When Prices Begin to Reflect Risk: Evidence From A Reform in Hazard Insurance for U.S. Mortgages

Matthew Carey Ricardo Correa Anil Jain†

June 2026

Abstract

Changes in the price of risk can influence credit and investment even when the underlying risk is unchanged. We study this mechanism using a nationwide shift to more granular flood-insurance pricing that raised premiums sharply in some locations, correcting long-standing underpricing. Using detailed insurance records matched to mortgage-level data, we compare census tracts experiencing larger premium increases with nearby tracts that did not. These higher-premium areas saw declines in new mortgages, smaller loan sizes, and shifts toward borrowers with weaker credit, alongside reductions in new construction and lower house prices. The reduction in mortgage credit supply is accompanied by an increase in loan denial rates and a larger reduction in credit for lenders exposed to affected areas. The results show that correcting underpriced insurance reshapes household budgets and lender behavior, including credit standards.

Keywords: Hazard Risk, Insurance, Mortgages
JEL Codes: G12, G21, E31

† All authors are at the Board of Governors of Federal Reserve System. The authors have no relevant financial or non-financial interests to disclose; nor did the authors receive support from any organization for the submitted work. The findings and conclusions in this paper are the sole responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or the views of any other person associated with the Federal Reserve System.

1 Introduction

How do markets respond when the price of risk changes but the underlying hazard stays the same? Addressing this question is key to understanding how adjustments in the cost of insurance may affect credit outcomes, yet empirical evidence is limited. We provide such evidence using a 2021 federal reform that repriced flood insurance nationwide, even with no changes in a household's physical flood exposure. The Federal Emergency Management Agency (FEMA) replaced its coarse, zone-based pricing formula with a granular property-level system, raising premiums in some locations while lowering them in others based on previously unmeasured property characteristics. This reform creates material quasi-random variation in premium *changes* across neighboring properties facing identical flood risk. Linking detailed insurance records to confidential mortgage data, we show that repricing reshaped credit markets: lenders contracted mortgage supply in areas facing more premium increases, borrowers shifted toward weaker credit profiles, and new housing construction declined. The results reveal how correcting long-standing price distortions generates real economic effects through household budgets and lender behavior, independent of changes in underlying risk.

This question has immediate relevance for housing and mortgage markets, and complements a large literature on how disaster risk affects property values. Persistent underpricing of flood insurance has distorted residential investment patterns, inflated property values by over $120 billion (Gourevitch et al., 2023; Hino and Burke, 2020; Niu et al., 2025; Doerner et al., 2025), and concentrated financial risk, potentially affecting both household balance sheets and lender portfolios. Information frictions, adverse selection, and cross-subsidization have created allocative inefficiencies in both insurance and housing markets (Bradt et al., 2021; Zhang, 2023; An et al., 2023). When premiums rise to reflect actual risk, do these distortions unwind smoothly, or does the correction itself create disruptions? Prior research on flood risk has mostly focused on information shocks, that is, how markets respond when beliefs about risk change following disasters, map revisions, or projections about the likelihood of future natural disasters (Keys and Mulder, 2020; Hino and Burke, 2020; Bakkensen and Barrage, 2022; Skouralis et al., 2024; Billings et al., 2025). We provide complementary evidence on price corrections by studying how markets adjust when the price of a known risk changes through policy reform. The distinction matters for policy design because abrupt changes to prices due to policy reforms may materially affect housing markets and households.

We study the 2021 Risk Rating 2.0 reform (RR2), which comprehensively repriced flood insurance premiums administered via the National Flood Insurance Program (NFIP). The reform replaced a zone-based formula that was mostly unchanged since 1968 with property-specific rates incorporating building characteristics, proximity to water, expected flood frequency, and replacement cost. Importantly, FEMA's new formula used information that was always observable (structure elevation, distance from shoreline, foundation type), but not previously incorporated into pricing. The reform generated substantial variation in premium changes across properties within the same neighborhood: some households saw premiums fall by 30 percent or more, while others faced large increases, which were capped at 18 percent per year by statute.[1] This variation, driven by features of FEMA's pricing formulas rather than by differences in actual flood exposure, enables clean identification of repricing effects.

Linking detailed insurance records to comprehensive mortgage data, we find that the repricing of flood insurance policies was associated with a contraction in mortgage lending and residential investment in areas facing the largest premium increases. Mortgage originations fell by approximately 1 percent for each 10 percentage point increase in a tract's exposure to premium hikes, translating to 2-3 fewer mortgages per year in heavily affected neighborhoods relative to nearby areas with smaller increases. Property values associated with mortgage transactions declined by 1.8 percent in the most exposed tracts. The composition of borrowers shifted markedly: new mortgages went to borrowers with credit scores 3 points lower on average, the share of owner-occupied properties rose, and the proportion of first-time homeowners rose. These patterns are consistent with tightening credit access pushing marginal buyers toward lower-cost properties while investors exit high-premium areas.

We corroborate the tightening of credit standards in affected areas by testing for changes in loan application denial and withdrawal rates after RR2 implementation. We find an increase in these loan application adverse outcomes for depository and non-depository lenders. Consistent

[1]Nationwide, roughly two-third of policyholders saw premium increases capped at 18 percent annually—even with these 18 percent annual increases, FEMA estimates that by 2035, 10 percent of policies will still be paying below actuarially-fair priced insurance premiums.

with potential supply side constraints, we find that lenders with larger exposures to affected areas, especially depository institution, contract credit more after the policy reform.

The implementation of RR2 is also associated with changes in investment. New residential construction permits fell by nearly 20 percent in tracts with a 100 percent cap share relative to tracts in the same county with no capped policies, though county-level aggregates show no effect, indicating localized reallocation rather than broad-based contraction.

These patterns reflect a key institutional feature of the reform: new policies originating after October 2021 immediately pay "full-risk rates," that is, the insurance premium consistent with the underlying flood risk of their property, while existing policies face annual increases capped at 18 percent per year, creating a "glide path" of premium increase over potentially many years. As a result, the reform generates mortgage lock-in, reducing transaction volumes and reshaping who transacts, while largely sparing existing borrowers from immediate financial stress. Consistent with this mechanism, effects are concentrated on the extensive margin (who buys homes and at what volume), while existing mortgage performance shows limited response.

Our identification strategy exploits cross-sectional variation in treatment intensity, comparing census tracts which were more versus less exposed to premium increases as a result of the reform. We measure exposure using the share of policies in a tract that hit the statutory 18 percent cap at their first renewal. This "cap share"—the fraction of policies hitting the ceiling within each tract—identifies areas that were most heavily subsidized under the old formula and face the highest proportion of policies facing compounding increases going forward. We compare outcomes across tracts with different cap shares using difference-in-differences specifications with county-by-time fixed effects, absorbing all time-varying local economic conditions and policy changes. Event study specifications confirm parallel pre-trends in mortgage originations, supporting a causal interpretation. We combine four data sources: enhanced NFIP policy records with reliable identifiers linking policies across renewals and accurate census tract coding (unavailable in public data); the National Mortgage Database, a five percent nationally representative sample covering both conforming and jumbo mortgages; confidential Home Mortgage Disclosure Act (HMDA) data for robustness checks and analysis of mortgage denials; and Cotality real estate data capturing property transactions and building permits. Our final sample comprises 1.8 million unique flood insurance policies across 28,781 census tracts, merged with mortgage and housing market outcomes from 2016 to 2024.

This paper makes several contributions to the literatures on disaster risk pricing, insurance market reform, and mortgage credit supply.

First, we provide comprehensive empirical evidence on how insurance repricing affects credit markets and residential investment using realized premium changes rather than projections. Prior work on NFIP reforms has focused on the 2012 Biggert-Waters Act, which targeted a narrower set of properties and was partially reversed (Indaco et al., 2019; Ge et al., 2025a; Collier et al., 2023); RR2 represents a permanent, program-wide repricing with different treatment patterns. Recent concurrent work by Ortega and Petkov (2024) examines RR2's effects on insurance take-up, while Fabian (2024) analyzes construction responses using predicted premium changes in two states. We extend this work by documenting mortgage market responses nationwide and employing realized rather than predicted premium variation.

Second, we contribute a novel continuous treatment measure based on the intensity of pre-reform subsidization. The "cap share" variable captures the pervasiveness of subsidization within small geographic areas, providing additional information compared to binary indicators for FEMA flood zone designation or discrete flood map revision events used in prior studies (Blickle and Santos, 2022; Indaco et al., 2019; Mulder, 2024). Our access to enhanced NFIP data with reliable policy identifiers linked over time represents a substantial improvement over public OpenFEMA records.

Third, we document credit market adjustments to improved risk pricing, distinct from previous findings on flood map revisions (Blickle and Santos, 2022; Arseneau and Kara, 2025) and other information effects (Skouralis et al., 2024; Gillespie et al., 2025; Fairweather et al., 2024; Billings et al., 2025; Miller and Pinter, 2022; Pollack et al., 2023; Sun et al., 2022; Aiba et al., 2025). Blickle et al. (2024) show that lenders already respond to flood risk beyond official FEMA maps, reducing approval rates and loan-to-value ratios for properties with high third-party flood risk scores. We demonstrate that when official insurance pricing changes to better reflect that underlying risk, mortgage markets respond on both the extensive margin (fewer originations) and the intensive margin (compositional shifts toward riskier borrowers), with effects persisting even after controlling for local economic trends. These findings suggest that flood insurance premium increases operate through affordability channels constraining borrower budgets, complementing recent work by Ge et al. (2025b) linking homeowner insurance premium increases to mortgage delinquency.

Fourth, we provide causal evidence on housing investment responses to improved disaster risk pricing. While prior work documents that flood events, sea level rise projections, and belief changes about disaster risk affect property values (Keys and Mulder, 2020; Hino and Burke, 2020; Bakkensen and Barrage, 2022), less is known about how correcting persistent mispricing through insurance reform affects new construction (Peralta and Scott, 2024; Ostriker and Russo, 2024). Our finding of reduced building permit issuance in high-exposure tracts suggests that improved price signals may partially correct inefficient development patterns in flood-prone areas, though the gradual glide path to full-risk rates likely attenuates near-term effects.

The remainder of the paper proceeds as follows. Section 2 describes the institutional features of the NFIP and the RR2 reform. Section 3 presents our data sources, sample construction, and measurement of treatment intensity. Section 4 outlines our empirical strategy, identification assumptions, and results. Section 5 concludes.

2 Institutional Setting

The NFIP was established in 1968 in response to private insurers' retreat from the flood insurance market. Administered through FEMA, the program provides federally backed insurance policies for residential and commercial properties and engages in floodplain risk management. There are over four million active residential flood insurance policies nationwide, corresponding to a national take-up rate just under 4 percent. However, because flood insurance is voluntary for most properties, take-up rates are much lower outside of FEMA's mapped zones of elevated danger, known as Special Flood Hazard Areas (SFHAs). Bradt et al. (2021) estimate that just over 2 percent of homes outside of SFHAs have flood insurance, compared to just under 50 percent of homes within SFHAs.

