

Flood Exposure and Asset Repricing: Evidence from FEMA LOMRs and ZIP-Code Home Values

Dartmouth College
Topics in Money and Finance

Cameron Keith*

March 15, 2026

Abstract

Do coastal housing markets reprice properties when official flood-map designations change? I study this question using FEMA Letters of Map Revision (LOMRs), which are localized revisions to National Flood Insurance Program (NFIP) flood maps. In a quarterly ZIP-level event study across U.S. coastal counties from 2009 to 2022, with ZIP and county-by-year fixed effects, I find that home values show near-zero response at impact but decline by about 2.8 percent four or more years after a LOMR. The decline is larger in an interaction specification with pre-treatment flood-insurance penetration. An exploratory political interaction suggests weaker repricing in Republican-leaning counties, although this comparison does not pass the pre-trend test. These findings are consistent with slow, uneven capitalization of regulatory flood-risk information rather than immediate repricing.

Keywords: Flood risk; property values; FEMA; capitalization; flood insurance; event study

JEL Codes: R31; Q54; G22; H84

*I thank Professor Apoorv Gupta, Professor of Economics at Dartmouth College, for his guidance and helpful comments on this project. I used Claude Opus 4.6 to assist with code development as well as grammar and spelling checks. All analysis, interpretation, and errors are my own.
Email: cameron.s.keith.26@dartmouth.edu.

1 Introduction

Flood risk affects housing markets through official designations and realized disasters. Federal Emergency Management Agency (FEMA) flood maps assign properties to regulatory flood-risk zones, which can affect insurance requirements, disclosure, underwriting, and buyer beliefs. The most economically salient threshold is often the Special Flood Hazard Area (SFHA), but Letters of Map Revision (LOMRs) can also matter more broadly because they update the official map that private actors use in transactions. I study whether housing markets reprice when those official designations change through FEMA LOMRs.

These stakes are large because the NFIP is both economically important and fiscally strained. As of 2025, the program covers more than 4.7 million policies with roughly \$1.3 trillion in insured value, yet it still carries \$22.5 billion in debt to the U.S. Treasury after Congress canceled \$16 billion of NFIP debt in 2017 ([Congressional Research Service, 2025a,b](#)). Prior work also suggests that flood risk is not fully incorporated into housing prices: [Hino and Burke \(2021\)](#) estimate that floodplain properties are collectively overvalued by about \$44 billion and that rezoning into the floodplain is associated with roughly a 2 percent price discount.

These stakes make it economically important to understand how housing markets respond when FEMA revises the official map itself. To study that question, I construct a quarterly ZIP-level panel for coastal U.S. counties from 2009 to 2022 that combines Zillow home values with FEMA LOMRs, National Flood Insurance Program (NFIP) data, and county controls ([Federal Emergency Management Agency, 2024a,b](#); [Zillow Research, 2024](#)). A further contribution is the construction of a linked panel that brings together home values, official flood-map revisions, insurance exposure, and local controls in one ZIP-quarter dataset. I run a staggered event study with ZIP and county-by-year fixed effects on first in-window LOMRs. I then interact treatment with pre-treatment NFIP policy penetration to

analyze whether repricing is larger where mapped flood risk is already economically relevant.

Home values do not fall sharply after a LOMR. In the baseline specification, the estimated effect is close to zero at treatment and is consistent with a decline of about 2.8 percent four or more years later. The response is stronger in the continuous insurance-intensity specification. An exploratory political split also suggests weaker repricing in ZIP codes located in more Republican-leaning counties, consistent with heterogeneity in attention to official flood-risk information.

I study four related hypotheses. First, LOMRs are associated with lower home values. Second, if housing markets adjust with frictions, that response should emerge gradually after treatment rather than on impact. Third, if flood insurance mainly shields owners from repricing, higher pre-treatment NFIP exposure should weaken the response, while if salience dominates it should strengthen it. Fourth, if political beliefs shape attention to official flood-risk information, repricing should be weaker in more Republican counties. I test the first two with the baseline event study, the third with the intensity and insurance-outcome specifications, and the fourth with a political interaction that I treat as exploratory because its pre-trends are weaker.

The paper contributes to the flood-risk literature in three ways. First, it studies localized official map revisions instead of realized flood events, broad climate-risk disclosures, or static cross-sectional floodplain comparisons. Among prior work, [Shr and Zipp \(2019\)](#) is closest in studying flood-zone remapping, while [Weill \(2023\)](#) studies FEMA map updates as information shocks in the insurance market. Second, the staggered ZIP-level panel allows me to study the timing of capitalization after a revision becomes effective, which helps distinguish gradual repricing from an immediate market-wide reset. Third, by interacting treatment with pre-treatment NFIP penetration, I test whether the same map change carries more pricing content where flood insurance was already more financially material.

[Section 2](#) describes the institutional setting, [Section 3](#) reviews the related literature, and

Section 4 and Section 5 present the data and design. Section 6 reports the main results, Section 7 presents exploratory heterogeneity and robustness checks, Section 8 discusses interpretation and limitations, and Section 9 provides areas for future research. Supporting insurance-market and subsample analyses appear in Appendix A.

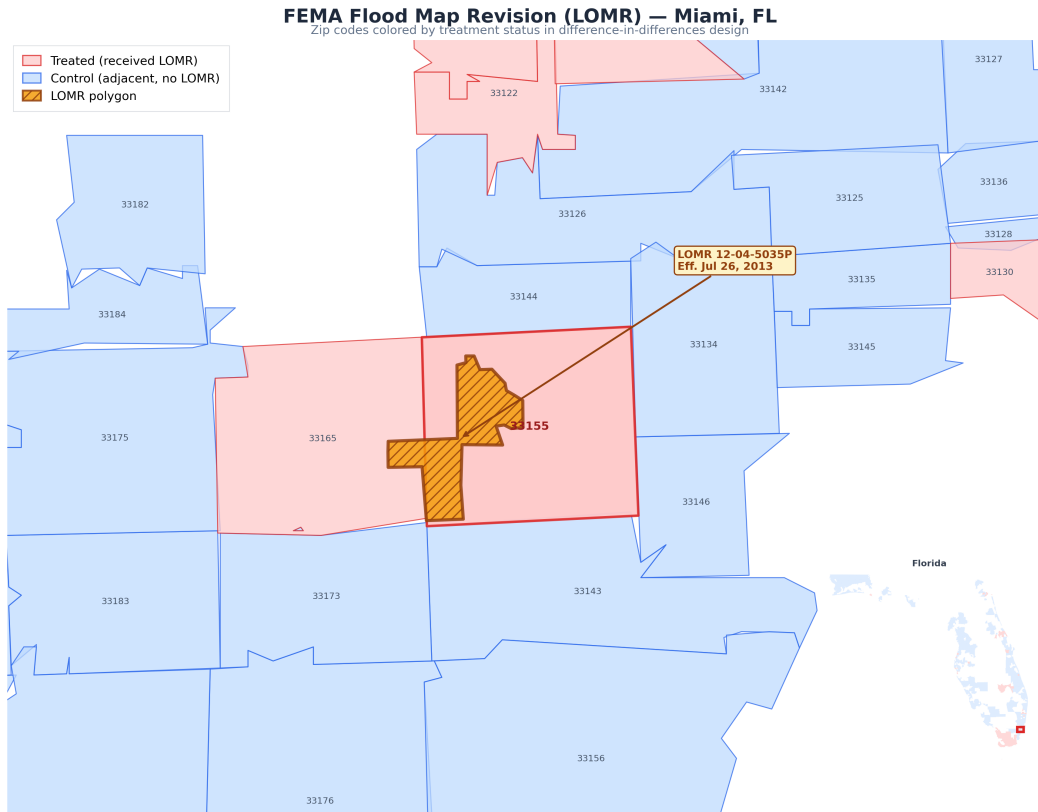
2 Institutional Background

2.1 FEMA Flood Maps and LOMRs

FEMA’s Flood Insurance Rate Maps (FIRMs) assign properties to regulatory flood-risk zones and provide related hazard information for lenders, insurers, local governments, and buyers. Letters of Map Revision (LOMRs) update those classifications after new topographic, hydrologic, engineering, or physical information. A LOMR can move land into or out of the Special Flood Hazard Area (SFHA), revise the location of the mapped floodplain boundary, or change other map-based flood hazard information that matters for underwriting, disclosure, and insurance. The SFHA remains the most consequential threshold because federally backed mortgages in participating NFIP communities generally require flood insurance for improved structures inside that zone. However, the broader point for this paper is that LOMRs change the regulatory baseline that lenders, insurers, and buyers rely on, even when the revision does not map cleanly into a simple up-or-down SFHA crossing or follow a recent flood that has already drawn attention to the area.

