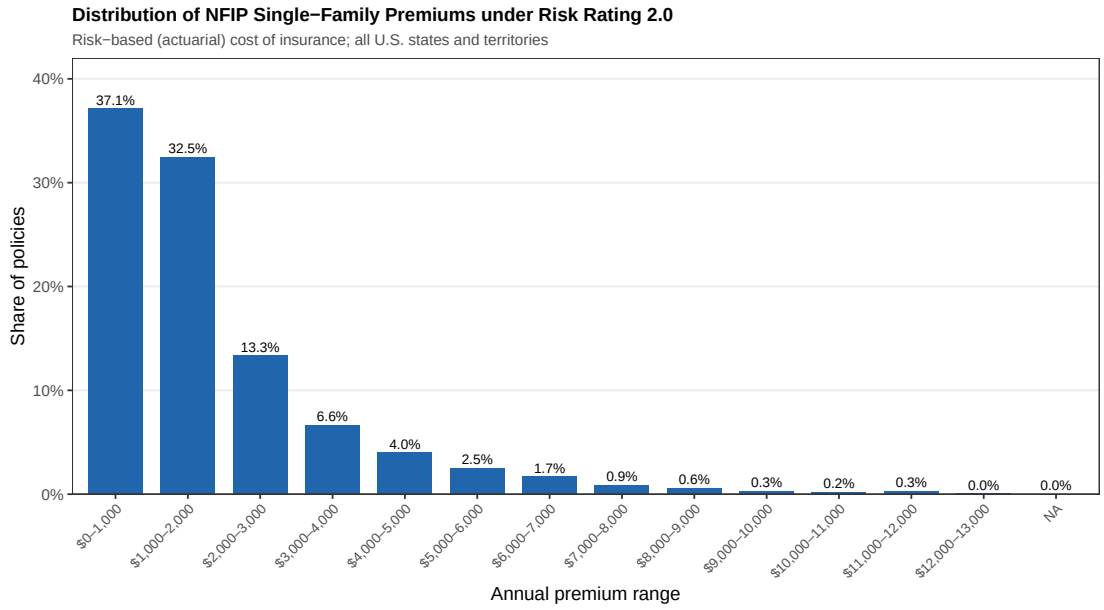
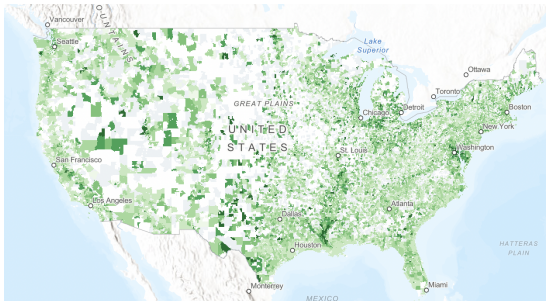


Figure 1: Distribution of single-family NFIP premiums under Risk Rating 2.0 (risk-based target cost)

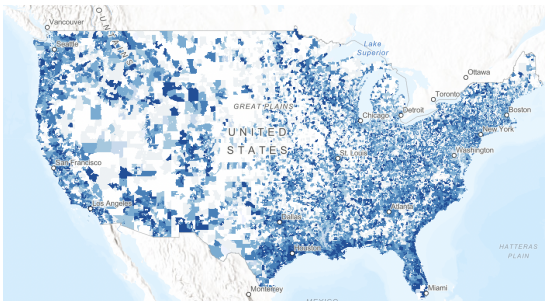
6.2 Geographic Variation

The four maps in figure 2 visualize the geographic variation in premium change using the projected premium changes dataset, each representing one of the categories: decreased monthly premiums, increase of $0-$10 per month, increase of $10-$20 per month, increase of $20+ per month.