While private insurance companies are beginning to re-enter the flood insurance market, the NFIP still accounts for the vast majority of policies in force at the time of writing, especially those for residential properties. Approximately two-thirds of these residential NFIP policies are for single-family homes, which form the basis of our analysis sample. Because the program (rather than private competition) sets the price of flood risk for the typical mortgaged household, a change in NFIP pricing maps directly into the cost of owning an at-risk home. That affordability margin is large: Amornsiripanitch et al. (2025) estimate that nearly 90% of at-risk households are under-insured for their flood risk by an average of $7,000 per year, yet over 70% of those households would still benefit from purchasing coverage even after the premium increases RR2 introduced.[2]

Unlike homeowners' insurance, flood insurance is typically not mandatory in the United States. The exception to this rule is homes in SFHAs with active mortgages from federally backed lenders. However, monitoring compliance with this requirement falls to lenders, who often fail to strictly enforce continuous coverage. By tracking the percentage of new home purchases within SFHAs from 2014 to 2018 which maintained flood insurance policies in successive years after their insurance origination, GAO (2021) places a lower bound on compliance rates at approximately 90% after the first renewal year and below 80% after the third renewal year.[3]

The NFIP is funded primarily by the premiums it receives from policyholders. Historically, premium revenue has not been sufficient to cover the costs of paying claims and running the program, leading to a $22.5 billion lifetime budget shortfall owed to the US Treasury. The majority of this shortfall was accumulated in crisis years, such as Hurricane Katrina in 2005 and Hurricane Sandy in 2012, when the sheer volume of claims from catastrophic events far outweighed premium receipts. Revenue generation has been hampered by extensive

[2]Underinsurance is in some cases caused by the NFIP's strict coverage limits of $250,000 for a residential structure and $100,000 for its eligible contents. These limits have not been updated since 1994, leading to a substantial erosion of real coverage value over time.

[3]The issue of noncompliance is not new. Michel-Kerjan et al. (2012) documented substantial policy lapse rates between 2001 and 2009: homeowners maintained NFIP coverage for a median of only 2 to 4 years, well below the median housing tenure of 5 to 6 years during the same period. This gap indicates that many policyholders dropped coverage while still occupying their homes. We replicate their analysis using residential NFIP policy data from 2011 to 2024 and find that the pattern persists: median policy tenure remained around 3 years throughout the 2010s, even as median housing tenure increased.

subsidization of policies throughout the NFIP's lifetime, both from actuarially warranted discounts and from grandfathering practices. Previous reform attempts which attempted to cut or phase out some subsidization, such as the Biggert-Waters Act of 2012, were later scaled back due to affordability concerns.

Before the reform, premiums were often imperfectly related to floor risk, both because the pricing formula was coarse and because the maps underlying it were often badly out of date. The old formula primarily relied on the location of a structure on FEMA's official flood maps inside or outside of an SFHA to determing premiums. These flood maps are still a primary method by which an existing or prospective owner can assess the flood risk of a property. While each map is meant to be updated every five years, most of the United States is covered by older maps; in 2019, 75% of flood maps were over five years old, and 11% were over thirty-five years old (Eby, 2019). Additionally, communities or property holders can pursue an appeals process to alter their maps, which may downgrade the mapped flood risk of their area and which further delays map updates.

In October 2021, the NFIP began implementing RR2, the first comprehensive update to its insurance pricing models and policies since the program's inception. Spurred by a growing recognition that its budget shortfall would not be sustainable in the long term, the program reworked its pricing formulas to incorporate many more granular building characteristics and to remove flood map zoning as a direct determinant of premium cost. Also, for the first time, premiums now take into account the replacement cost for a covered building.

The new pricing formula was applied to new policies starting on October 1, 2021, and to existing policies upon renewal between April 1, 2022 and March 31, 2023. New policies pay the rate applied by the new formula, referred to as the 'full risk rate' by FEMA, immediately. However, nearly all existing policies are subject to a statutorily mandated 18% cap on premium increases, meaning that if a policy's full risk rate tripled under the new formula, the policyholder would only see an 18% premium increase at their first renewal after RR2, and subsequent compounding 18% increases along a 'glide path' until reaching their full risk rate.

These two features—immediate full-risk pricing for new policies and a binding cap for existing ones—shape our empirical design. The cap censors the premium change we observe in heavily subsidized areas, where many policies simply hit the 18% ceiling regardless of how far their full risk rate lies above it. We therefore measure a tract's exposure to the reform by its 'cap share,' which we define as the fraction of existing policies in that tract bound by the cap at their first renewal. This proxies for the *pervasiveness* of pre-reform subsidization within a neighborhood.

The extent of NFIP subsidization prior to RR2 is evident when considering FEMA's own projections of the time horizon after which existing policies will pay actuarially fair rates under the new pricing formula. As of December 2022, only one-third of NFIP policies were at their full risk rate, and by 2027, five years after RR2's full implementation, only half of policies will reach their full risk rate. FEMA further estimates that in 2035, 10 percent of policies will still be below their full risk rate, even after a cumulative 914% increase in premiums accrued since RR2's implementation (GAO, 2023). A policy with $500 premiums in 2021 that remains on the 'glide path' will cost $1,150 in 2026, and $5,070 in 2035. On aggregate, RR2 lowered premiums for approximately 18% of policyholders at the time of its implementation, but it is projected to at least double premiums for nearly one-third of existing policyholders.

We find strong evidence that the reform's effect on policy take-up was concentrated among new policies in the most subsidized tracts. Figure 1 indexes new-policy counts (those in their first twelve months after origination) to September 2021 within cap share bins. New policy take-up in high cap share tracts plummeted once RR2 took effect for new policies in October 2021 and has remained depressed since, while take-up in low cap share tracts rose as homeowners were induced to purchase insurance by relatively cheaper premiums. This pattern is consistent with Ortega and Petkov (2024), who find that renewal rates fell more sharply in the periphery of the flood zone—where cap shares are highest, since these residents benefited most from the old location-based formula—than within the flood zone itself.

Taken together, these institutional features motivate the empirical design that follows, and

Figure 1: New NFIP residential policies in force over time, by cap share

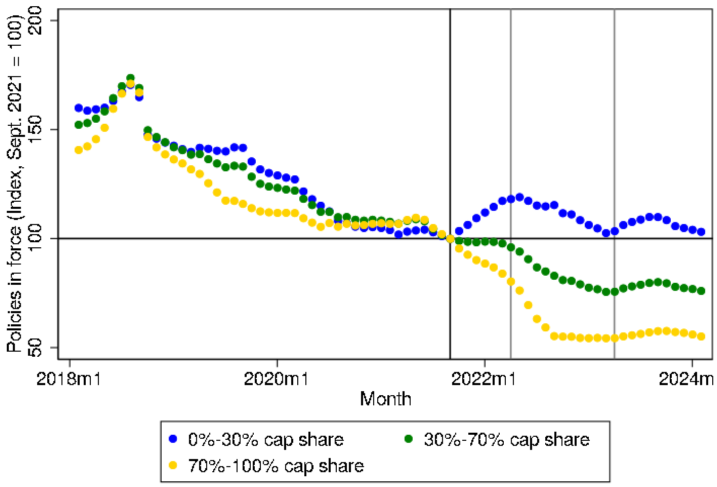

Note: Data from the OpenFEMA NFIP Redacted Policies set, enhanced and maintained by the Federal Reserve Bank of Philadelphia. The sample consists of new residential policies, defined as those in their first twelve months after origination. A policy is considered active for the twelve months following its origination or renewal unless recorded as canceled prior to the twelve-month mark, in which case it is considered active through the month of its cancellation. Policy counts are indexed within cap share bins to September 2021 = 100.

each maps to a specific set of outcomes. First, the structure of the reform (immediate full-risk pricing for new policies alongside a binding 18% cap for existing ones) concentrates the near-term effects of RR2 on the flow of new mortgages rather than on the stock of incumbent borrowers. This leads us to anticipate meaningful effects on the *extensive* margin of mortgage activity (the number of new originations) and on the *composition* of who transacts, while predicting muted effects on the credit performance of existing borrowers, whom the cap shields from large premium increases. Second, because the cap also creates option value that is destroyed upon sale, that same feature predicts reduced mobility, or lock-in, among incumbent homeowners in (previously) heavily subsidized areas. Third, the fact that flood risk reaches lenders through the mandatory-purchase requirement implies that any contraction should operate in part through credit *supply* and should be visible in application denials. Fourth, because the reform shifted prices while leaving FEMA's flood maps unchanged, observed responses are more plausibly attributed to the price of risk than to revised beliefs about the risk itself, a distinction that also shapes how we interpret the response of new residential construction. Finally, the pervasiveness of pre-reform subsidization—evident in FEMA's own glide-path projections—is precisely what our 'cap share' treatment is designed to capture, and it is what we expect to be capitalized into house prices and investment. The remainder of the paper takes up each of these predictions in turn.

3 Data

We obtain comprehensive information on the universe of NFIP policies from 2009 to 2024 from a nonpublic, enhanced version of the FEMA NFIP Redacted Policies dataset maintained by the Federal Reserve Bank of Philadelphia. This data includes a breakdown of policy coverage and costs into premia and other fees, as well as building and policy characteristics such as elevation, community rating, and a mandatory purchase dummy. The key advantage of this vintage is that it contains reliable native identifiers that connect policies over time – since standard flood insurance policies are renewed yearly, a unique policy in effect for n years will have n observations in the data, one for each renewal. These identifiers are not present in the publicly available OpenFEMA vintage, forcing researchers to create their own using nominally time-invariant characteristics such as census block group, base floor elevation, and policy original new business date (see Mulder and Kousky (2023), Ortega and Petkov (2024)). Critically, census tract and block group data is not reliable over time in the public vintage, which has not previously been noted in the literature. As a result, existing algorithms for linking policies over time may substantially mis-identify policy-year observations as distinct when there are discrepancies in the reporting of census location measures.

Our main analysis sample consists of residential policies on single-family homes with policy origination dates prior to October 1, 2021 and at least one renewal on or after April 1,

2022. These conditions only admit policies that were both purchased under FEMA's old pricing regime and renewed at least once under RR2. We further restrict to policies that did not change their coverage amount at the occasion of their first renewal under RR2, and to census tracts with at least eight policies in force in January 2021. Summary statistics for the relevant years of the analysis sample are in Table 1; a broader pre- and post-reform summary snapshot is in Appendix Table A1. The sample includes over 1.8 million unique policies, 87% of which are for primary residences and 42% of which are located in a mapped flood zone. Median coverage for the sample policies is at the $350,000 combined coverage limit for a building and its contents.