Figure 1 shows a LOMR in Miami. The example makes the treatment structure concrete: the LOMR polygon overlaps two ZIP codes.

Figure 1: Illustrative LOMR and Treated ZIP Code



Note: Deliberately large illustrative FEMA LOMR in Miami, Florida, chosen for visibility rather than representativeness. The shaded polygon is the revision footprint, the two ZIP codes (33165 and 33155) are treated in the event-study design, and adjacent ZIP codes remain untreated in that event window. The inset shows the location within Florida.

2.2 Why LOMRs Should Affect Housing Markets

I expect LOMRs to affect housing prices because they revise the regulatory designation that shapes expected ownership costs. Lenders, insurers, appraisers, and buyers all rely on the official map in transactions, so a revision can have pricing content even when no flood occurs nearby.

A lender's flood determination during underwriting may reflect the updated map, which can change insurance requirements, documentation, or the buyer's perception of risk before closing. Prices may then adjust as the new information diffuses through transactions, listings, appraisals, and comparable sales.

3 Related Literature

A large literature documents that flood risk is capitalized into property values, especially after important events. Floodplain discounts appear after hurricanes (Bin and Polasky, 2004), near-miss events (Hallstrom and Smith, 2005), and severe flood experience (Kousky, 2010), and a meta-analysis estimates an average inland floodplain discount of about 4.6 percent, though the same study finds that coastal floodplain properties can command a premium when amenities correlate with flood exposure (Beltrán et al., 2018). More recent work extends those findings to broader climate-risk signals: Hino and Burke (2021) show that regulatory flood-risk information can affect property values even without contemporaneous damage, and Bernstein et al. (2019) document discounts for sea-level-rise exposure, although Murfin and Spiegel (2020) provide more skeptical evidence.

Whether risk is fully capitalized depends in part on beliefs and attention. Bakkensen and Barrage (2022), Gibson and Mullins (2020), and Baldauf et al. (2020) emphasize that capitalization varies with local climate beliefs rather than with physical exposure alone. On the disclosure side, Troy and Romm (2004) and Pope (2008) find larger floodplain discounts under stronger disclosure rules. Weill (2023) is especially close in spirit because it studies the effects of official FEMA map updates on flood-insurance take-up. I complement that literature by asking whether localized official map revisions are also capitalized into housing values.

4 Data

4.1 Core Panel Structure

My main estimation sample is a ZIP-by-quarter panel for U.S. coastal ZIP codes from 2009 to 2022. I begin with a monthly panel, because I observe Zillow home values, NFIP

insurance records, and several controls at that frequency. Then, I collapse the data to quarters by averaging the monthly log values within each quarter, because the Zillow Home Value Index (ZHVI) is a three-month smoothed series and monthly observations inherit substantial serial correlation from that construction. After the sample restrictions described below, the quarterly panel contains 4,272 ZIP codes and 228,005 ZIP-quarter observations (228,002 in regressions after the fixed-effects estimator drops three singleton groups).

Those 4,272 ZIP codes include 398 single-LOMR treated ZIP codes within the estimation window, 3,711 never-treated ZIP codes, and 163 not-yet-treated ZIP codes whose first LOMR occurs after the sample window and that contribute control observations during the sample period.

4.2 Data Sources

The primary outcome is Zillow’s ZIP-level Home Value Index (ZHVI), using the all-homes (single-family residences plus condo/co-op), middle-tier, smoothed, seasonally adjusted series, which I deflate to December 2022 dollars using the Consumer Price Index for All Urban Consumers (CPI-U) ([Federal Reserve Bank of St. Louis, 2024](#)) and transform into $\ln(\text{real ZHVI})$ for the main specifications. Zillow’s coverage determines the first major sample screen because a nontrivial share of coastal ZIP codes does not have usable home-value data over the full period.

Treatment comes from FEMA’s S-LOMR polygon layer in the National Flood Hazard Layer ([Federal Emergency Management Agency, 2024a](#)). I spatially overlay those polygons onto Census ZCTA boundaries ([U.S. Census Bureau, 2024](#)) to identify whether a ZIP code is ever exposed to a LOMR and whether the first effective LOMR falls inside the 2009–2022 analysis window. The overlay supports the geographic-intensity extension directly. For the zoning-direction analysis, I use post-LOMR changes in SFHA share, defined as the fraction

of a ZIP code’s NFIP policies whose rated flood zone falls in the SFHA (A or V zones) in a given month, as a proxy for whether local flood exposure shifts upward or downward.

I merge those treatment data with OpenFEMA NFIP policy and claims records ([Federal Emergency Management Agency, 2024b](#)), which I aggregate to the ZIP-quarter level. The insurance data provide policy counts, average premiums, the share of policies rated in SFHA zones, and claims activity. Those variables serve two purposes. They measure pre-treatment exposure to flood insurance, and they also provide outcomes for the mechanism tests. From the policy counts I also construct *policy intensity*, defined as the pre-LOMR average NFIP policy count divided by ZIP population, which serves as the main treatment-interaction variable in the intensity specification.

I add county-level unemployment rates from BLS Local Area Unemployment Statistics (LAUS) ([U.S. Bureau of Labor Statistics, 2024](#)) as a local macro control, and I merge county presidential vote returns from the MIT Election Data and Science Lab ([MIT Election Data and Science Lab, 2017](#)) to build the political heterogeneity measure. Finally, I hand-code state flood-disclosure regimes into strict and broad indicators. That combination gives me a panel that links local home values, official flood-map revisions, insurance exposure, and institutional context in a common ZIP-quarter frame.

4.3 Sample Construction

[Table 1](#) summarizes the sample construction. I start from 33,782 U.S. ZIP codes, restrict to 5,450 ZIP codes in NOAA coastal counties, and then merge Zillow coverage, which leaves 4,836 coastal ZIP codes with usable home-value data. The LOMR–ZCTA overlay identifies 1,125 ever-treated ZIP codes in that sample, including 855 whose first observed LOMR falls inside the 2009–2022 window.

Then, I impose the event-study restrictions. First, I drop ZIP codes that were already

treated before 2009 because they do not have a clean pre-period. Second, I drop ZIP codes with multiple LOMRs because overlapping revisions make event time difficult to interpret. I also winsorize $\ln(\text{real ZHVI})$ at the 1st and 99th percentiles to limit the influence of extreme observations. After these restrictions, the final regression panel contains 4,272 ZIP codes. I impose these restrictions because clean treatment timing is worth more than the additional observations those ZIP codes would contribute.