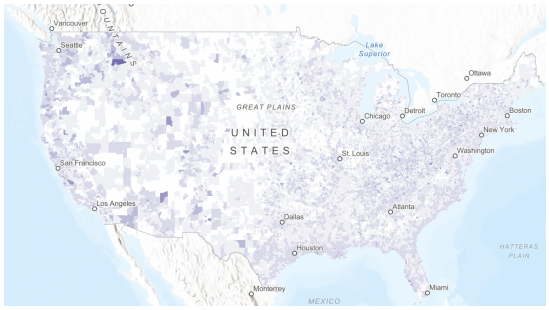
The implicit subsidy varies substantially across states. Figure 3 plots the median risk-



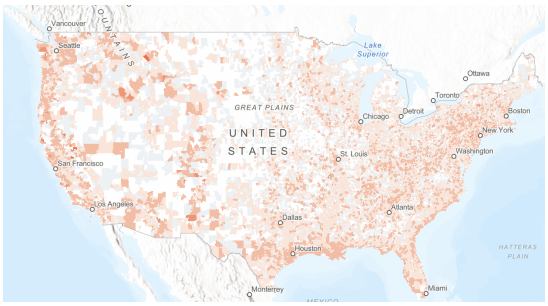
(a) Decreased



(b) Increase \$0-\$10



(c) Increase \$10-\$20



(d) Increase \$20+

Figure 2: Projected Premium Changes by Zip Code

based target premium versus the median currently-paid premium for the 20 states with the largest NFIP enrollment. Florida (894,619 policies), Texas (586,680), and Louisiana (397,824) account for over half of all single-family NFIP policies nationally. Among high-enrollment states, New Jersey (51.0%), New York (48.6%), and Louisiana (46.5%) have the largest percentage subsidies. Hawaii faces the largest absolute implicit subsidy. This geographic variation implies that RR2.0 creates a differential cost shock across states and, within states, across properties with different risk profiles and pre-reform subsidy histories.

6.3 County-Level Treatment Intensity

The FEMA county-level projected premium change files (2021) provide the key treatment variation for identification. Figure 4 shows the distribution across 2,786 counties of the share of single-family NFIP policies projected to face monthly increases exceeding \$10 under RR2.0. The distribution is highly skewed: on average 7.0% of a county's NFIP policies face a >\$10/month increase, while the median county has 33.1% of policies projected to receive decreases. The mass of counties has relatively low treatment intensity, but there is meaningful right-tail variation—several hundred counties have 20%–50% of policies facing substantial increases.

Charlotte County, FL stands out with 53.3% of its 25,846 NFIP policies facing monthly increases exceeding \$10. Other high-intensity counties include Sierra County, CA (57.4%