Table 1: Descriptive statistics for the main flood policy sample

	Pre-reform		Post-reform	
	Mean	Median	Mean	Median
Insurance premium (annual)	653.03	442.00	646.11	522.00
Total coverage	311324.31	350000.00	311324.31	350000.00
Building replacement cost	348690.90	250000.00	348690.90	250000.00
Cost per $1000 coverage	2.50	1.36	2.38	1.49

	Pre-reform	Post-reform
Located in SFHA	0.424	0.424
Mandatory purchase	0.163	0.163
Primary residence	0.866	0.866
In CRS community	0.729	0.733
N	1,813,760 (50.00%)	1,813,760 (50.00%)

Note: Data from the Federal Reserve Bank of Philadelphia's vintage of the NFIP Redacted Policies set, at the policy-year level. The policy sample is balanced across the pre- and post-reform periods, including only residential policies for single-family homes in tracts with at least eight policies in force as of January 2021. Pre-reform refers to the last renewal year before the implementation of RR2, and post-reform refers to the first renewal year after the implementation of RR2. SFHA: Special Flood Hazard Area. CRS: Community Rating System.

For each census tract, we generate the distribution of premium changes in the first renewal year under RR2, obtaining statistics for all policies within a tract and the subsets of mandatorily purchased policies and voluntarily purchased policies. We offer descriptive statistics for our 'cap share' analysis variable in Table 2, for both individual policies and tract aggregations. The median policy in our analysis sample was capped in the first renewal year, and the median census tract had a 'cap share' of 50%.

Table 2: Descriptive statistics for the 'cap share' treatment variable

	Policies			Tracts		
	Mean	Median	Std. Dev.	Mean	Median	Std. Dev.
First-year premium change	0.096	0.179	0.192			
Cap dummy	0.626	1	0.484			
Cap share				0.473	0.5	0.282
Observations		1,813,760			28,781	

Note: Data from the Federal Reserve Bank of Philadelphia's vintage of the NFIP Redacted Policies set, at the policy-year (left columns) and census tract (right columns) levels. Policy-level data is from the first renewal of each policy under the RR2 pricing formula. First-year premium change is given as a percentage; "cap dummy" is an indicator taking the value of 1 if a given policy had its premium increase by 18% year-over-year around the implementation of RR2.

These tract-level measures are then merged with loan-level mortgage data. We use the National Mortgage Database (NMDB), a nationally representative five percent sample of mortgages across the United States maintained by the Federal Housing Finance Agency (FHFA). The NMDB allows us to track loan- and borrower-level data both at the time of origination and across the life of a loan, augmenting our analysis by incorporating additional data on non-mortgage credit outcomes for borrowers. Importantly, the NMDB includes both conforming and jumbo loans, providing a broad view of the mortgage market at all price points. Descriptive statistics for our analysis sample of NMDB loans are in Table 3. We validate our use of the NMDB by replicating our central results with loan-level confidential Home Mortgage Disclosure Act (HMDA) data, the standard mortgage data set used in the literature.

Table 3: Summary statistics for NMDB loans and census tracts

	(1)			
	Cap share...			
	below median	above median	Diff.	SE
ARM Dummy	0.03	0.02	0.01	(0.01)
GSE Dummy	0.44	0.45	-0.01	(0.02)
Owner Dummy	0.90	0.90	-0.00	(0.01)
Term (months)	331.58	331.48	0.11	(2.52)
First-time Owner	0.41	0.37	0.04*	(0.02)
Credit Score	731.49	735.31	-3.82	(2.23)
Age	42.49	43.47	-0.98*	(0.44)
DTI	35.55	36.21	-0.66*	(0.33)
Interest Rate	4.21	4.21	0.01	(0.03)
Value (Ln)	12.32	12.40	-0.08***	(0.02)
Tract median income (Ln)	11.42	11.44	-0.02	(0.01)
Policies in force (Jan. 2021)	68.69	141.80	-73.11***	(6.70)
Total policies in cap share sample	50.75	110.66	-59.91***	(5.42)
Mandatorily purchased policies	8.47	16.26	-7.79***	(1.24)
Total loans in 2020	50.78	50.50	0.28	(1.10)

Note: Loans and census tracts are split by their tract cap share being above or below the nationwide median of 50%. Data are from the NMDB and NFIP Redacted Policies data sets. NMDB data spans 2016-2024; NFIP statistics calculated from policy-year observations between April 2021 and March 2023. The 'Diff.' column reports the results of a two-sample t-test for each variable; 'SE' is the standard error of that t-test.

Further analysis is conducted using the Cotality Owner Transfer, Property Basic, and Building Permit data sets. All three contain information collected from municipalities on the status, value, and characteristics of properties and permits. National coverage is uneven over time, peaking at approximately 2000 jurisdictions across the United States. We merge building permit and property transaction data from 2018 to 2024 on shared identifiers with property basic data to connect permits with their census tracts. While this data has some inconsistencies due to its sourcing from municipalities with idiosyncratic reporting conventions, we only utilize fields that have been cleaned or standardized.

4 Results

The main hypothesis that we test is whether changes in flood insurance premiums—unrelated to the risk profile of properties—trigger a change in mortgage and housing market conditions. In addition, we further explore whether lenders with different characteristics adjust their credit supply differentially after the flood insurance reform.

The statutory 18 percent annual cap on premium increases creates a truncation that shapes our measurement strategy. Policies requiring large upward adjustments are compressed into the same 18 percent increase at their first renewal—regardless of whether their actuarially fair premiums would rise by 80, 150, or 200 percent. This truncation means that in tracts with heavily subsidized policies, we observe only that premiums hit the cap, not the underlying variation in the extent of policy mispricing. We exploit this feature of the policy to define our main treatment measures. Cap share, our treatment measure, captures the prevalence of pre-reform subsidization within a tract by measuring the fraction of policies constrained by the 18 percent cap.

We test our main hypothesis using this baseline specification:

$$y_{tract,month} = \beta \times post \times \text{cap share} + \alpha_{geo \times time} + \alpha_{tract} + \varepsilon_{tract,time} \tag{1}$$

where the outcome is measured at the tract-time level ($y_{t,m}$)—for example, the number of new mortgage originations. We include either state-by-time or county-by-time fixed effects alongside tract fixed effects, and cluster standard errors at the tract level. This specification absorbs all common shocks within states or counties over time, so that identification comes from comparing tracts with different exposure to pre-reform subsidization within the same state or county in the same time period.[4] We estimate Poisson models for count outcomes (for example, mortgage originations or new listings) and OLS for continuous outcomes.

In this specification, the coefficient of interest is β, which captures the change in the dependent variable for tracts with a higher share of capped flood policies after the reform, relative to other tracts. In the following sections, we will explore the effect of the reform on mortgage credit, house prices, and borrower characteristics.

4.1 Extensive Margin: Mortgage originations across tracts

We begin by examining how mortgage origination activity responds to the intensity of RR2 within a census tract. Figure 2 plots coefficients from a Poisson event study, showing the evolution of new mortgage counts for tracts with higher cap shares relative to tracts with lower cap shares, all within the same county-month. The specification includes census tract

[4]Our estimates are unaffected by FEMA's Community Rating System (CRS), which offers premium discounts of up to 45 percent to communities that adopt flood mitigation measures. Because CRS discounts apply proportionally to both pre- and post-reform premiums, they do not alter the change in premiums induced by Risk Rating 2.0. Moreover, CRS participation is absorbed by tract fixed effects.

and county-month fixed effects, so identification comes from within-county differences in pre-reform subsidization. The figure shows a clear drop in originations beginning immediately after implementation, with no evidence of differential pre-trends. Table 4 quantifies this relationship: tracts with more capped policies see fewer originations relative to other tracts, including those in the same state (column 2) or county (column 3) following the reform. The point estimate in column 3 implies that a tract with a 10 percent cap share experiences roughly 1 percent fewer new mortgages after the reform than a tract in the same county with no capped policies.

Table 4: New mortgage originations across tracts by cap share

	(1) Total Loans	(2) Total Loans	(3) Total Loans
Post x Share Capped	-0.0603***	-0.0966***	-0.0972***
	(0.0197)	(0.0211)	(0.0222)
Observations	3040740	3040216	2713534
Month FE	Yes	N/A	N/A
Tract FE	Yes	Yes	Yes
State x Month FE	No	Yes	No
County x Month FE	No	No	Yes
Cluster	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Poisson regressions on aggregated NMDB data spanning 2016-2024. The outcome variable is the number of new mortgage originations in a given census tract and year. Regressions include progressively more stringent geography-by-time fixed effects; standard errors clustered at the tract level.

Figure 2: Evolution of new mortgage originations across tracts by cap share

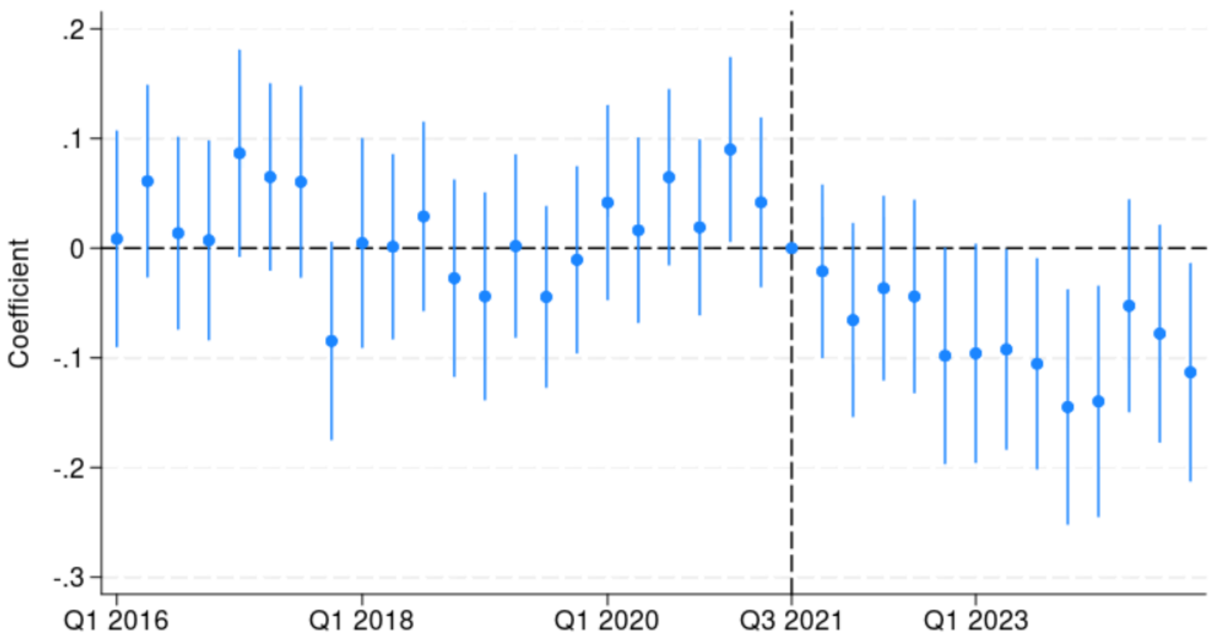

Note: Poisson regression on NMDB data spanning 2016-2024, aggregated to the tract-quarter level. Outcome variable: new mortgage originations. Explanatory variable: cap share of policies interacted with quarter dummies. The dashed line denotes the last quarter before the implementation of Risk Rating 2.0 for new policies. Regression includes county × quarter fixed effects; standard errors clustered at the tract level.