4.4 Summary Statistics

[Table 2](#) reports descriptive moments for the estimation sample. The treatment variables are highly skewed: the mean policy-intensity measure is 5.93 policies per 1,000 residents, while the median is only 0.68. Most coastal ZIP codes have limited flood-insurance exposure, while a smaller subset is much more exposed. That pattern suggests that average treatment effects may mask larger responses in places where flood insurance is already economically meaningful.

Table 1: Sample Construction: From US Zip Codes to Estimation Panel

Step	Description	Zip Codes
<i>Panel A: Geographic sample</i>		
1	All US zip codes	33,782
2	Restrict to NOAA coastal counties (ocean-bordering)	5,450
3	Merge with Zillow ZHVI (drop zips without coverage)	4,836
<i>Panel B: Treatment classification (within 4,836)</i>		
	LOMR–ZCTA spatial overlay pairs	2,936
	Zip codes with ≥ 1 effective LOMR	1,125
	First LOMR in window (all treated zips)	855
	Single LOMR, first in window (estimation treated group)	398
	Multi-LOMR zips	550
	Already treated before 2009	54
	First LOMR after sample window	216
	Never-treated zip codes	3,711
<i>Panel C: Event-study restrictions</i>		
4	Drop already-treated (no clean pre-period)	–54
	Drop multi-LOMR (overlapping treatment timing)	–510
	Final estimation sample	4,272
	Single-LOMR treated (in window)	398
	Not-yet-treated (first LOMR after window)	163
	Never-treated controls	3,711
5	Quarterly panel (after collapse)	228,005 obs

Notes: Sample begins with all US zip codes from the SimpleMaps database. Step 2 restricts to zip codes within NOAA-designated coastal counties that border the ocean (excluding Alaska, Hawaii, and US territories). Step 3 drops zips lacking Zillow Home Value Index coverage. Treatment is defined by spatial overlay of FEMA LOMR polygons onto Census ZCTA boundaries. Categories in Panel B are not mutually exclusive: 40 zips are both multi-LOMR and already-treated. Panel C drops are applied sequentially; the multi-LOMR drop count (510) is net of those 40 overlapping zips already removed in the prior step (550 total multi-LOMR minus 40 overlap = 510 additional drops). Not-yet-treated zips (163) retain single LOMRs arriving after 2022 and contribute control observations during the sample period. The quarterly panel collapses monthly observations by averaging log values within each quarter.

Table 2: Summary Statistics: Estimation Sample (2009–2022)

	count	mean	sd	min	p25	p50	p75	max
<i>Panel A: Outcomes</i>								
Home Value Index (Dec 2022 \$)	228,005	486,214.62	416,676.37	32,424.71	234,258.39	374,152.97	594,862.94	8,125,586.00
ln(Real ZHVI)	228,005	12.85	0.67	11.36	12.36	12.83	13.30	14.56
ln(NFIP Policies + 1)	228,005	2.46	1.80	0.00	0.98	2.08	3.71	7.87
ln(NFIP Claims + 1)	228,005	0.11	0.46	0.00	0.00	0.00	0.00	8.16
<i>Panel B: Treatment Variables</i>								
Post-LOMR	228,005	0.04	0.19	0.00	0.00	0.00	0.00	1.00
Ever Treated	228,005	0.13	0.34	0.00	0.00	0.00	0.00	1.00
Policy Intensity (per 1,000)	228,005	5.93	21.07	0.00	0.18	0.68	4.13	716.99
Treatment Intensity (per 1,000)	228,005	3.60	32.54	0.00	0.00	0.00	0.00	944.01
Upzoned (into SFHA)	228,005	0.05	0.21	0.00	0.00	0.00	0.00	1.00
Downzoned (out of SFHA)	228,005	0.03	0.18	0.00	0.00	0.00	0.00	1.00
SFHA-Crossing LOMR	228,005	0.08	0.27	0.00	0.00	0.00	0.00	1.00
<i>Panel C: Controls</i>								
County Unemp. Rate (%)	228,005	6.47	2.89	1.70	4.23	5.80	8.27	29.80
NFIP Policies (qtr avg)	228,005	65.21	164.11	0.00	1.67	7.00	39.67	2,624.67
NFIP Avg Premium (\$)	228,005	710.36	631.82	0.00	357.67	565.67	914.10	19,570.33
NFIP Claims (qtr avg)	228,005	1.13	25.22	0.00	0.00	0.00	0.00	3,498.00
<i>Panel D: Heterogeneity Measures</i>								
Republican County	222,846	0.50	0.50	0.00	0.00	1.00	1.00	1.00
R Two-Party Vote Share	222,846	0.42	0.15	0.06	0.32	0.43	0.52	0.89
Disclosure: Strict (9 states)	228,005	0.29	0.45	0.00	0.00	0.00	1.00	1.00
Disclosure: Broad (13 states)	228,005	0.61	0.49	0.00	0.00	1.00	1.00	1.00
<i>Panel E: Sample Characteristics</i>								
Zip Population	228,005	19,717.91	17,846.46	25.00	5,236.00	15,229.00	29,309.00	109,931.00
Zip Pop. Density	228,005	4,960.31	11,251.76	0.90	225.60	1,571.80	4,837.00	143,070.70
SFHA Zone Share	228,005	0.35	0.31	0.00	0.00	0.33	0.60	1.00

4.5 Balance Table

Table 3 compares pre-treatment means across treated and control (never-treated and not-yet-treated) ZIP codes. Pre-treatment home values are similar, while treated ZIP codes show higher NFIP exposure, reinforcing the case for within-ZIP identification.

Table 3: Balance Table: Pre-Treatment Characteristics

	Control	Treated	Difference
Home Value (Dec 2022 \$)	478,848.23	474,199.50	-4,648.73
Population	18,542.78	21,918.10	3,375.32***
Pop. Density (per sq mi)	5,027.05	3,123.25	-1,903.80***
County Unemp. Rate (%)	6.45	6.80	0.35***
NFIP Policies (qtr avg)	59.98	84.55	24.57***
NFIP Avg Premium (\$)	690.42	781.76	91.35***
SFHA Zone Share	0.34	0.40	0.06***
NFIP Claims (qtr avg)	1.07	1.35	0.28

Treated: zips with single LOMR during 2009-2022.

Control: zips with no LOMR.

Pre-treatment means reported for treated zips.

Difference = Treated – Control. Welch t-test.

5 Empirical Strategy

5.1 Baseline Event-Study Specification

I estimate dynamic effects with a staggered event-study difference-in-differences design because the institutional argument in Section 2.2 motivates a dynamic specification that can trace the adjustment path after a LOMR. The unit of observation is the ZIP-quarter. I assign treatment timing with the first effective LOMR in the 2009–2022 window, and I exclude ZIP codes with earlier treatment or multiple LOMRs (as described in Section 4.3).

$$\ln(\text{Real ZHVI})_{z,t} = \alpha_z + \delta_{c(z),y(t)} + \sum_{\tau \neq -1} \beta_\tau \cdot \mathbf{1}[t - E_z = \tau] + \gamma' \mathbf{X}_{z,t} + \varepsilon_{z,t} \quad (1)$$

In Equation 1, α_z denotes ZIP fixed effects, $\delta_{c(z),y(t)}$ denotes county-by-year fixed effects, and E_z is the effective date of the first LOMR for ZIP z . The coefficients of interest, β_τ , compare treated ZIP codes at each event-time bin to the omitted year immediately before treatment, $\tau = -1$. I group event time into nine annual bins from $\tau = -4$ through $\tau = +4$: an endpoint bin for three or more years before treatment, two interior pre-treatment bins, the omitted reference year immediately before treatment, four post-treatment year bins, and an endpoint bin for four or more years after treatment. The main control vector $\mathbf{X}_{z,t}$ includes county unemployment and contemporaneous NFIP policy counts.¹ I winsorize $\ln(\text{real ZHVI})$ at the 1st and 99th percentiles to limit the influence of outliers. I weight observations by ZIP population and cluster standard errors at the county level because treatment timing and housing-market shocks are unlikely to be independent within counties.