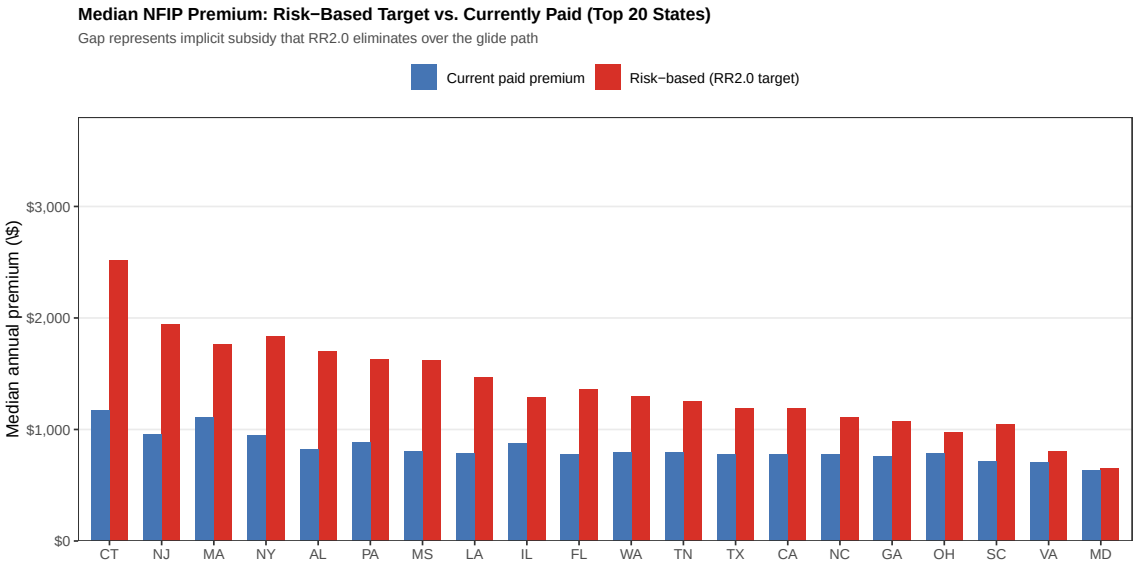


Figure 3: Median NFIP single-family premium

of 54 policies) and several rural inland counties. This variation—concentrated in coastal high-risk areas but with substantial heterogeneity even within states—is what the DiD research design exploits for identification.

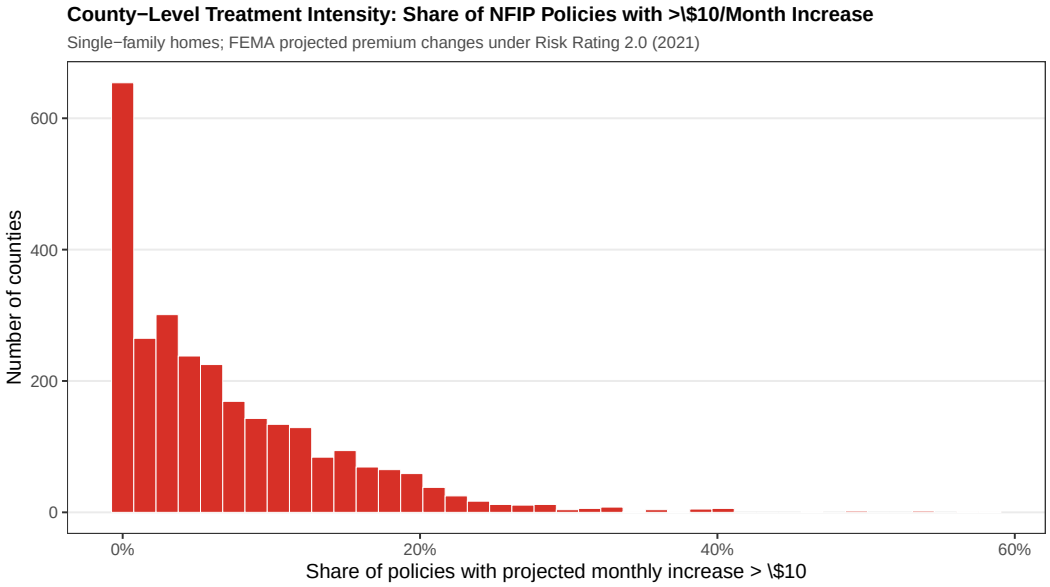


Figure 4: County-level treatment intensity

6.4 Property Insurance Cost Trends

Figure 5 documents the aggregate property insurance cost shock using the Keys and Mulder (2024) mortgage escrow data. Weighted-mean annual property insurance premiums were remarkably stable from 2014–2020 ($1,657 to $1,864), then accelerated sharply:

$2,096 in 2022, $2,433 in 2023, and $2,712 in 2024—a 30.5% nominal increase from 2020 to 2023 alone. The timing aligns closely with the October 2021 RR2.0 Phase I implementation and the April 2022 Phase II rollout. While the Keys–Mulder data captures total homeowners insurance rather than flood-specific premiums, the acceleration reflects the broader repricing of climate risk that RR2.0 is part of.