4.2 House price changes after RR2 implementation

We next ask whether the repricing of flood insurance was capitalized into house prices. We estimate transaction-level versions of equation (1) using the universe of arm's-length sales with recorded sale amounts in the Cotality Owner Transfer data. We use the log sale amount as the outcome variable and include tract and county-by-quarter fixed effects. Standard errors are clustered at the tract level.

Table 5 presents the results. In the baseline specification (column 1), sale prices in a tract with a 100 percent cap share fall by approximately 1.8 percent after the reform relative to tracts in the same county with no capped policies. Because the reform also changes *which* homes transact, we progressively saturate the specification with property characteristics—home age, parcel acreage, bedrooms, bathrooms, square footage, and number of stories. The coefficient of interest remains between −1.4 and −1.8 percent throughout (columns 2–6), indicating that the price decline reflects genuine capitalization rather than a shift in the composition of transacted properties. Restricting the sample to exclusively mortgage-financed transactions yields nearly identical estimates (Appendix Table A2). The compositional shift is nonetheless informative in its own right: homes transacting in exposed tracts after the reform are older and somewhat smaller (Appendix Table A3), consistent with marginal buyers being pushed toward lower-cost properties.

These price declines also reflect a demand-side mechanism: the reform's structure creates lock-in among existing homeowners, who face capped premiums that rise at most 18 percent annually while new buyers would pay full risk-based rates, reducing the pool of potential buyers and putting downward pressure on prices. These lock-in effects are examined in detail in Section 4.5, where we show that properties in high cap-share areas became significantly less likely to transact after the reform (Table 12).

4.3 Intensive Margin: Composition of lending

The NMDB also allows us to examine how RR2 affected the intensive margin of mortgage supply by studying the characteristics of newly originated loans. After the reform, census tracts with a higher share of capped policies exhibit a noticeable shift in the composition of new borrowers relative to other tracts in the same county, with lower average credit scores and a higher share of first-time owners (Table 6). Quantitatively, a tract with a 100 percent cap share experienced an average credit score decline of just over three points among new borrowers compared with a tract in the same county with no capped policies. Homes are more likely to be owner-occupied, consistent with investors exiting more expensive high cap share locations, and the share of first time owners rises. Taken together with the point estimates for DTI and interest rates, these patterns suggest that lenders may have tightened underwriting standards in more affected areas, requiring larger down payments or offering

Table 5: House price changes after RR2 implementation

	ln(Sale Amount)					
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Share Capped	-0.0181***	-0.0180***	-0.0186***	-0.0179***	-0.0145***	-0.0145***
	(0.00300)	(0.00292)	(0.00287)	(0.00285)	(0.00274)	(0.00274)
Age of home (years)		-0.00603***	-0.00620***	-0.00575***	-0.00339***	-0.00338***
		(0.0000501)	(0.0000547)	(0.000152)	(0.0000393)	(0.0000395)
ln(Parcel acreage)			0.168***	0.161***	0.0754***	0.0751***
			(0.00211)	(0.00329)	(0.00120)	(0.00120)
Number of bedrooms				0.00187	-0.00560**	-0.00561**
				(0.00858)	(0.00219)	(0.00220)
Number of bathrooms				0.0349***	0.0110***	0.0110***
				(0.0112)	(0.00320)	(0.00322)
ln(Square footage)					0.647***	0.649***
					(0.00451)	(0.00442)
Number of stories						-0.00304**
						(0.00134)
Constant	12.62***	12.86***	13.08***	12.96***	7.970***	7.961***
	(0.000453)	(0.00205)	(0.00361)	(0.0458)	(0.0297)	(0.0291)
Observations	8238585	8238585	8238585	8238585	8238585	8238585
Adjusted R^2	0.658	0.684	0.707	0.714	0.774	0.774
Tract FE	Yes	Yes	Yes	Yes	Yes	Yes
County × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Cluster	Tract	Tract	Tract	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Data from the Cotality Owner Transfer set. The sample spans 2016-2024, and includes all arm's-length transactions regardless of whether or not sales have an associated mortgage. Only properties with the full slate of relevant characteristics are included in all samples for consistent comparison. The outcome variable is the log sale amount for the transaction. Columns progressively add controls for property characteristics (home age, parcel acreage, bedrooms, bathrooms, square footage, number of stories). All regressions include tract and county × quarter fixed effects; standard errors clustered at the tract level.

credit only on smaller loan amounts to borrowers with weaker observable credit profiles. We turn to direct evidence on this supply-side interpretation in the next section.

Table 6: Characteristics of new mortgage loans in exposed tracts

	(1)	(2)	(3)	(4)	(5)
	Credit Score	Owner-Occupied	First-time Owner	DTI	Interest Rate
Post x Share Capped	-3.440***	0.0127***	0.0198***	-0.0917	0.0131
	(0.858)	(0.00391)	(0.00666)	(0.142)	(0.0120)
Constant	740.4***	0.912***	0.424***	37.41***	4.461***
	(0.136)	(0.000620)	(0.00106)	(0.0225)	(0.00190)
Observations	530711	530717	530717	530717	530717
Adjusted R^2	0.143	0.110	0.108	0.074	0.721
Tract FE	Yes	Yes	Yes	Yes	Yes
County x Month FE	Yes	Yes	Yes	Yes	Yes
Cluster	Tract	Tract	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Data are from the NMDB and span 2016-2024. The outcome variable for columns (2) and (3) is a dummy variable indicating the status of the given characteristic for each new loan. All regressions include tract and county × month fixed effects; standard errors clustered at the tract level.

4.4 Mechanisms: credit supply and risk transfer

The contraction in originations and the compositional shifts documented above could in principle reflect either reduced demand from prospective buyers facing higher insurance costs or reduced supply from lenders managing their exposure to repriced flood risk. Both channels appear to be at work. On the supply side, we use loan-level cHMDA data, which records both successful and unsuccessful applications along with lender identity, to assess whether lenders tightened credit standards in exposed areas. On the demand side, the reform creates lock-in among existing homeowners, reducing mobility and shrinking the pool of transacting properties. We begin with the supply-side channel, returning to direct evidence on lock-in in Section 4.5.

Adverse outcomes. We first examine application outcomes at the loan level, regressing indicators for an adverse outcome (denial or withdrawal) on the interaction of the post dummy and tract cap share. We include county-by-quarter, lender-by-quarter, and tract-by-lender fixed effects in these regressions. These saturated specifications compare outcomes for applications to the *same lender* in the *same tract* before and after the reform, netting out lender-specific time trends. Table 7 shows that the probability of an adverse outcome rose by 0.43 percentage points in a tract with 100 percent cap share, with denials alone rising by 0.20 percentage points—an approximately 2 percent increase relative to the mean denial rate. Conditional on applying, borrowers in exposed tracts became *more* likely to be turned down after the reform.[5] These results are consistent with a decrease in lenders' mortgage supply in areas with broader increases in flood insurance premiums, as reflected by our cap share measure.

Lender heterogeneity. We further explore the mechanism driving the contraction of credit in affected areas by testing whether there are differences in credit supply across depository and non-depository institutions, and within depository institutions, across smaller and larger

[5]Tract-level denial *counts* fall in high cap share tracts after the reform (Appendix Table A5). This reflects the contraction in application volume rather than looser standards: with fewer applications submitted, fewer are denied in absolute terms even as the denial rate conditional on applying rises.

Table 7: Adverse outcomes in loan applications

	(1) Adverse Outcome	(2) Denial	(3) Withdrawal
Post × Share Capped	0.00425**	0.00200**	0.00160
	(0.00205)	(0.000920)	(0.00162)
Constant	0.306***	0.0989***	0.178***
	(0.000419)	(0.000188)	(0.000331)
Observations	16970401	16970401	16970401
Adjusted R^2	0.140	0.162	0.0641
County × Quarter FE	Yes	Yes	Yes
Lender × Quarter FE	Yes	Yes	Yes
Tract × Lender FE	Yes	Yes	Yes
Clusters (Separate)	Tract, Lender, Quarter	Tract, Lender, Quarter	Tract, Lender, Quarter

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Loan-level cHMDA regressions of application-outcome indicators on Post × Share Capped. The sample is all home purchase loan applications recorded in cHMDA, for properties in census tracts with nonmissing cap share, from 2018 to 2024. An "adverse outcome" is a denial or a withdrawal of an application. Withdrawal can occur during pre-approval, prior to a credit decision, or after approval but prior to origination. All regressions include county × quarter, lender × quarter, and tract × lender fixed effects; standard errors clustered separately at the tract, lender, and quarter levels.

banks. We conjecture that depository institutions may be constrained by their risk management posture and regulatory compliance, while non-depository institutions may adjust their lending depending on the demand from the securitization market.

In Table 8, we test whether there are differences across these institutions in loan originations to borrowers in areas with higher cap shares after the implementation of RR2. Column (1) reports the coefficient on a triple interaction term between the post dummy, cap share, and an indicator variable equal to one for depository institutions. We find that depository institutions lend more after the implementation of RR2 in affected areas relative to non-depository institutions. However, as shown in columns (2) and (3), the number of new loans decreases for both types of institutions in affected areas. Focusing on depository institutions, we find that smaller banks do not change their lending differently compared to larger depository institutions in areas with higher cap shares (column 4). These results suggest that depository institutions are able to accommodate the risk of flood insurance repricing better than the non-depository counterparts.

We also test for changes in adverse outcomes across institutions to determine the adjustment in credit standards after the implementation of RR2. As reported in column (1) of Table 9, we find no significant difference in the change of adverse outcomes in affected areas in the post period. We do find a statistically significant increase in adverse outcomes for depository and non-depository institutions in affected areas, consistent with the findings in Table 7. Within banks, we find no material difference between smaller and larger banks in terms of adverse outcomes. These results suggest that credit decisions become more scrutinized in affected areas after the reform, but this applies across all institutions.

Lastly, if the credit contraction operates through lenders' management of their own risk exposure, it should be concentrated among lenders for whom that exposure is more important. Table 10 provides supporting evidence that pre-reform exposure to high cap-share areas is a key determinant of subsequent credit supply.

Lenders with the highest pre-reform exposure contracted lending the most in affected tracts following the implementation of RR2 (column 3). However, this effect differs markedly by lender type: depository institutions with high exposure pull back sharply, while non-depository institutions with comparable exposure show no statistically significant reduction (columns 3 and 4). This pattern reflects their distinct business models. Non-depository lenders primarily "originate-to-distribute," thereby retaining little credit risk on their balance sheets, so their lending decisions depend less on their own portfolio concentration and more on downstream demand from securitization markets. Depository institutions, in contrast, face balance sheet constraints—potentially driven by regulatory or supervisory considerations—that make portfolio concentration more binding, leading them to rebalance away from areas where repriced flood risk increases borrower credit risk and expected losses.