I prefer this specification to a single post-treatment indicator because a pooled TWFE coefficient would average together near-zero early effects and more negative late effects,

¹Including NFIP policies as a control raises a potential bad-control concern because Appendix A later treats policy counts as a mechanism outcome. I retain them because the event-study coefficients are nearly identical with and without NFIP controls (Column 1 vs. Column 2 of Table 4).

hiding the timing pattern that matters for the capitalization question.

5.2 Intensity Specification

I next ask whether the same map revision has a larger price effect where flood-insurance exposure was already more salient before treatment. I expect larger responses where pre-treatment NFIP participation is higher because more households already sit near the margin where insurance costs matter. The relevant margin here is not the size of the LOMR polygon itself. It is the extent to which households in the ZIP already sat close to the insurance and disclosure margin before the map change arrived.

$$\ln(\text{Real ZHVI})_{z,t} = \alpha_z + \delta_{c(z),y(t)} + \sum_{\tau \neq -1} \beta_\tau \cdot \mathbf{1}[t - E_z = \tau] \times \underbrace{\frac{\text{Policies}_z^{\text{pre}}}{\text{Pop}_z}}_{\text{Intensity}_z} + \gamma' \mathbf{X}_{z,t} + \varepsilon_{z,t} \quad (2)$$

Here Intensity_z is the ZIP’s pre-LOMR NFIP policy count per resident. I construct it from pre-treatment policy levels so the interaction captures pre-existing exposure rather than post-treatment behavioral adjustment. That distinction matters because I want the interaction to proxy for salience, not to fold an outcome of the treatment back into the treatment itself. The coefficients in [Equation 2](#) therefore show how the dynamic effect varies with baseline insurance penetration. I interpret those coefficients at realistic values of intensity, such as the sample mean or the quartiles, because an intensity value of one lies far from the center of the data.

5.3 Identification

Identification comes from comparing within-ZIP changes in home values around the first LOMR to changes in never-treated and not-yet-treated ZIP codes observed in the same

county-year. ZIP fixed effects absorb time-invariant differences across local markets, such as baseline housing quality, persistent amenity value, and long-run flood exposure. County-by-year fixed effects absorb annual shocks that operate at the county level, including broad labor-market conditions and county housing cycles.

The key assumption is that, absent the LOMR, treated ZIP codes would have followed the same trend as comparison ZIP codes within the same county-year cell. I cannot test that assumption directly, and I do not want to state it too strongly. A LOMR can follow local development, drainage work, or engineering review, and those processes may coincide with housing-market changes for reasons other than flood-risk information. That concern does not invalidate the design, but it remains the main threat to identification. I address it with the sample restrictions above, rich fixed effects, and explicit pre-trend tests. Recent work by [Goodman-Bacon \(2021\)](#), [Sun and Abraham \(2021\)](#), and [Callaway and Sant’Anna \(2021\)](#) shows that staggered TWFE event studies can be misleading when treatment timing and effects vary across cohorts. I address those concerns with dynamic estimates, cautious interpretation of pooled coefficients, and a Callaway and Sant’Anna heterogeneity-robust check in [Section 7.2](#).


5.4 Parallel Trends

The most direct diagnostic is the behavior of the lead coefficients ([Figure 2](#)). If treated ZIP codes were already on a different trend before the LOMR, the post-treatment profile would be harder to interpret as capitalization around the revision itself. In the baseline specification, the three pre-treatment coefficients are jointly insignificant, with $F(3, 338) = 1.04$ and $p = 0.373$. Therefore, the pre-treatment path is consistent with parallel trends at the annual-bin level.

However, I treat that result as necessary evidence, not a guarantee. The test has limited

power, and it cannot rule out slow-moving local shocks that line up with the engineering or administrative process behind a map revision. I read the pre-trend evidence as supportive because it reduces one obvious concern, but it does not eliminate the broader identification problem on its own.

Figure 2: Parallel Trends: Pre-Treatment Event-Time Coefficients



Note: Mean real ZHVI (December 2022 dollars) by year for treated ZIP codes (first LOMR in 2009–2022 window) and control ZIP codes (never-treated and not-yet-treated). Unconditional means, population-weighted.

6 Main Results

6.1 Baseline Effects of LOMRs

Table 4 reports the main test of H1 and H2 through baseline event-study regressions for $\ln(\text{real ZHVI})$. The dependent variable is log real ZIP-level home value, and the key regressors are event-time indicators for years relative to the first in-window LOMR. Column 1 is the total-effect specification with no controls beyond fixed effects, which avoids conditioning on NFIP policy counts that may themselves respond to treatment. Column 2 adds

county unemployment and NFIP policy controls, and Column 3 adds premiums and claims. I present all three to show whether the treatment path is sensitive to those additions. [Figure 3](#) plots the same coefficients as a dynamic path. The three lead coefficients are small, and the joint pre-trend test from [Section 5.4](#) does not reject. The first two post-treatment bins are also close to zero and statistically imprecise. However, the path becomes more negative over time. In the main specification, the coefficient at three to four years after treatment is -0.0139 with a standard error of 0.0071 and is significant at the 10 percent level. The coefficient at four or more years after treatment is -0.0287 with a standard error of 0.0113 , a 95 percent confidence interval of $[-0.0508, -0.0066]$, and a p -value of 0.011 .

Economically, the late-post-treatment estimate is associated with about a 2.8 percent decline in real home values relative to the year before treatment. At the sample mean real home value in the estimation sample, that magnitude corresponds to roughly \$13,800. That magnitude is more consistent with gradual repricing that accumulates as the map revision enters transactions, underwriting, and buyer expectations than with a large immediate shock.

The estimates also move very little across columns. The $\tau+4$ coefficient ranges from -0.0285 without controls to -0.0287 with the main controls and -0.0287 again with the fuller insurance-control set. That stability matters because it suggests the late negative effect is not especially sensitive to whether I control only for unemployment and NFIP policies or also add premiums and claims.

ZIP-level home values reflect a sequence of transactions, so the adjustment path should be slower than a single market-clearing price would suggest. I cannot rule out coincident local changes tied to development or engineering review, but the absence of differential pre-treatment decline and the later arrival of negative effects support capitalization of the map revision.

Table 4: ln(Real ZHVI) on LOMR Treatment

	(1)	(2)	(3)
	No Controls	Main Controls	Full Controls
$\tau = -4$ (3+ yrs pre)	-0.00840 (0.00765)	-0.00839 (0.00765)	-0.00836 (0.00765)
$\tau = -3$ (2-3 yrs pre)	-0.000235 (0.00472)	-0.000353 (0.00470)	-0.000299 (0.00470)
$\tau = -2$ (1-2 yrs pre)	-0.00131 (0.00304)	-0.00138 (0.00303)	-0.00134 (0.00303)
$\tau = 0$ (0-1 yr post)	0.0000564 (0.00230)	-0.0000485 (0.00231)	-0.0000269 (0.00231)
$\tau = +1$ (1-2 yrs post)	-0.00435 (0.00397)	-0.00454 (0.00400)	-0.00457 (0.00399)
$\tau = +2$ (2-3 yrs post)	-0.00908 (0.00595)	-0.00926 (0.00599)	-0.00923 (0.00599)
$\tau = +3$ (3-4 yrs post)	-0.0139* (0.00717)	-0.0139* (0.00714)	-0.0139* (0.00714)
$\tau = +4$ (4+ yrs post)	-0.0285** (0.0113)	-0.0287** (0.0113)	-0.0287** (0.0113)
Observations	228,002	228,002	228,002
Within R ²	0.0015	0.0029	0.0030

Standard errors in parentheses

Zip and county×year fixed effects.