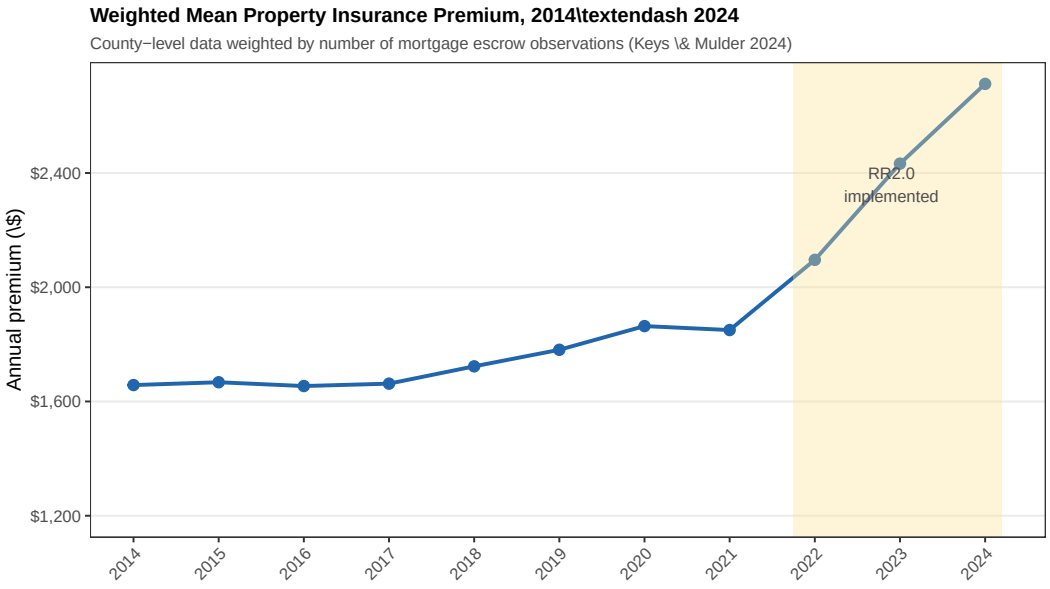


Figure 5: Weighted mean annual property insurance premium, 2014–2024

6.5 RR2.0 Adoption in NFIP Policy Data

Figure 6 shows the rapid adoption of RR2.0 pricing in the OpenFEMA redacted policies data, based on a systematic 1:20,000 sample. The share of policies rated under the new `RatingEngine` method was essentially zero in 2020, rose to 3.4% in 2021 (new policies only, per Phase I), jumped to 82.4% in 2022 (Phase II renewals), and reached 99.5% by 2023. This clean phase-in creates a sharp, policy-driven event for the event-study design.

In the sample, RR2.0 policies carry a higher median annual premium ($617) than pre-reform rated policies ($392), consistent with the removal of historical subsidies. Among the sample, 56.8% of policies are in SFHA flood zones, 17.9% carry the pre-FIRM subsidized rate type (`subsidizedRateType == "P"`)—the set facing the largest potential increases under the glide path.

6.6 CoreLogic Transaction Panel: Summary Statistics

Table 1 summarizes the CoreLogic panel used in estimation. The main price estimation sample—after excluding non-disclosure states, the Phase I transition window (2021Q4–2022Q1), and corporate buyers—contains 5,658,411 transactions from 4,849,126 unique