Table 8: New loan originations by lender type

	New Mortgage Loan Originations			
	(1)	(2)	(3)	(4)
Post × Share Capped	-0.144***	-0.129***	-0.119***	-0.138***
	(0.0264)	(0.0268)	(0.0236)	(0.0312)
Post × Depository × Share Capped	0.0584**			
	(0.0272)			
Post × Small Bank × Share Capped				0.0544
				(0.0346)
Constant	1.288***	1.404***	1.084***	1.084***
	(0.00394)	(0.00461)	(0.00363)	(0.00359)
Observations	5263043	3180021	2082517	2082517
Sample	All Lenders	Nonbanks	Banks	Banks
Pseudo R^2	0.485	0.507	0.443	0.443
County × Year FE	Yes	Yes	Yes	Yes
Lender × Year FE	Yes	Yes	Yes	Yes
Tract × Lender FE	Yes	Yes	Yes	Yes
Clusters (separate)	Tract, Lender, Year	Tract, Lender, Year	Tract, Lender, Year	Tract, Lender, Year

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Poisson regressions of new mortgage originations on the triple interaction of a post dummy, tract-level cap share, and lender type. The data are from cHMDA, span 2018-2024, and are collapsed to the tract-lender-year level. Information on lenders is obtained from the Federal Reserve Bank of Philadelphia Consumer Finance Institute's HMDA Lender File (the 'Avery file'). "Depository" indicates a depository-institution HMDA filer. "Small Banks" are depository institutions with fewer than $10 billion in total assets as of December 31, 2024. Standard errors clustered separately at the tract, lender, and year levels.

Table 9: Adverse outcomes by lender type

	Adverse Outcome			
	(1)	(2)	(3)	(4)
Post × Share Capped	0.00729***	0.0104***	0.00692**	0.00393
	(0.00140)	(0.00260)	(0.00197)	(0.00352)
Post × Depository × Share Capped	-0.00424			
	(0.00310)			
Post × Small Bank × Share Capped				0.00209
				(0.00587)
Constant	0.302***	0.303***	0.298***	0.298***
	(0.000223)	(0.000467)	(0.000312)	(0.000500)
Observations	13242285	8561641	4680655	4680578
Sample	All Lenders	Nonbanks	Banks	Banks
Adjusted R^2	0.133	0.153	0.0856	0.0908
Tract FE	N/A	Yes	Yes	N/A
Year FE	N/A	Yes	Yes	N/A
County × Year FE	Yes	No	No	Yes
Lender × Year FE	Yes	Yes	Yes	Yes
Tract × Lender FE	Yes	Yes	Yes	Yes
Clusters (separate)	Tract, Lender, Year	Tract, Lender, Year	Tract, Lender, Year	Tract, Lender, Year

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Loan-level cHMDA regressions of application-outcome indicators on a post dummy, tract-level cap share, and lender type. The data are from cHMDA, span 2018-2024 and are loan-level. Information on lenders is obtained from the Federal Reserve Bank of Philadelphia Consumer Finance Institute's HMDA Lender File (the 'Avery file'). "Depository" indicates a depository-institution HMDA filer. "Small Banks" are depository institutions with fewer than $10 billion in total assets as of December 31, 2024. Standard errors clustered separately at the tract, lender, and year levels.

4.5 Existing borrowers and the glide path

The structure of the reform generates a sharp prediction: because the 18 percent annual cap shields existing policyholders from immediate premium shocks, the financial position of *current* borrowers should respond far less than the flow of *new* lending in the short term. The data bear this out. We find no statistically significant effects on the credit outcomes of existing mortgage borrowers in either high–cap-share tracts or tracts with widespread premium decreases. Although the point estimates in Table 11 suggest modest declines in credit scores and increases in credit card debt for borrowers in high–cap-share tracts, with the opposite pattern in areas experiencing premium reductions, none of these estimates are

Table 10: Lender pre-reform exposure and new loan originations

	New Mortgage Loan Originations				
	(1)	(2)	(3)	(4)	(5)
Post × Share Capped	-0.144***		-0.0397	-0.0806	0.0350
	(0.0264)		(0.0392)	(0.0614)	(0.0371)
Post × Depository × Share Capped	0.0584**				
	(0.0272)				
Post × Lender Exposure		-0.335			
		(0.338)			
Post × Share Capped × Lender Exposure			-0.340**	-0.203	-0.571***
			(0.144)	(0.259)	(0.129)
Constant	1.288***	1.249***	1.288***	1.404***	1.086***
	(0.00394)	(0.0273)	(0.00407)	(0.00469)	(0.00369)
Observations	5263043	5362975	5263043	3180021	2082517
Sample	All Lenders	All Lenders	All Lenders	Nondepository only	Depository only
Pseudo R^2	0.485	0.461	0.485	0.507	0.443
County × Year FE	Yes	Yes	Yes	Yes	Yes
Lender × Year FE	Yes	No	Yes	Yes	Yes
Tract × Lender FE	Yes	Yes	Yes	Yes	Yes
Clusters (separate)	Tract, Lender, Year	Tract, Lender, Year	Tract, Lender, Year	Tract, Lender, Year	Tract, Lender, Year

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Poisson regressions of new mortgage originations on the triple interaction of a post dummy, tract-level cap share, and lender exposure. The data are from cHMDA, span 2018-2024, and are collapsed to the tract-lender-year level. Information on lenders is obtained from the Federal Reserve Bank of Philadelphia Consumer Finance Institute's HMDA Lender File (the 'Avery file'). "Depository" indicates a depository-institution HMDA filer; "Lender Exposure" is the lender's pre-reform (2018-2020) share of originations in capped census tracts, weighted by dollar volume. "Small Banks" are depository institutions with fewer than $10 billion in total assets as of December 31, 2024. Standard errors clustered separately at the tract, lender, and year levels.

statistically distinguishable from zero. Rather than a non-result, this null is a confirmation of the reform's design: the glide path concentrates the near-term incidence of RR2 on the extensive margin (who buys homes, at what price, and with what financing) while leaving incumbents largely insulated.

Table 11: Credit outcomes for existing mortgage borrowers

	(1) Credit Score	(2) Overdue Rate	(3) Credit Card Debt	(4) Credit Score	(5) Overdue Rate	(6) Credit Card Debt
Post x Share Capped	-0.502	-0.00324	101.4			
	(0.955)	(0.00209)	(201.4)			
Post x Share $\geq$ 20% Decrease				1.416	0.000756	-355.7*
				(0.993)	(0.00212)	(213.7)
Constant	727.7***	0.0196***	9562.4***	727.4***	0.0187***	9636.8***
	(0.243)	(0.000521)	(51.70)	(0.134)	(0.000281)	(29.10)
Observations	460195	1407635	445189	460195	1407635	445189
Tract FE	Yes	Yes	Yes	Yes	Yes	Yes
County x Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Data Frequency	Quarterly	Monthly	Quarterly	Quarterly	Monthly	Quarterly
Cluster	Tract	Tract	Tract	Tract	Tract	Tract

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Loan-level regressions on NMDB data spanning 2016-2024. "Share $\geq$ 20% Decrease" is a continuous variable between 0 and 1, reporting the fraction of persisting flood insurance policies in a given census tract that experienced premium declines of at least 20 percent year-on-year across RR2's introduction. County $\times$ time fixed effects are monthly or quarterly as indicated by 'Data Frequency.' All regressions include county $\times$ time and time-invariant tract fixed effects; standard errors clustered at the tract level.

While existing policyholders are insulated from immediate financial distress, the glide path introduces a different distortion: homeowner lock-in. Properties in high cap-share areas became less likely to transact after the reform (Table 12). The estimated effect is economically meaningful: a property in a tract with a 10 percentage point higher cap share is roughly 3 basis points less likely to change hands in a given year than a comparable property in the same county. This lock-in effect persists when controlling for property characteristics such as age, size, and structure type. The mechanism is straightforward: existing homeowners face capped premiums that rise at most 18 percent annually, while new buyers would pay full risk-based rates. Moving destroys this option value, reducing mobility among incumbent homeowners in high-risk areas.

4.6 Investment

Thus far, we have examined the effect of RR2 on the existing stock of houses. We now turn to new residential construction. If the repricing corrects inefficient capital allocation driven by pre-reform subsidization, building permit issuance should decline in high cap-share areas. Developers face higher insurance costs that reduce buyer willingness to pay, while the reform's structure—capped premiums for existing homes but full-risk pricing for new construction—creates a price wedge that disadvantages new builds. We test this using building permits for single-family homes issued between 2016 and 2024 in jurisdictions covered by Cotality.[6]

Our specification is a Poisson regression of the following form:

$$\text{Permits}_{tract,t} = \beta \times \mathit{post} \times \text{cap share} + \alpha_{geo \times t} + \alpha_{tract} + \varepsilon_{tract,t} \tag{2}$$

[6] We retain only the most recent entry for each permit with positive approval status, dropping duplicates from multiple contractors and permits that were later revoked or canceled. Permit data is aggregated to the census tract and quarter level and merged with our tract-level flood insurance statistics, yielding an unbalanced panel of 658,000 tract-quarter observations across 13,449 census tracts in 49 states.

Table 12: Property transaction probability and insurance lock-in

	Transacted		
	(1)	(2)	(3)
Post × Share Capped	-0.00342***	-0.00316***	-0.00315***
	(0.000310)	(0.000350)	(0.000350)
Age of home (years)		-0.0000686***	-0.0000492***
		(0.00000235)	(0.00000243)
ln(Parcel acreage)			-0.00596***
			(0.0000879)
ln(Square footage)			0.00217***
			(0.000454)
Number of bedrooms			-0.000106
			(0.000138)
Number of bathrooms			0.000937***
			(0.000217)
Number of stories			0.000747***
			(0.000215)
Constant	0.0461***	0.0501***	0.0235***
	(0.0000510)	(0.000117)	(0.00264)
Observations	262818450	171002885	171002885
Adjusted R^2	0.0158	0.00723	0.00773
Property FE	Yes	No	No
Tract FE	No	Yes	Yes
County × Year FE	Yes	Yes	Yes
Cluster	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Regressions on a property-year panel expanded from Cotality data, spanning 2016-2024. The outcome variable, "Transacted," takes the value of 1 if a property is recorded as having transferred ownership in a given year, and 0 otherwise. All regressions include county × year fixed effects; column (1) includes time-invariant property fixed effects, and columns (2) and (3) include time-invariant tract fixed effects; standard errors are clustered at the tract level.

where "Permits$_{tract,t}$" is the number of new construction building permits for single-family homes that were issued in a given census tract and quarter; $\alpha_{geo \times t}$ is either state-quarter or county-quarter fixed effects; and standard errors are clustered at the census tract level. The beta coefficient, when exponentiated, measures the incidence rate ratio for building permit issuance in a census tract with a higher cap share after RR2's implementation.