Standard errors clustered at county level.

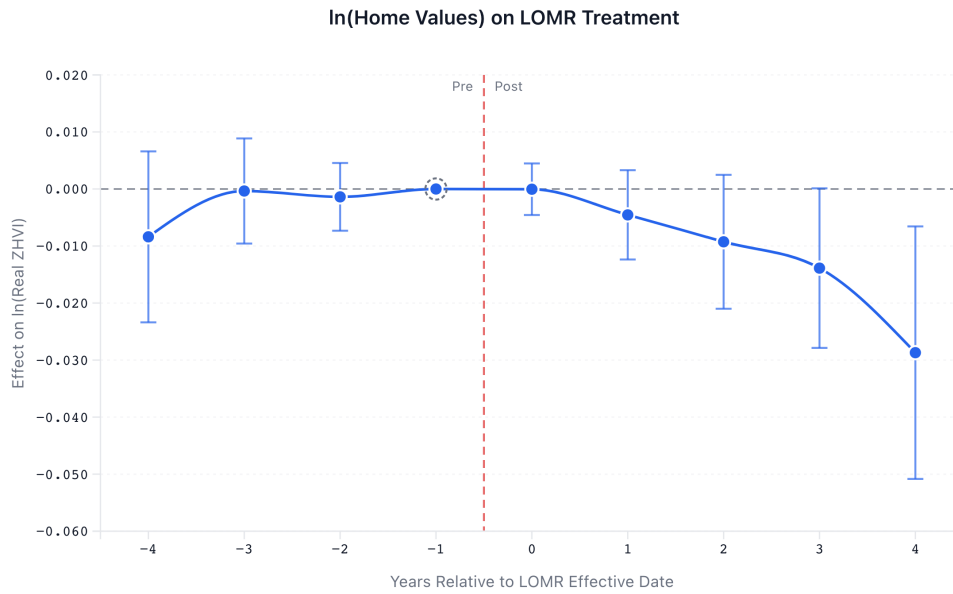
Reference period: 12-0 months before LOMR ($\tau = -1$).

Already-treated zips (LOMR before 2009) excluded.

(1) No controls. (2) Unemployment rate, NFIP policies. (3) Adds avg premium, claims.

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

Figure 3: Event Study: LOMR Effect on Home Values



Note: Coefficient estimates from the main specification (Equation 1) with 95% confidence intervals. Reference period: $\tau = -1$. Controls: county unemployment, NFIP policies. ZIP and county-by-year fixed effects. Standard errors clustered at county level. Population-weighted.

6.2 Heterogeneity by Pre-Treatment Insurance Exposure

Table 5 provides the primary test of H3 and keeps $\ln(\text{real ZHVI})$ as the dependent variable while comparing two sets of regressors. Column 1 repeats the binary event-time indicators from the baseline design, and Column 2 replaces them with event-time indicators interacted with pre-treatment NFIP policy intensity. Figure 4 plots the interaction coefficients as a dynamic path. The comparison tests whether the average treatment effect masks larger declines in places where flood insurance was already more prominent before treatment. In the continuous interaction specification, the later post-treatment coefficients are more negative than in the baseline binary model. The estimate at four or more years after treatment is $\hat{\beta}_{\tau+4} = -7.15^{***}$ with a standard error of 2.54. The coefficient is large in absolute value because I measure intensity as policies per resident (mean = 0.0059, or equivalently 5.93 per 1,000 residents as reported in Table 2). Evaluated at the sample mean, that coefficient implies about a 4.2 percent decline in real home values, roughly 1.5 times the baseline binary estimate. The joint pre-trend test yields $F(3, 338) = 1.47$ with $p = 0.223$, which is broadly supportive, although the interaction magnitudes deserve more caution than the baseline event study because the heterogeneous-effects design is less stable overall.

I interpret the intensity results as evidence that pre-treatment insurance penetration amplifies the price effect of a LOMR because it proxies for flood-risk salience, not protection from repricing. If flood insurance mainly insulates owners from revised risk, higher NFIP penetration should weaken the housing-price response. I find the opposite pattern, consistent with Wagner (2022), who shows that flood-insurance participation reflects pre-existing observable differences related to flood-risk exposure rather than simply muting price responses. The data favor the salience interpretation over a moral-hazard or already-informed alternative, though the design does not cleanly distinguish between them.

Table 5: $\ln(\text{Real ZHVI})$ on $\text{LOMR} \times \text{Policy Intensity}$

	(1)	(2)
	Binary Treatment	Intensity-Weighted
$\tau = -4$ (3+ yrs pre)	-0.00839 (0.00765)	
$\tau = -3$ (2-3 yrs pre)	-0.000353 (0.00470)	
$\tau = -2$ (1-2 yrs pre)	-0.00138 (0.00303)	
$\tau = 0$ (0-1 yr post)	-0.0000485 (0.00231)	
$\tau = +1$ (1-2 yrs post)	-0.00454 (0.00400)	
$\tau = +2$ (2-3 yrs post)	-0.00926 (0.00599)	
$\tau = +3$ (3-4 yrs post)	-0.0139* (0.00714)	
$\tau = +4$ (4+ yrs post)	-0.0287** (0.0113)	
$\tau = -4 \times \text{intensity}$		0.411 (0.574)
$\tau = -3 \times \text{intensity}$		-0.130 (0.279)
$\tau = -2 \times \text{intensity}$		-0.302 (0.194)
$\tau = 0 \times \text{intensity}$		-0.00919 (0.287)
$\tau = +1 \times \text{intensity}$		-0.811** (0.398)
$\tau = +2 \times \text{intensity}$		-1.651* (0.881)
$\tau = +3 \times \text{intensity}$		-3.288*** (1.030)
$\tau = +4 \times \text{intensity}$		-7.152*** (2.541)
Observations	228,002	228,002
Within R^2	0.0029	0.0039

Standard errors in parentheses


Zip and county \times year FE. SE clustered at county level.

Intensity = pre-LOMR NFIP policies / population.

Proxy for pre-treatment NFIP penetration.

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

Figure 4: Intensity-Weighted Event Study



Note: Coefficient estimates from the intensity specification (Equation 2) with 95% confidence intervals. Treatment intensity = pre-LOMR NFIP policies per capita. ZIP and county-by-year fixed effects. Standard errors clustered at county level. Population-weighted.

A quartile decomposition of the same interaction shows that repricing becomes detectable once a ZIP reaches moderate insurance exposure (Q2), although the gradient across upper quartiles is non-monotonic and imprecise (Table A1, Figure A2). The continuous intensity specification therefore serves as the main evidence, and the quartile design as a descriptive check.