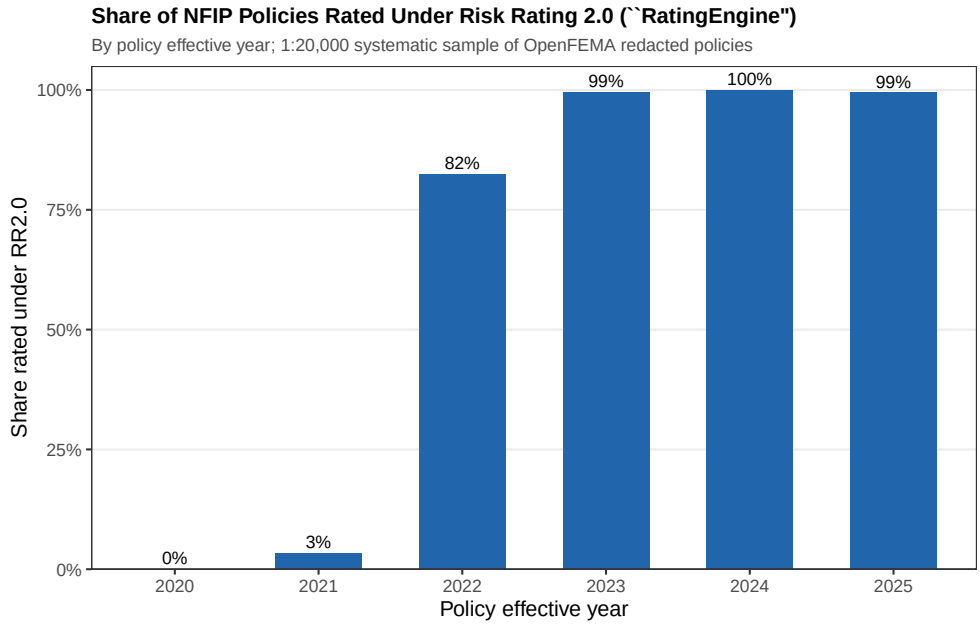


Figure 6: Share of NFIP policies rated under Risk Rating 2.0 by policy effective year

properties across 14,428 ZIP codes.

The median sale price in the estimation sample is $240,000 (log price 12.39), with the 10th and 90th percentiles at $53,000 and $652,000, respectively, reflecting the substantial right skew typical of residential housing markets. Transaction volume is concentrated in the warmer quarters, with 2022Q2 being the highest-volume single quarter in the post-period (438,809 transactions) before the rate-hike-driven cooling of 2022H2. Median prices rose from $185,000 in 2018Q1 to a peak of $310,000 in 2023Q2, consistent with the national housing boom.

Treatment intensity (Δl_z, the share of NFIP policies in each ZIP projected to face monthly increases exceeding $10) is highly right-skewed. The median ZIP has only 1.3% of NFIP policies facing large increases, while the 75th and 90th percentiles are 8.1% and 18.6%, respectively, with the maximum at 85.7%. The large mass at zero treatment intensity reflects the many NFIP ZIP codes where most policyholders face modest changes; identification comes from the meaningful right-tail variation in high-exposure coastal and riverine communities. Out-of-ZIP buyers—buyers whose mailing address differs from the property ZIP, a proxy for in-migration—account for 32.5% of transactions on average and will serve as an early demand-side outcome in the sorting analysis.

6.7 Price Effects of Risk Rating 2.0

Table 2 reports estimates of equation (1) with log sale price as the outcome. The five columns vary the time fixed effect structure and treatment variable, holding the property fixed effects and ZIP-clustered standard errors constant throughout.

Table 1: CoreLogic Transaction Panel: Sample Coverage and Summary Statistics

Panel A: Sample coverage

	Full panel	Estimation sample
Transactions	8,563,362	5,658,411
Unique properties	6,771,758	4,849,126
Unique ZIP codes	16,617	14,428
Unique counties	2,612	2,084
States	49	39
Date range	2018Q1–2024Q1	2018Q1–2024Q1

Panel B: Variable summary statistics (estimation sample)

Variable	N	Mean	SD	P10	P25	P50	P75	P90
Sale price ($000s)	5,658,411	335.86	492.25	52.5	132	239.9	400	652.12
Log sale price	5,658,411	12.25	1.05	10.87	11.79	12.39	12.9	13.39
Treatment intensity Δl_z (%)	5,658,411	6.29	11.1	0	0	1.26	8.11	18.6
Share with premium decrease (%)	5,658,411	29.31	22.76	2.5	10.78	25	42.99	62.65
Out-of-ZIP buyer share (%)	5,658,411	32.46	46.82	0	0	0	100	100

Notes: The full panel covers all CoreLogic residential transactions in ZIP codes with active NFIP policies, 2018Q1–2024Q1. The estimation sample excludes non-disclosure states, the Phase I transition window (2021Q4–2022Q1), and corporate buyers. Sale price is reported in thousands of dollars. Treatment intensity Δl_z is the ZIP-level share of NFIP policies projected to face monthly premium increases exceeding $10 under RR2.0, reported as a percentage. Share with premium decrease is the ZIP-level share projected to receive any decrease. Out-of-ZIP buyer is an indicator for buyers whose mailing address ZIP differs from the property ZIP, a proxy for in-migration used in the sorting analysis.