Table 13: New construction permits fell in exposed tracts

	New Construction Permits	
	(1)	(2)
Post × Share Capped	-0.198***	-0.198***
	(0.0706)	(0.0733)
Constant	3.442***	3.474***
	(0.0108)	(0.0112)
Observations	313593	293750
Sample	All Tracts	Non-coastal Tracts
Pseudo R^2	0.723	0.725
Tract FE	Yes	Yes
County × Quarter FE	Yes	Yes
Cluster	Tract	Tract

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Poisson regressions on Cotality Building Permits data collapsed to tract-quarter aggregates. The permit sample is all approved, issued, final, or complete permits for new construction on single-family home properties from 2016 to 2024. Coastal census tracts are removed from the sample in column (2). Both regressions include county × quarter and time-invariant tract fixed effects; standard errors are clustered at the tract level.

We find that a census tract with a 100 percent cap share saw a nearly 20 percent decrease in the issuance of building permits for the new construction of single-family homes after the reform, compared to a census tract in the same county with a 0 percent cap share (Table

13, column (1)). This result is robust to the removal of coastal census tracts (column (2)), which are subject to idiosyncratic demand forces around vacation properties and amenity value.

4.7 Robustness

Finally, we conduct several robustness checks on our results thus far, using alternate data sources, treatment measures, and estimators. Our findings are confirmed by each test.

Using cHMDA for main specifications: To validate our choice of the NMDB as our primary analysis data set, we replicate our central results with HMDA data. We utilize a confidential vintage of HMDA which includes more precise information on loan and applicant characteristics, and certain fields such as borrower credit score which are removed from the public vintage. Our cHMDA headline analysis sample includes over 25 million successful conforming loan originations between 2018 and 2024. As with the NMDB, we retain only cHMDA loans in census tracts with at least eight NFIP policies in force during January 2021. cHMDA data prior to 2022 uses the 2010 Census definitions for reporting census tracts; we use a Census-provided land area crosswalk from 2010 to 2020 tract definitions to accurately match derived NFIP policy statistics over time. We further discard loans in 2020-definition census tracts which did not have a majority 'feeder' 2010 tract (e.g. if the land area of 2020 tract 005 was split equally between 2010 tracts 002, 003, and 004).

Figure 3: Evolution of new mortgage loan originations using HMDA data

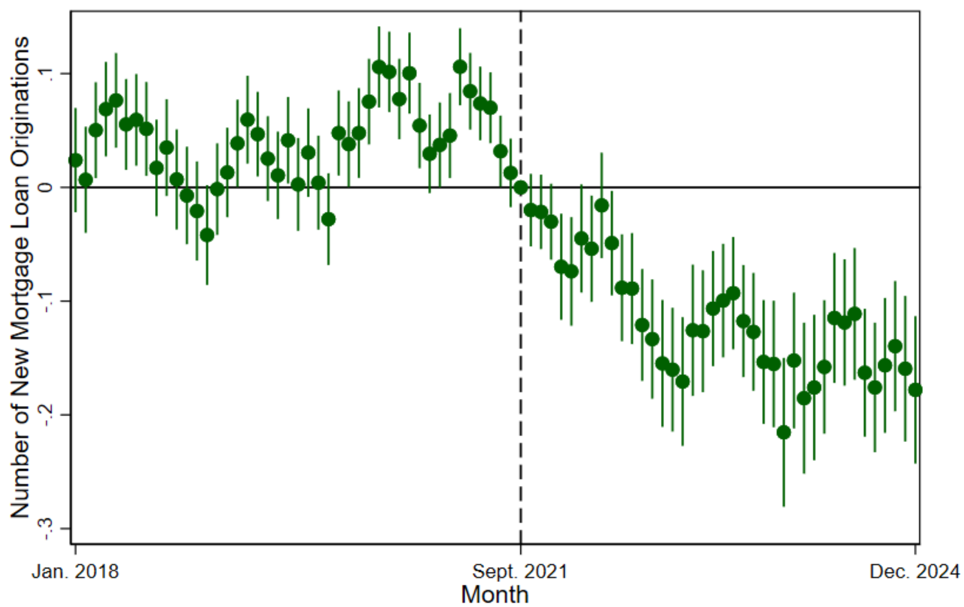

Note: Data are from cHMDA and span 2018-2024. The sample includes all new home purchase mortgage loans successfully originated in census tracts with a nonmissing cap share value, aggregated to the tract-month level. Outcome variable: the number of new mortgage loan originations in a given tract and month. Explanatory variable: post × cap share. The dashed line denotes the last quarter before the implementation of Risk Rating 2.0 for new policies. The regression includes county × month fixed effects and time-invariant tract fixed effects; standard errors clustered at the tract level.

Replication of our headline specifications with cHMDA data (Figure 3, Table 14) produces qualitatively and quantitatively similar results to those obtained with the NMDB. Census tracts with a higher cap share experience fewer new home purchase loans than other tracts with relatively lower cap shares, persisting through the incorporation of state-month and county-month fixed effects. Classifying loan time periods by application date instead of action date does not affect the results (Appendix Table A4, Appendix Figure A1).

Alternate treatment measures: To confirm the validity of cap share as our main treatment

Table 14: New mortgage loan originations using HMDA data

	New Mortgage Loan Originations		
	(1)	(2)	(3)
Post × Share Capped	-0.214***	-0.202***	-0.154***
	(0.0159)	(0.0176)	(0.0187)
Constant	3.279***	3.279***	3.282***
	(0.00318)	(0.00351)	(0.00374)
Observations	723782	723782	710830
Pseudo R^2	0.661	0.665	0.678
Tract FE	Yes	Yes	Yes
Quarter FE	Yes	N/A	N/A
State × Quarter FE	No	Yes	No
County × Quarter FE	No	No	Yes
Cluster	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Data are from cHMDA and span 2018-2024. The sample includes all new home purchase mortgage loans successfully originated in census tracts with a nonmissing cap share value, aggregated to the tract-month level. Outcome variable: the number of new mortgage loan originations in a given tract and month. Explanatory variable: post × cap share. Regressions include progressively more stringent geography × time fixed effects and time-invariant tract fixed effects; standard errors clustered at the tract level.

measure, we construct two alternate treatment measures at the tract level: the difference in the logarithm of the average (median) cost of coverage before and after RR2. While cap share is constructed by observing the set of policies on single-family homes that persisted *through* the policy change, these alternate measures compare the non-symmetric sets of SFR policies that were renewed between October 1, 2020 and September 30, 2021 (the pre-period) and the set of new SFR policies that were originated between October 1, 2021 and June 1, 2024 (the post-period), taking only the first policy-year price of the latter group. We align these comparison metrics by retaining the fixed $193 'expense and loss constant' and variable 'increased cost of compliance premium' in the new premium pricing structure, by adjusting coverage amounts for deductibles, and by standardizing costs to reflect paid rates per $1,000 of coverage.

Cap share is highly correlated with both of these alternate treatment intensity measures across the full sample, and the relationship is modestly stronger for those tracts that are within coastal counties (Appendix Table A6). Results are very similar when substituting these alternate measures into our headline and intensive margin analysis in place of cap share (Appendix Tables A7, A8, A9).

Alternate estimators: In recent years, numerous researchers have documented that standard two-way fixed-effects regressions can produce unreliable results when treatment exposure is heterogeneous, whether in terms of intensity between groups or the timing of application.[7] These lines of papers have proposed new tests and estimators to handle such situations. We show that our headline results are robust to the use of multiple alternate estimators proposed by de Chaisemartin et al. (2025) and de Chaisemartin et al. (2024): one designed to handle differences-in-differences estimations without strictly untreated units, and another with continuous rather than binary treatment exposure. We utilize cHMDA for these tests, collapsing new mortgage loan originations to a two-period tract-level structure with symmetrical 36-month periods around the implementation of RR2. This allows us to avoid the complication of interpreting the treatment as continuing at one level over multiple periods, instead isolating heterogeneous treatment exposures in an otherwise classical two-period DiD setup.

We consider two circumstances under which these alternate estimators may apply:

- The subset of census tracts with at least eight flood policies in 2021, and therefore with nonmissing cap share values, can be considered to be a group with no strictly untreated units, but with "quasi-stayers," tracts which had a cap share of zero. We utilize the "did_had" estimation command for this sample, which also allows for the treatment effect to vary between units.

[7] See, for example, Borusyak et al. (2024), Goodman-Bacon (2021), de Chaisemartin and D'Haultfœuille (2020), Callaway and Sant'Anna (2021), Sun and Abraham (2021).

- By taking as our sample all census tracts and by recoding missing cap share values to zero, we can consider those census tracts without cap share as "untreated," and utilize the "did_multiplegt_stat" estimation command for a static design with heterogeneous treatment exposure in the form of a continuous treatment variable.

Under both circumstances, we adjust the cap share value for those tracts with an original cap share of zero upwards to 0.0001. In the former case, "did_had" only admits strictly positive treatment values; in the latter case, we make the adjustment to differentiate between meaningful zeros representing tracts with flood policies which are all uncapped, and defined zeros representing tracts which are originally missing cap share. Both alternate estimators produce similar results to a naive two-way fixed effects design on the same sample (Table 15).

Table 15: New loan originations using alternative estimators

	New Mortgage Loan Originations		
	(1)	(2)	(3)
	Standard TWFE	HAD Quasi-Stayers	Static; Heterogeneous
Post × Share Capped	-64.143***	-65.495***	
	(3.789)	(4.374)	
Post × Share Capped;			-58.745***
No cap share := 0			(3.155)
Constant	225.978***	N/A	N/A
	(0.875)		
N	49300	49300	143718
Estimation Command	reghdfe	did_had	did_multiplegt_stat
Adjusted R^2	0.7528	N/A	N/A
Tract FE	Yes	N/A	N/A
Period FE	Yes	N/A	N/A
Cluster	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Regressions run on cHMDA data collapsed to symmetrical 36-month periods around treatment, and census tract level. The 'pre-period' is October 2018 through September 2021, and the 'post-period' is October 2021 through September 2024. In columns (2) and (3), tracts with cap share = 0 have their cap share values adjusted to 0.0001 to distinguish between treatment effect zeroes and untreated zeroes. The coefficients reported in columns (2) and (3) are the weighted average slope (WAS) parameters as described in de Chaisemartin et al. (2025) and de Chaisemartin et al. (2024).

5 Conclusion

This paper asks how markets respond when the price of a risk changes while the underlying hazard does not. The 2021 RR2 reform offers a clean setting to test this question. RR2 replaced FEMA's coarse zone-based formula with property-level pricing, generating large differences in premium *changes* across neighboring properties facing identical flood risk. Using a tract-level cap share measure of pre-reform subsidization, we find that, relative to a tract in the same county where no policies are capped, a fully capped tract sees roughly 9 percent fewer mortgage originations after the reform, prices about 1.8 percent lower, new borrowers with credit scores about three points lower and more first-time and owner-occupant buyers, and single-family construction permits down about 18 percent. These effects emerge immediately, show no differential pre-trends, and survive confidential HMDA data, alternate treatment measures, and heterogeneous-treatment estimators; because flood maps were unchanged, they reflect the price of a known risk, not revised beliefs. Two mechanisms drive them. On the supply side, adverse application outcomes rise even within the same lender and tract, concentrated among depository institutions with high pre-reform exposure while comparably exposed non-depository lenders do not retrench—consistent with balance-sheet constraints binding more for portfolio lenders than originate-to-distribute ones. On the demand side, the 18 percent cap shields incumbents while new buyers pay full rates, so moving forfeits the option value of a grandfathered policy.