7 Exploratory Heterogeneity and Robustness

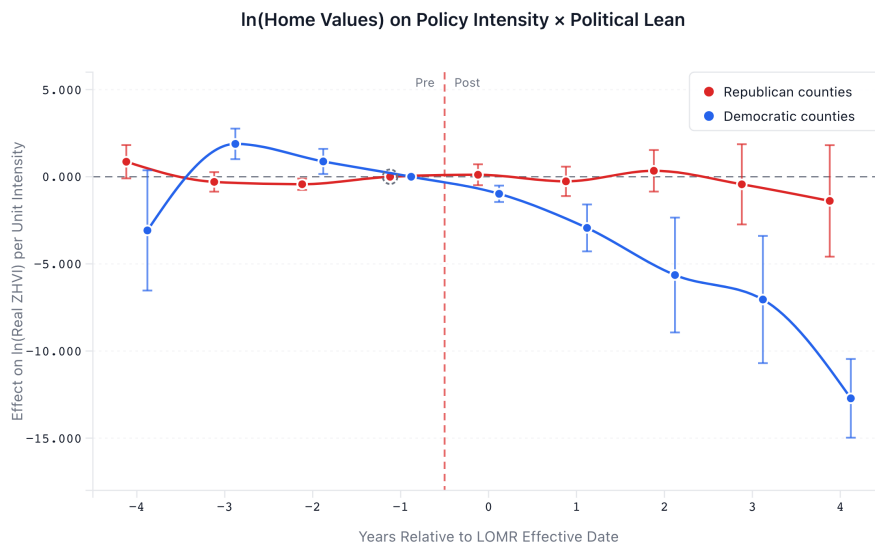
7.1 Exploratory Heterogeneity

This section turns to H4 and supporting heterogeneity checks that I treat as exploratory. The overall panel remains large, but several splits leave only about 100 to 225 treated ZIP codes, which reduces precision, and some diagnostics weaken. Splitting by proxy SFHA-share

direction and by state disclosure regime does not produce individually precise estimates, and neither set of interaction terms is statistically distinguishable from zero (Table A3, Figure A4, Table A4, Figure A5). The null on disclosure is compatible with the LOMR itself serving as the relevant information event, although low power in the strict-disclosure subsample limits that interpretation.

Table 6 returns to $\ln(\text{real ZHVI})$ as the dependent variable, but the key regressors now split the intensity-weighted event-study coefficients into a baseline path for Democratic-leaning counties and a differential interaction for Republican-leaning counties. Figure 5 plots the implied event-study paths.

Figure 5: Event Study by Political Lean



Note: Intensity-weighted event-study coefficients split by county-level Republican vote share (above/below median). Democratic-leaning counties (blue) and Republican-leaning counties (red). Pre-treatment interaction F-test rejects parallel trends ($p < 0.001$); interpret as descriptive.

These results present a sharper initial contrast on political heterogeneity. Democratic-leaning counties show a large negative baseline intensity effect, while the Republican interaction term is positive and similar in magnitude, which implies a near-zero net effect in

Table 6: $\ln(\text{Real ZHVI})$ on Policy Intensity \times Political Lean

	(1)
	Intensity \times Republican Interaction
$\tau = -4 \times \text{intensity}$	-3.080* (1.760)
$\tau = -3 \times \text{intensity}$	1.885*** (0.447)
$\tau = -2 \times \text{intensity}$	0.877** (0.368)
$\tau = 0 \times \text{intensity}$	-0.979*** (0.240)
$\tau = +1 \times \text{intensity}$	-2.937*** (0.687)
$\tau = +2 \times \text{intensity}$	-5.640*** (1.681)
$\tau = +3 \times \text{intensity}$	-7.046*** (1.862)
$\tau = +4 \times \text{intensity}$	-12.72*** (1.152)
$\tau = -4 \times \text{intensity} \times \text{Republican}$	3.942** (1.826)
$\tau = -3 \times \text{intensity} \times \text{Republican}$	-2.181*** (0.531)
$\tau = -2 \times \text{intensity} \times \text{Republican}$	-1.308*** (0.403)
$\tau = 0 \times \text{intensity} \times \text{Republican}$	1.092*** (0.389)
$\tau = +1 \times \text{intensity} \times \text{Republican}$	2.672*** (0.811)
$\tau = +2 \times \text{intensity} \times \text{Republican}$	5.981*** (1.789)
$\tau = +3 \times \text{intensity} \times \text{Republican}$	6.610*** (2.201)
$\tau = +4 \times \text{intensity} \times \text{Republican}$	11.33*** (1.997)
Observations	222,843
Within R ²	0.0058

Standard errors in parentheses

Sample: all treated zips + never-treated controls (full event study sample).

ibin = intensity LOMR effect (Dem-leaning counties). rep = differential Republican effect.

Republican = above-median county Republican two-party vote share.

Intensity = pre-LOMR NFIP policies / population.

Zip and county \times year FE. SE clustered at county level.

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

more Republican counties. That pattern is in line with belief heterogeneity dampening the capitalization of official flood-risk information. However, the pre-treatment interaction test rejects parallel trends, with $F(3, 331) = 12.05$ and $p < 0.001$, so the partisan gap may reflect pre-existing differential trends rather than a differential response to the LOMR itself. Therefore, I present the partisan split as descriptive evidence that motivates further investigation rather than as a causal claim.

7.2 Robustness Checks

The Goodman-Bacon decomposition (Figure A8) is reassuring on one dimension: 95.4 percent of the identifying weight comes from never-treated versus treated comparisons. However, the simple annual TWFE overall DD estimate underlying that decomposition is positive (+0.020), consistent with the positive sign in the unweighted and C&S specifications and opposite to the main event-study result. The Callaway and Sant’Anna check (Figure A9) shows a positive pre-treatment average ($Pre_{avg} = 0.0059$, $p = 0.002$) and an imprecise post-treatment average ($Post_{avg} = 0.0073$, $p = 0.299$), in tension with the baseline; I note that this estimator uses annual data, a balanced panel, no controls, and no weights, features that align with the attenuated unweighted column in Table 7. The unemployment placebo (Figure A10) fails the pre-trend test ($F(3, 338) = 7.22$, $p = 0.0001$), which implies that LOMR-generating processes may correlate with local economic changes not fully absorbed by a linear unemployment control. Those checks suggest the baseline result depends in part on population weighting and may reflect some correlation between LOMR timing and local economic dynamics.

Table 7 keeps $\ln(\text{real ZHVI})$ as the dependent variable and changes the treatment specification across columns. Column 1 reproduces the weighted baseline event study, Column 2 removes the population weights, and Column 3 replaces the insurance-intensity interaction

with a geographic-intensity interaction based on LOMR overlap area. I use the table to show which parts of the baseline result survive those alternative choices.

Excluding Florida reduces the $\tau+4$ estimate from -0.029 to -0.017 and removes significance, while state-level clustering preserves the baseline result, with a 95 percent confidence interval of $[-0.0545, -0.0029]$ (Appendix B). In Table 7, unweighted estimates are close to zero, and the geographic-intensity interaction is null, even though the policy-intensity interaction in Section 6.2 is negative and significant. That contrast points toward salience rather than a simple area-overlap gradient.² Overall, the baseline pattern is present but not universal across every alternative specification.

²The population-threshold subsample ($\geq 1,000$ residents) leaves the baseline essentially unchanged ($\tau+4 = -0.025$, SE = 0.0115). The SFHA-crossing subsample (Table A6) remains negative but insignificant ($\tau+4 = -0.019$, SE = 0.012).