The main result, in column (3), shows that a one-unit increase in treatment intensity reduces log sale price by 0.148 (SE = 0.043, $p < 0.01$) in the post-Phase II period, controlling for property fixed effects and county × quarter fixed effects. In economic terms, a one-standard-deviation increase in treatment intensity (11 percentage points) is associated with a 1.6% price decline, or approximately $3,900 at the median transaction price of $240,000. The result is robust to including the Phase I transition window (column 4: $\hat{\beta} = -0.154$, SE = 0.044) and barely changes when the event boundary is relaxed in this way, confirming that the Phase I exclusion does not drive the result.

A notable finding is that the column (1) estimate—using only quarter fixed effects rather than county × quarter fixed effects—is positive and significant ($\hat{\beta} = +0.260$). This sign reversal highlights an important confound: coastal and high-flood-exposure areas experienced strong appreciation during the 2021–2022 housing boom for reasons unrelated to RR2.0 (remote-work demand, low interest rates). The county × quarter fixed effects in column (3) absorb these common county-level trends and isolate the within-county differential between high- and low-treatment-intensity ZIP codes. Column (2) with state × quarter fixed effects attenuates the estimate to near zero, suggesting the identifying variation is largely within-county rather than within-state.

Column (5) provides a direction-reversal falsification test: replacing Δl_z with the share of NFIP policies projected to receive premium *decreases*. If price effects are driven by the insurance cost channel, this coefficient should be positive (lower insurance costs raise housing values). The estimated coefficient is +0.091 (SE = 0.021, $p < 0.01$), consistent with the prediction and providing additional confidence that the main result reflects genuine insurance cost capitalization rather than an artifact of sample selection or unobserved ZIP-level trends.

Table 2: Effect of RR2.0 Treatment Intensity on Log Sale Price

	Dependent variable: log(sale price)				
	(1)	(2)	(3)	(4)	(5)
$\Delta l_z \times$ Post	0.2600***	0.0023	−0.1476***	−0.1540***	
	(0.0314)	(0.0309)	(0.0434)	(0.0435)	
$\Delta l_z^{\text{dec}} \times$ Post					0.0913***
					(0.0209)
Observations	5,658,411	5,658,411	5,658,411	6,216,192	5,658,411
R^2	0.973	0.973	0.976	0.974	0.976
Adjusted R^2	0.810	0.814	0.825	0.824	0.825
Property FE	✓	✓	✓	✓	✓
Quarter FE	✓				
State × Quarter FE		✓			
County × Quarter FE			✓	✓	✓
Phase I window included				✓	

Notes: Dependent variable is log sale price. Standard errors clustered by ZIP code in parentheses. Sample excludes non-disclosure states and corporate buyers. Columns (1)–(3) and (5) exclude the Phase I transition window (Oct 2021–Mar 2022). Column (4) includes the Phase I window as a robustness check. Column (5) replaces the treatment variable with the share of NFIP policies projected to receive premium *decreases* under RR2.0 (placebo test). *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Figure 7 presents the event-study version of equation (2). The figure plots the interaction coefficients $\hat{\beta}_\tau$ for each quarter $\tau \in [-8, +7]$ relative to Phase II implementation, with $\tau = -3$ (2021Q3) as the reference period.