These findings speak directly to NFIP reform. The cap, meant to protect affordability, delays full actuarial pricing for the most subsidized properties by a decade or more and concentrates credit-constrained borrowers in high-risk areas—borrowers who, as Rickard (2025) notes, are less likely to adapt their properties, raising future losses for households and the NFIP alike. Because construction responds at the tract but not the county level, the reform appears to redirect residential investment toward safer locations rather than eliminate it, leaving policymakers to weigh faster repricing against these distributional and lock-in costs. These lessons reach well beyond the NFIP: wherever insurers reprice previously mispriced hazards such as wildfire and wind, the same mechanisms should emerge, making continued monitoring essential (Sastry et al., 2024).

References

Aiba, Ikuto, Daisuke Hasegawa, and Hiroaki Shirai, "The Effects of Flood Risk Mandatory Disclosure on Housing Markets," Working Paper, Available at SSRN 2025.

Amornsiripanitch, Natee, Siddhartha Biswas, John Orellana-Li, and David Zink, "Measuring flood underinsurance in the USA," *Nature Climate Change*, 2025, pp. 1–7.

An, Xudong, Yongheng Deng, and Dayin Zhang, "Imperfect Flood Insurance Enforcement and Business Misallocation," Working Paper, Available at SSRN 2023.

Arseneau, David M and Gazi Kara, "Do Banks Price Flood Risk in Mortgage Origination: Evidence from a Natural Experiment in New Orleans," Working Paper 2025-081, Federal Reserve Board of Governors 2025.

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

Billings, Stephen B., Sophie Calder-Wang, and Weiling Liu, "The price and distributional impact of flood risk disclosure: Evidence from US housing platforms," Working Paper, Available at SSRN 2025.

Blickle, Kristian and Joao AC Santos, "Unintended consequences of" mandatory" flood insurance," Staff Report 1012, Federal Reserve Bank of New York 2022.

_ , **Evan Perry, and João AC Santos**, "Do mortgage lenders respond to flood risk?," Staff Report 1101, Federal Reserve Bank of New York 2024.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *The Review of Economic Studies*, 11 2024, *91* (6), 3253–3285.

Bradt, Jacob T, Carolyn Kousky, and Oliver EJ Wing, "Voluntary purchases and adverse selection in the market for flood insurance," *Journal of Environmental Economics and Management*, 2021, *110*, 102515.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230. Themed Issue: Treatment Effect 1.

Collier, Benjamin, Tobias Huber, Johannes Gerd Jaspersen, and Andreas Richter, "Homeowners' willingness to hedge flood risks as prices increase," Working Paper, Available at SSRN 2023.

de Chaisemartin, Clement, Diego Ciccia, Xavier D'Haultfoeuille, and Felix Knau, "Difference-in-Differences Estimators When No Unit Remains Untreated," Technical Report, Working Paper 2025.

__ , Xavier D'Haultfoeuille, Felix Pasquier, Doulo Sow, and Gonzalo Vazquez-Bare, "Difference-in-Differences for Continuous Treatments and Instruments with Stayers," Technical Report, Working Paper 2024.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, September 2020, *110* (9), 2964–96.

Doerner, William M, Michael J Seiler, and Matthew Suandi, "When risk does not discount: Flood history and rising property valuations," *Real Estate Economics*, 2025.

Eby, Matthew, "Understanding FEMA Flood Maps and Limitations," Technical Report, First Street Foundation 2019.

Fabian, Jacob, "The price of risk: Flood insurance premium reform and local development," Working Paper, Working paper 2024.

Fairweather, Daryl, Matthew E Kahn, Robert D Metcalfe, and Sebastian Sandoval Olascoaga, "Expecting climate change: A nationwide field experiment in the housing market," Working Paper 33119, National Bureau of Economic Research 2024.

GAO, "National Flood Insurance Program: Congress Should Consider Updating the Mandatory Purchase Requirement (GAO-21-578)," Technical Report, US Government Accountability Office 2021.

_ , "Flood Insurance: FEMA's New Rate-Setting Methodology Improves Actuarial Soundness but Highlights Need for Broader Program Reform (GAO-23-105977)," Technical Report, US Government Accountability Office 2023.

Ge, Shan, Ammon Lam, and Ryan Lewis, "The effect of insurance premiums on the housing market and climate risk pricing," Working Paper, Available at SSRN 2025a.

_ **, Stephanie Johnson, and Nitzan Tzur-Ilan**, "The Hidden Effects of Climate Risk: Rising Insurance Premiums Increase Mortgage Delinquency and Drive Relocation to Safer Areas," Working Paper, Available at SSRN 2025b.

Gillespie, Tom, Ronan C Lyons, and Thomas KJ McDermott, "Estimating the Flood Risk Discount: Evidence From a One-off National Information Shock," *Environmental and Resource Economics*, 2025, pp. 1–18.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277. Themed Issue: Treatment Effect 1.

Gourevitch, Jesse D, Carolyn Kousky, Yanjun Liao, Christoph Nolte, Adam B Pollack, Jeremy R Porter, and Joakim A Weill, "Unpriced climate risk and the potential consequences of overvaluation in US housing markets," *Nature climate change*, 2023, *13* (3), 250–257.

Hino, Miyuki and Marshall Burke, "Does information about climate risk affect property values?," Working Paper 26807, National Bureau of Economic Research 2020.

Indaco, Agustín, Francesc Ortega, and Süleyman Taşpınar, "The effects of flood insurance on housing markets," *Cityscape*, 2019, *21* (2), 129–156.

Keys, Benjamin J and Philip Mulder, "Neglected no more: Housing markets, mortgage lending, and sea level rise," Working Paper 27930, National Bureau of Economic Research 2020.

Michel-Kerjan, Erwann, Sabine Lemoyne de Forges, and Howard Kunreuther, "Policy tenure under the US national flood insurance program (NFIP)," *Risk Analysis: An International Journal*, 2012, *32* (4), 644–658.

Miller, Ryan G and Nicholas Pinter, "Flood risk and residential real-estate prices: evidence from three US counties," *Journal of Flood Risk Management*, 2022, *15* (2), e12774.

Mulder, Philip, "Mismeasuring risk: The welfare effects of flood risk information," Working Paper, Available at SSRN 2024.

— **and Carolyn Kousky**, "Risk Rating without Information Provision," *AEA Papers and Proceedings*, 05 2023, *113*, 299–303.

Niu, Dongxiao, Nils Kok, Juan Palacios, and Siqi Zheng, "Climate Risk and Collateral Misreporting," Technical Report 25/01, Massachusetts Institute of Technology Center for Real Estate 2025.

Ortega, Francesc and Ivan Petkov, "To improve is to change? the effects of risk rating 2.0 on flood insurance demand," Working Paper 17021, IZA Institute of Labor Economics 2024.

Ostriker, Abigail and Anna Russo, "The effects of floodplain regulation on housing markets," Working Paper, Working paper 2024.

Peralta, Abigail and Jonathan B Scott, "Does the National Flood Insurance Program drive migration to higher risk areas?," *Journal of the Association of Environmental and Resource Economists*, 2024, *11* (2), 287–318.

Pollack, Adam B, Douglas H Wrenn, Christoph Nolte, and Ian Sue Wing, "Potential benefits in remapping the special flood hazard area: evidence from the US housing market," *Journal of Housing Economics*, 2023, *61*, 101956.

Rickard, Natalie, "Sheltering from Climate Risks," Technical Report, JMP 2025.

Sastry, Parinitha, Ishita Sen, Ana-Maria Tenekedjieva, and Therese C Scharlemann, "The limits of insurance demand and the growing protection gap," Working Paper, Available at SSRN 2024.

Skouralis, Alexandros, Nicole Lux, and Mark Andrew, "Does flood risk affect property prices? Evidence from a property-level flood score," *Journal of Housing Economics*, 2024, *66*, 102027.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199. Themed Issue: Treatment Effect 1.

Sun, Pin, Douglas S Noonan, and Lilliard Richardson, "The Effect of Flood Zoning Policies on Housing Markets: Evidence from Texas," Working Paper, Available at SSRN 2022.

Zhang, Dayin, "How Are Flood Risks Managed in the United States?," *Oxford Research Encyclopedia of Economics and Finance*, 2023.

A Online Appendix

Figure A1: Evolution of loan originations using application date

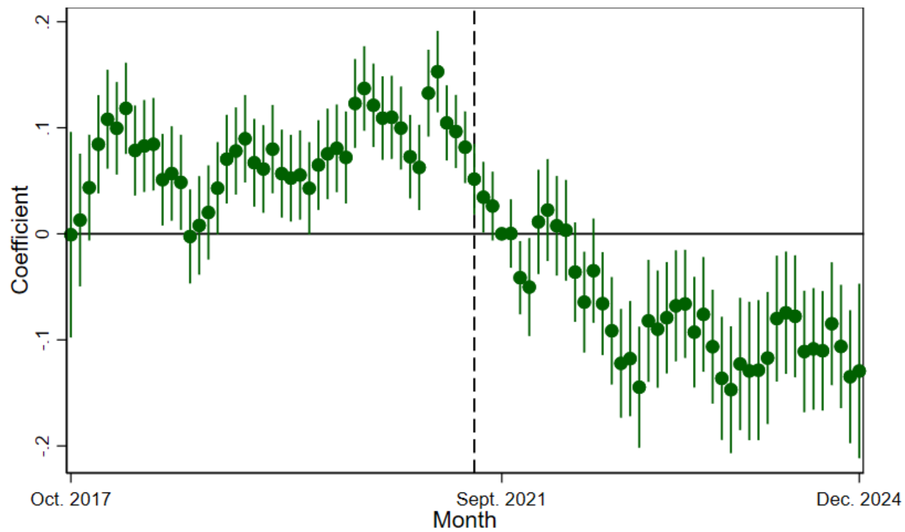

Note: Data are from cHMDA and span 2018-2024. The sample includes all new home purchase mortgage loans successfully originated in census tracts with a nonmissing cap share value, aggregated to the tract-month level. In this figure, mortgages are assigned to the month of application submission, rather than application decision. Outcome variable: the number of new mortgage loan originations in a given tract and month. Explanatory variable: post $\times$ cap share. The dashed line denotes the last quarter before the implementation of Risk Rating 2.0 for new policies. The regression includes county $\times$ month fixed effects and time-invariant tract fixed effects; standard errors clustered at the tract level.