Table 7: ln(Real ZHVI) on LOMR (Robustness)

	(1)	(2)	(3)
	Main (weighted)	Unweighted	Geographic Intensity
$\tau = -4$ (3+ yrs pre)	-0.00839 (0.00765)	-0.00513 (0.00583)	
$\tau = -3$ (2-3 yrs pre)	-0.000353 (0.00470)	-0.00113 (0.00358)	
$\tau = -2$ (1-2 yrs pre)	-0.00138 (0.00303)	-0.00165 (0.00231)	
$\tau = 0$ (0-1 yr post)	-0.0000485 (0.00231)	0.000768 (0.00204)	
$\tau = +1$ (1-2 yrs post)	-0.00454 (0.00400)	-0.000575 (0.00357)	
$\tau = +2$ (2-3 yrs post)	-0.00926 (0.00599)	-0.00257 (0.00515)	
$\tau = +3$ (3-4 yrs post)	-0.0139* (0.00714)	-0.00386 (0.00639)	
$\tau = +4$ (4+ yrs post)	-0.0287** (0.0113)	-0.00785 (0.0120)	
geo_m4			-0.0751 (0.0461)
geo_m3			-0.0201 (0.0341)
geo_m2			-0.0329 (0.0256)
geo_p0			0.0144 (0.0209)
geo_p1			-0.000154 (0.0318)
geo_p2			0.0135 (0.0395)
geo_p3			0.0267 (0.0520)
geo_p4			0.0536 (0.0654)
Observations	228,002	228,002	228,002
Within R ²	0.0029	0.0014	0.0017

Standard errors in parentheses

Zip and county \times year fixed effects.

Standard errors clustered at county level.

Geographic intensity = LOMR polygon area / ZCTA area.

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

8 Discussion

8.1 Interpretation

Taken together, the results suggest that LOMRs operate as slow-moving regulatory information shocks. The continuous intensity specification points to larger declines where pre-treatment NFIP penetration is higher, consistent with the idea that a map revision has more pricing content where more households, lenders, and insurers already sit near the margin where flood designation affects expected costs.

The weaker results limit how far that interpretation extends. Insurance take-up shows limited movement around the LOMR, and claims are largely unaffected, pointing to an informational shock rather than a realized disaster (Table A2). The policy-intensity interaction is significant but the geographic-overlap measure is null, which favors economic salience over simple area overlap. Disclosure-law interactions are weak (Table A4), the up/down price splits are imprecise (Table A3), and the political split fails the pre-trend standard (Table 6). Therefore, I read the core result as capitalization of revised official risk information, with supporting evidence on salience but only mixed evidence on the exact transaction channel.

8.2 Limitations

Several limitations remain. The analysis uses ZIP-level panels rather than parcel-level or transaction-level data, so I cannot observe which homes actually cross the SFHA boundary or how quickly lenders and buyers respond inside each ZIP. The baseline specification also pools revisions that may raise or lower mapped exposure because my ZIP-level treatment construction does not directly identify direction from the LOMR overlay alone. I proxy for directional heterogeneity with changes in NFIP SFHA share around the LOMR, but that measure is an exposure proxy, not a direct parcel-level map transition. The heterogeneity

exercises therefore remain suggestive. The robustness checks in [Section 7.2](#) also narrow the causal claim, particularly the sensitivity to population weighting and to Florida.

9 Conclusion

This paper asks whether official FEMA flood-map revisions affect local housing prices. Using a ZIP-by-quarter panel for coastal U.S. ZIP codes from 2009 to 2022 and a staggered event-study design, I find that first in-window LOMRs are associated with gradual declines in home values rather than an immediate one-period drop. In the baseline specification, the effect reaches about 2.8 percent four or more years after treatment. In the intensity specification, the implied effect at the sample mean insurance exposure is about 4.2 percent. Those results are consistent with the idea that map revisions matter most where flood risk is already economically salient.

Mechanism evidence on insurance take-up is limited, and the exploratory heterogeneity results are mixed, so I do not claim a single channel with confidence. Still, the main pattern is consistent across the preferred specification. Official flood-risk reclassification can be capitalized into housing markets over time even without a contemporaneous flood. Future work with parcel-level transactions, lender data, and direct map-transition measures could separate information, insurance, and financing channels more cleanly.

Data Availability

Replication files, code, and the constructed analysis dataset are available at <https://fear.camkeith.me>.

References

- Bakkensen, Laura A. and Lint Barrage (2022) “Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics,” *The Review of Financial Studies*, 35 (8), 3666–3709.
- Baldauf, Markus, Lorenzo Garlappi, and Constantine Yannelis (2020) “Does Climate Change Affect Real Estate Prices? Only If You Believe In It,” *The Review of Financial Studies*, 33 (3), 1256–1295.
- Beltrán, Allan, David Maddison, and Robert Elliott (2018) “Is Flood Risk Capitalised into Property Values?” *Ecological Economics*, 146, 668–685.
- Bernstein, Asaf, Matthew T. Gustafson, and Ryan Lewis (2019) “Disaster on the Horizon: The Price Effect of Sea Level Rise,” *Journal of Financial Economics*, 134 (2), 253–272.
- Bin, Okmyung and Stephen Polasky (2004) “Effects of Flood Hazards on Property Values: Evidence Before and After Hurricane Floyd,” *Land Economics*, 80 (4), 490–500.
- Callaway, Brantly and Pedro H. C. Sant’Anna (2021) “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 225 (2), 200–230.
- Congressional Research Service (2025a) “Introduction to the National Flood Insurance Program (NFIP),” Report R44593, Library of Congress, <https://www.congress.gov/crs-product/R44593>.
- (2025b) “National Flood Insurance Program Borrowing Authority,” Insight IN10784, Library of Congress, <https://www.congress.gov/crs-product/IN10784>.
- Federal Emergency Management Agency (2024a) “Letters of Map Revision (LOMRs),” FEMA.
- (2024b) “National Flood Insurance Program,” FEMA.
- Federal Reserve Bank of St. Louis (2024) “Consumer Price Index for All Urban Consumers (CPI-U),” FRED Economic Data.
- Gibson, Matthew and Jamie T. Mullins (2020) “Climate Risk and Beliefs in New York Floodplains,” *Journal of the Association of Environmental and Resource Economists*, 7 (6), 1069–1111.
- Goodman-Bacon, Andrew (2021) “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 225 (2), 254–277.
- Hallstrom, Daniel G. and V. Kerry Smith (2005) “Market Responses to Hurricanes,” *Journal of Environmental Economics and Management*, 50 (3), 541–561.
- Hino, Miyuki and Marshall Burke (2021) “The Effect of Information about Climate Risk on Property Values,” *Proceedings of the National Academy of Sciences*, 118 (17), e2003374118.

- Kousky, Carolyn (2010) “Learning from Extreme Events: Risk Perceptions after the Flood,” *Land Economics*, 86 (3), 395–422.
- MIT Election Data and Science Lab (2017) “County Presidential Election Returns,” Harvard Dataverse.
- Murfin, Justin and Matthew Spiegel (2020) “Is the Risk of Sea Level Rise Capitalized in Residential Real Estate?” *The Review of Financial Studies*, 33 (3), 1217–1255.
- Pope, Jaren C. (2008) “Do Seller Disclosures Affect Property Values? Buyer Information and the Hedonic Model,” *Land Economics*, 84 (4), 551–572.
- Shr, Yau-Huo J. and Katherine Y. Zipp (2019) “The Aftermath of Flood Zone Remapping: The Asymmetric Impact of Flood Maps on Housing Prices,” *Land Economics*, 95 (2), 174–192, [10.3368/le.95.2.174](https://doi.org/10.3368/le.95.2.174).
- Sun, Liyang and Sarah Abraham (2021) “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 225 (2), 175–199.
- Troy, Austin and Jeff Romm (2004) “Assessing the Price Effects of Flood Hazard Disclosure Under the California Natural Hazard Disclosure Law (AB 1195),” *Journal of Environmental Planning and Management*, 47 (1), 137–162.
- U.S. Bureau of Labor Statistics (2024) “Local Area Unemployment Statistics (LAUS),” Bureau of Labor Statistics.
- U.S. Census Bureau (2024) “ZIP Code Tabulation Areas (ZCTAs),” U.S. Census Bureau.
- Wagner, Katherine R. H. (2022) “Adaptation and Adverse Selection in Markets for Natural Disaster Insurance,” *American Economic Journal: Economic Policy*, 14 (3), 380–421.
- Weill, Joakim A. (2023) “Flood Risk Mapping and the Distributional Impacts of Climate Information,” Technical Report 2023-066, Board of Governors of the Federal Reserve System (U.S.), [10.17016/FEDS.2023.066](https://doi.org/10.17016/FEDS.2023.066), Forthcoming in *American Economic Journal: Economic Policy*.
- Zillow Research (2024) “ZHVI User Guide: Home Value Data Documentation,” Zillow Research.