Two findings are noteworthy. First, the pre-treatment coefficients ($\tau = -8$ through $\tau = -4$) are uniformly small and statistically indistinguishable from zero, with a joint pre-trend Wald test of $F(5, 271{,}069) = 0.888$ ($p = 0.488$). This provides strong evidence in favor of the parallel trends assumption: before Phase II, high- and low-treatment-intensity ZIP codes within the same county were on statistically identical price trajectories. Second, the post-treatment coefficients become increasingly negative beginning at $\tau = 0$, accelerating through $\tau = 3$–4 (mid-2023, approximately $\hat{\beta} \approx -0.40$) before partially stabilizing. The dynamic pattern—growing price effects for six to seven quarters after the reform—is consistent with the 18%/year statutory glide path: as more policies

reach their actuarial target at each successive annual renewal, the full insurance cost shock materializes and progressively capitalizes into prices.

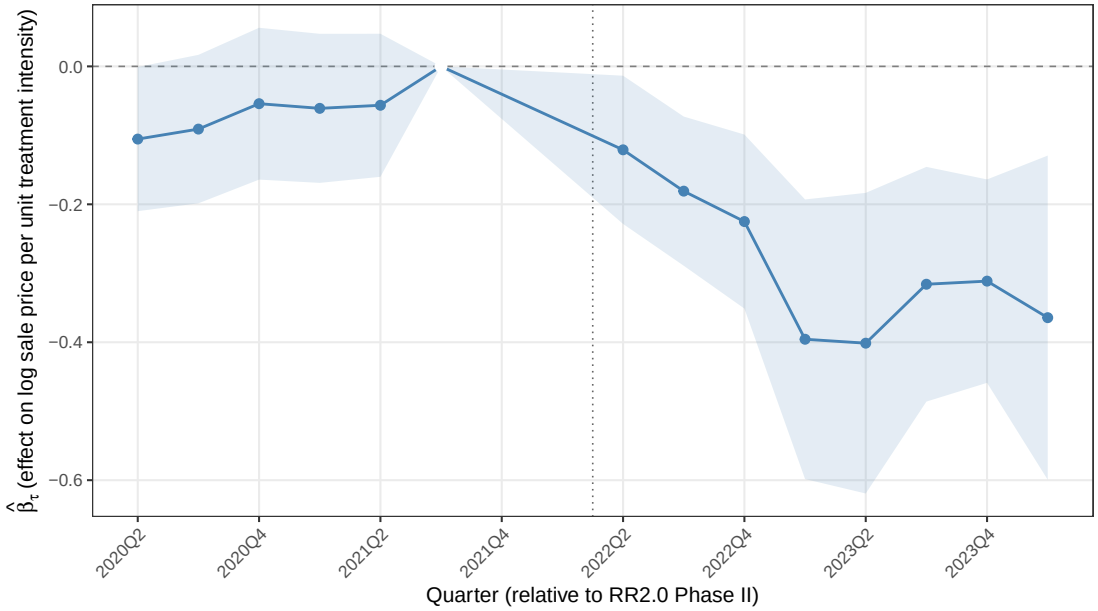


Figure 7: Event study: effect of treatment intensity on log sale price ($\tau = -8$ to $+7$, reference $\tau = -3$)

7 Next Steps

The price results in Section 6 establish the first link in the causal chain. The immediate next step is to extend the reduced-form analysis to the distributional sorting outcomes that constitute the second and third links.

Sorting outcomes. The core distributional question is whether RR2.0 reshapes the income and racial composition of buyers across ZIP codes with different flood exposure. Using the CoreLogic deed panel, I will estimate equation (1) with the following outcome variables: (i) the income percentile of the buyer, imputed from mortgage origination records linked to HMDA filings; (ii) the racial and ethnic composition of in-movers at the ZIP-quarter level, from ACS estimates; and (iii) the share of out-of-ZIP buyers (buyers whose mailing address differs from the property ZIP), as a proxy for in-migration from other markets. These outcomes are estimated using the ZIP-quarter collapsed version of equation (1) with state $\times$ quarter fixed effects. I will also examine heterogeneity by pre-reform neighborhood income, racial composition, SFHA status, and housing supply elasticity (Saiz, 2010) to test whether climate gentrification effects are concentrated in the neighborhoods the theoretical mechanism predicts.