Table A1: Descriptive statistics for unbalanced flood policy sample

	Pre-reform		Post-reform	
	Mean	Median	Mean	Median
Insurance premium (annual)	587.83	395.00	686.96	562.00
Total coverage	306289.21	350000.00	311934.37	350000.00
Cost per $1000 coverage	2.29	1.26	2.50	1.75
Building replacement cost	335908.60	250000.00	492066.78	250000.00
	Pre-reform		Post-reform	
Located in SFHA		0.44		0.43
Mandatory purchase		0.12		0.16
Primary residence		0.86		0.87
In CRS community		0.72		0.72
N		10099221.00 (74.11%)		3527223.00 (25.89%)

Note: Data from the Federal Reserve Bank of Philadelphia's vintage of the NFIP Redacted Policies set, at the policy-year level. The policy sample is unbalanced across the pre- and post-reform periods, but includes only residential policies for single-family homes in tracts with at least eight policies in force as of January 2021. SFHA: Special Flood Hazard Area. CRS: Community Rating System. "Pre-reform" covers April 2015 through March 2022; "Post-reform" covers April 2022 - March 2025.

Table A2: House prices for home sales with associated mortgages

	ln(Sale Amount)					
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Share Capped	-0.0176***	-0.0180***	-0.0185***	-0.0180***	-0.0162***	-0.0162***
	(0.00264)	(0.00258)	(0.00248)	(0.00245)	(0.00230)	(0.00230)
Age of home (years)		-0.00529***	-0.00549***	-0.00507***	-0.00267***	-0.00267***
		(0.0000461)	(0.0000509)	(0.000151)	(0.0000328)	(0.0000328)
ln(Parcel acreage)			0.173***	0.166***	0.0811***	0.0810***
			(0.00210)	(0.00334)	(0.00117)	(0.00117)
Number of bedrooms				-0.000841	-0.00475***	-0.00476***
				(0.00672)	(0.00172)	(0.00173)
Number of bathrooms				0.0331***	0.00976***	0.00979***
				(0.01000)	(0.00246)	(0.00247)
ln(Square footage)					0.623***	0.624***
					(0.00418)	(0.00408)
Number of stories						-0.00185*
						(0.000954)
Constant	12.68***	12.88***	13.11***	13.00***	8.174***	8.168***
	(0.000384)	(0.00182)	(0.00351)	(0.0462)	(0.0268)	(0.0263)
Observations	6243544	6243544	6243544	6243544	6243544	6243544
Adjusted R^2	0.708	0.732	0.760	0.768	0.834	0.834
Tract FE	Yes	Yes	Yes	Yes	Yes	Yes
County × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Cluster	Tract	Tract	Tract	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Data from the Cotality Owner Transfer set. The sample spans 2016-2024, and includes all arm's-length transactions which have an associated mortgage. Only properties with the full slate of relevant characteristics are included in all samples for consistent comparison. The outcome variable is the log sale amount for the transaction. Columns progressively add controls for property characteristics (home age, parcel acreage, bedrooms, bathrooms, square footage, number of stories). All regressions include tract and county × quarter fixed effects; standard errors clustered at the tract level.

Table A3: Characteristics of transacted homes

	(1)	(2)	(3)	(4)	(5)	(6)
	Home age	ln(Parcel acreage)	Bedrooms	Bathrooms	ln(Square footage)	Number of stories
Post × Share Capped	0.0229***	0.00355	-0.00300	-0.00496*	-0.00563***	-0.00355*
	(0.00311)	(0.00376)	(0.00244)	(0.00195)	(0.00151)	(0.00161)
Constant	3.828***	-1.242***	1.204***	0.937***	7.533***	0.324***
	(0.000471)	(0.000567)	(0.000367)	(0.000295)	(0.000228)	(0.000242)
Observations	8238575	8238585	8238585	8238585	8238585	8238585
Adjusted R^2		0.497			0.393	
Pseudo R^2	0.432		0.0196	0.0508		0.0218
Regression Type	Poisson	OLS	Poisson	Poisson	OLS	Poisson
Tract FE	Yes	Yes	Yes	Yes	Yes	Yes
County × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Cluster	Tract	Tract	Tract	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Data from the Cotality Owner Transfer set. The sample spans 2016 to 2024, and includes all arm's-length transactions regardless of whether or not sales have an associated mortgage. Only properties with the full slate of relevant characteristics are included in all samples for consistent comparison. The outcome variable is the log sale amount for the transaction. Columns progressively add controls for property characteristics (home age, parcel acreage, bedrooms, bathrooms, square footage, number of stories). Regression type (OLS or Poisson) is reported for each specification. All regressions include tract and county × quarter fixed effects; standard errors clustered at the tract level.

Table A4: Loan originations using application date

	New Mortgage Loan Originations		
	(1)	(2)	(3)
Post × Share Capped	-0.191***	-0.192***	-0.149***
	(0.0163)	(0.0180)	(0.0192)
Constant	3.237***	3.240***	3.244***
	(0.00304)	(0.00336)	(0.00359)
Observations	759711	759711	745859
Pseudo R^2	0.658	0.662	0.674
Tract FE	Yes	Yes	Yes
Quarter FE	Yes	N/A	N/A
State × Quarter FE	No	Yes	No
County × Quarter FE	No	No	Yes
Cluster	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Data are from cHMDA and span 2018-2024. The sample includes all new home purchase mortgage loans successfully originated in census tracts with a nonmissing cap share value, aggregated to the tract-quarter level. In this table, mortgages are assigned to the quarter of application submission, rather than application decision. Outcome variable: the number of new mortgage loan originations in a given tract and quarter. Explanatory variable: post × cap share. Regressions include progressively more stringent geography × quarter fixed effects and time-invariant tract fixed effects; standard errors clustered at the tract level.

Table A5: Loan denial *counts*

	New Mortgage Loan Denials		
	(1)	(2)	(3)
Post × Share Capped	-0.00144	-0.0802***	-0.0770***
	(0.0229)	(0.0247)	(0.0218)
Constant	1.274***	1.297***	1.528***
	(0.00512)	(0.00551)	(0.00488)
Observations	802760	801372	526118
Pseudo R^2	0.368	0.371	0.345
Tract FE	Yes	Yes	Yes
Quarter FE	Yes	N/A	N/A
State × Quarter FE	No	Yes	No
County × Quarter FE	No	No	Yes
Cluster	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Data are from cHMDA and span 2018-2024. The sample includes all new home purchase mortgage loan applications that were denied, from census tracts with a nonmissing cap share value. Data are aggregated to the tract-quarter level. A denial can be issued by a lender at the pre-approval or standard approval stages. Outcome variable: the number of denied new mortgage applications in a given tract and quarter. Explanatory variable: post × cap share. Regressions include progressively more stringent geography × quarter fixed effects and time-invariant tract fixed effects; standard errors clustered at the tract level.

Table A6: Alternate treatment measure correlations with cap share

Change in 'full risk' cost of $1000 of coverage	Full sample	Tracts in coastal counties	Tracts in non-coastal counties
Average	0.7125	0.7092	0.6699
Median	0.6576	0.6349	0.6407

Note: Alternate treatment measures are constructed as described in Section (4.7). Data are at the tract level and calculated from the NFIP Redacted Policies set.

Table A7: Headline regression with alternate treatment measures

	(1) Total Loans	(2) Total Loans	(3) Total Loans	(4) Total Loans
Post × average cost of coverage increase	-0.0987***			
	(0.0173)			
Post × median cost of coverage increase		-0.0851***		
		(0.0209)		
Post × Δ average cost of coverage			-0.0973***	
			(0.0151)	
Post × Δ median cost of coverage				-0.0675***
				(0.0124)
Constant	2.408***	2.407***	2.387***	2.391***
	(0.00655)	(0.00441)	(0.000817)	(0.00168)
Observations	877892	877892	877892	877892
Pseudo R^2	0.543	0.543	0.543	0.543
Tract FE	Yes	Yes	Yes	Yes
County × Month FE	Yes	Yes	Yes	Yes
Cluster	Tract	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Data are from cHMDA and span 2018-2024. The sample includes all new home purchase mortgage loans successfully originated in census tracts with a nonmissing cap share value, aggregated to the tract-month level. For the first two rows, the explanatory variable is the interaction of a post dummy and dummies taking the value of 1 if the observed average (median) cost of flood coverage in a given tract rose year-on-year across the implementation of RR2, and 0 otherwise. For the third and fourth rows, the explanatory variable is the interaction of a post dummy and the continuous alternate treatment measures described in Section (4.7). Regressions include county × month fixed effects and time-invariant tract fixed effects; standard errors clustered at the tract level.

Table A8: Characteristics of new mortgage loans with median coverage cost change as explanatory variable

	(1) Credit Score	(2) Owner-Occupied	(3) First-time Owner	(4) DTI	(5) Interest Rate
Post $\times$ Δ median cost of coverage	-1.663***	0.00431***	0.00261	0.0261	0.0125***
	(0.340)	(0.00158)	(0.00268)	(0.0562)	(0.00482)
Constant	739.9***	0.912***	0.426***	37.44***	4.462***
	(0.0358)	(0.000166)	(0.000283)	(0.00592)	(0.000508)
Observations	506735	506735	506735	506729	506735
Adjusted R^2	0.143	0.112	0.109	0.074	0.723
Tract FE	Yes	Yes	Yes	Yes	Yes
County $\times$ Month FE	Yes	Yes	Yes	Yes	Yes
Cluster	Tract	Tract	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Data are from the NMDB and span 2016-2024. The outcome variable for columns (2) and (3) is a dummy variable indicating the status of the given characteristic for each new loan. The explanatory variable is the log-difference in *median* flood insurance coverage cost per $1,000 of building coverage between pre-reform (October 2020–September 2021) and post-reform (October 2021–June 2024) policy cohorts. All regressions include tract and county $\times$ month fixed effects; standard errors clustered at the tract level.

Table A9: Characteristics of new mortgage loans with average coverage cost change as explanatory variable

	(1) Credit Score	(2) Owner-Occupied	(3) First-time Owner	(4) DTI	(5) Interest Rate
Post $\times$ Δ average cost of coverage	-1.917***	0.00215	0.00486*	-0.0476	0.0116**
	(0.357)	(0.00161)	(0.00279)	(0.0591)	(0.00518)
Constant	739.8***	0.913***	0.426***	37.44***	4.463***
	(0.00749)	(0.0000339)	(0.0000585)	(0.00124)	(0.000109)
Observations	506735	506735	506735	506729	506735
Adjusted R^2	0.143	0.112	0.109	0.074	0.723
Tract FE	Yes	Yes	Yes	Yes	Yes
County $\times$ Month FE	Yes	Yes	Yes	Yes	Yes
Cluster	Tract	Tract	Tract	Tract	Tract

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: Data are from the NMDB and span 2016-2024. The outcome variable for columns (2) and (3) is a dummy variable indicating the status of the given characteristic for each new loan. The explanatory variable is the log-difference in *average* flood insurance coverage cost per $1,000 of building coverage between pre-reform (October 2020–September 2021) and post-reform (October 2021–June 2024) policy cohorts. All regressions include tract and county $\times$ month fixed effects; standard errors are clustered at the tract level.