Structural estimation. The reduced-form estimates capture causal price and sorting effects but do not permit welfare analysis or counterfactual policy simulation. I will build a dynamic residential sorting model in the tradition of Bayer et al. (2016), extended with heterogeneous flood risk beliefs and insurance pricing, and estimate it using the pre-reform (2018–2021) portion of the CoreLogic panel. Estimation proceeds in two steps: (i) preference parameters and supply elasticities are estimated via BLP-style moment conditions using pre-reform data; (ii) the observed RR2.0 premium changes are imposed as a policy shock, and the predicted post-reform sorting equilibrium is compared to the reduced-form estimates as an out-of-sample validation. Three policy counterfactuals are of central interest:

1. **Means-tested premium caps:** Cap the annual premium increase for policyholders below 80% of area median income, as proposed in several Congressional reform bills, and quantify the fiscal cost and distributional consequences.
2. **Community flood adaptation investments:** Model the effect of reducing flood risk ρ_j in high-risk neighborhoods (e.g., levee upgrades, coastal protection) on both housing values and sorting patterns, relevant for evaluating FEMA's Community Rating System and BRIC grant program.
3. **Information provision:** Simulate the counterfactual in which all households have correct flood risk beliefs ($\pi = \rho_j$ for all θ), quantifying the additional sorting consequences of correcting belief heterogeneity (Bakkensen and Barrage, 2022; Baldauf et al., 2020).

References

Bakkensen, L. A. and Barrage, L. (2022). Going underwater? Flood risk belief heterogeneity and coastal home price dynamics. *Review of Financial Studies*, 35(8):3666–3709.

Bakkensen, L. A. and Ma, L. (2020). Sorting over flood risk and implications for policy reform. *Journal of Environmental Economics and Management*, 104:102362.

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

Banzhaf, H. S. and Walsh, R. P. (2008). Do people vote with their feet? An empirical test of Tiebout's mechanism. *American Economic Review*, 98(3):843–863.

Bayer, P., McMillan, R., Murphy, A., and Timmins, C. (2016). A dynamic model of demand for houses and neighborhoods. *Econometrica*, 84(3):893–942.

Bernstein, A., Gustafson, M. T., and Lewis, R. (2019). Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics*, 134(2):253–272.

Fabian, J. (2024). The price of risk: Flood insurance premium reform and local development. Technical report, Working paper.

Georgic, W. and Klaiber, H. A. (2022). Stocks, flows, and flood insurance. *Journal of Environmental Economics and Management*, 111:102567.

Gourevitch, J. D., Kousky, C., Liao, Y., Nolte, C., Pollack, A. B., Porter, J. R., and Weill, J. A. (2023). Unpriced climate risk and the potential consequences of overvaluation in us housing markets. *Nature climate change*, 13(3):250–257.

Gruber, J. and Solomon, A. (2026). Optimal flood insurance in a second-best world: Fiscal spillovers, reclassification risk and moral hazard. Technical report, Working paper.

Hino, M. and Burke, M. (2021). The effect of information about climate risk on property values. *Proceedings of the National Academy of Sciences*, 118(17):e2003374118.

Keenan, J. M., Hill, T., and Gumber, A. (2018). Climate gentrification: from theory to empiricism in Miami-Dade county, Florida. *Environmental Research Letters*, 13(5):054001.

Keys, B. J. and Mulder, P. (2024). Property insurance and disaster risk: New evidence from mortgage escrow data. Technical report, National Bureau of Economic Research.

Mulder, P. and Kousky, C. (2023). Risk rating without information provision. *AEA Papers and Proceedings*, 113:299–303.

Ortega, F. and Petkov, I. (2024). To improve is to change? The effects of Risk Rating 2.0 on flood insurance demand. Technical report, IZA Discussion Paper No. 17021.

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

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy*, 64(5):416–424.