A Additional Tables

This appendix collects supplementary estimates that I discuss more briefly in the main text. I include them here because they help show what becomes weaker, more fragile, or less precisely estimated once I move away from the baseline event-study design.

Table A1 uses $\ln(\text{real ZHVI})$ as the dependent variable with quartile-specific event-time interactions rather than one continuous intensity interaction. Each column corresponds to a different quartile of pre-treatment NFIP policy intensity among treated ZIP codes. The quartile specification does not show a smooth monotonic gradient. The lowest-exposure quartile is close to zero throughout. The second quartile shows the clearest negative response, with statistically significant coefficients at $\tau+2$, $\tau+3$, and $\tau+4$. The upper quartiles are generally negative, but the estimates are imprecise. That pattern suggests repricing appears once a ZIP reaches a moderate level of flood-insurance exposure, but it does not support the stronger claim that each additional increase in exposure produces a larger decline.

Table A1: $\ln(\text{Real ZHVI})$ by Policy Intensity Quartile

	(1)	(2)	(3)	(4)
	Q1 (Low)	Q2	Q3	Q4 (High)
$\tau = -4$ (3+ yrs pre)	0.00174 (0.0154)	-0.0194 (0.0159)	-0.0169 (0.0107)	-0.00284 (0.0190)
$\tau = -3$ (2-3 yrs pre)	0.00279 (0.00732)	0.000847 (0.00950)	-0.00771 (0.00939)	0.000271 (0.00605)
$\tau = -2$ (1-2 yrs pre)	-0.00243 (0.00474)	0.00484 (0.00431)	-0.00509 (0.00721)	-0.00487 (0.00318)
$\tau = 0$ (0-1 yr post)	0.000878 (0.00444)	0.00176 (0.00368)	-0.00339 (0.00550)	0.000415 (0.00448)
$\tau = +1$ (1-2 yrs post)	0.00210 (0.00505)	-0.0120 (0.00779)	-0.00726 (0.00912)	-0.00484 (0.00671)
$\tau = +2$ (2-3 yrs post)	0.00451 (0.00866)	-0.0301*** (0.0107)	-0.00960 (0.0116)	-0.0141 (0.0163)
$\tau = +3$ (3-4 yrs post)	0.00826 (0.0107)	-0.0396*** (0.0147)	-0.0243* (0.0131)	-0.0211 (0.0228)
$\tau = +4$ (4+ yrs post)	-0.00817 (0.0113)	-0.0479** (0.0186)	-0.0315 (0.0191)	-0.0574 (0.0472)
Observations	228,002			
Within R ²	0.0039			

Standard errors in parentheses

Zip and county \times year FE. SE clustered at county level.

Quartiles of pre-LOMR NFIP policy penetration among treated zips.

Q1 = lowest exposure, Q4 = highest exposure.

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

Insurance Market Responses

Table A2 shifts the dependent variable away from home values and onto insurance outcomes. Column 1 uses $\ln(\text{NFIP policies} + 1)$, Column 2 uses $\ln(\text{NFIP claims} + 1)$, and the key regressors remain the event-time indicators for years relative to the first LOMR. Figure A3 plots the policy coefficients. Policy counts serve as mechanism evidence: if LOMRs reduce housing values through immediate insurance adjustment, NFIP policy counts should move around the revision itself. In the aggregate policy event study, I see limited movement in $\ln(\text{NFIP policies} + 1)$, with at most a brief decline in the early post-treatment window. Claims serve a different, placebo-like role. A LOMR does not generate flooding on its own, and I find no post-treatment effects on $\ln(\text{NFIP claims} + 1)$ that are significant at the 5 percent level, although the $\tau+2$ coefficient is marginally significant at 10 percent. That pattern is consistent with an informational or regulatory shock rather than a realized disaster shock.

Table A2: $\ln(\text{NFIP Policies})$ on LOMR Treatment

	(1) $\ln(\text{Policies})$	(2) $\ln(\text{Claims})$
$\tau = -4$ (3+ yrs pre)	-0.0213 (0.0221)	-0.0262* (0.0152)
$\tau = -3$ (2-3 yrs pre)	-0.0234 (0.0165)	-0.0286 (0.0254)
$\tau = -2$ (1-2 yrs pre)	-0.0120 (0.00947)	-0.0177 (0.0181)
$\tau = 0$ (0-1 yr post)	-0.0160 (0.0121)	0.00154 (0.0195)
$\tau = +1$ (1-2 yrs post)	-0.0330** (0.0130)	0.00506 (0.0235)
$\tau = +2$ (2-3 yrs post)	-0.0291 (0.0187)	-0.0348* (0.0191)
$\tau = +3$ (3-4 yrs post)	-0.0111 (0.0264)	-0.0296 (0.0232)
$\tau = +4$ (4+ yrs post)	0.00315 (0.0387)	-0.0275 (0.0226)
Observations	228,002	228,002
Within R^2	0.0021	0.0002

Standard errors in parentheses

Zip and county \times year FE. SE clustered at county level.

Policies: NFIP policy count (mechanism). Claims: falsification.

Reference period: $\tau = -1$ (12-0 months before LOMR).

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

Table A3 keeps $\ln(\text{real ZHVI})$ as the dependent variable and splits the event-study coefficients into ZIP codes with proxy SFHA-share increases and decreases around the LOMR.

The key regressors are therefore the same event-time indicators as in the baseline specification, estimated separately for the two proxy zoning directions. Neither path is individually precise. At $\tau+4$, the upzoned estimate is -0.0142 with a standard error of 0.0170 , while the downzoned estimate is -0.0237 with a standard error of 0.0186 . Those estimates remain statistically indistinguishable from zero, so the sign pattern remains descriptive rather than probative.

Table A4 also uses $\ln(\text{real ZHVI})$ as the dependent variable, but the key regressors are event-time indicators interacted with state disclosure-law categories. The strict-disclosure category includes nine states with verified mandatory flood-zone disclosure on the seller form (CA, IL, IN, LA, MS, OR, SC, TX, WI), and the broad category adds four states with weaker or more recent requirements (FL, NC, NY, VA). The table asks whether the home-value response differs across these regimes. The late post-treatment interaction for the strict-disclosure specification is 0.0065 with a standard error of 0.0231 , while the corresponding broad-disclosure interaction is -0.0260 with a standard error of 0.0238 . Neither post-treatment interaction is statistically distinguishable from zero. However, the strict-disclosure lead at $\tau-4$ is positive, 0.0280 with a standard error of 0.0166 , which is another reason to avoid a strong causal interpretation of these subsample differences.

Table A5 changes the dependent variables to insurance outcomes and estimates them separately for the proxy upzoning and downzoning groups. The key regressors remain the event-time indicators, but the table shows whether the insurance response differs with the direction of the proxy map-status change. The clearest pattern is on the policy side in downzoned ZIP codes. The coefficients are -0.0648 with a standard error of 0.0311 at $\tau+1$, and -0.0924 with a standard error of 0.0429 at $\tau+2$. Upzoned policy effects remain close to zero throughout. I read this asymmetry as suggestive evidence that the removal of mandatory coverage may matter more clearly than the imposition of coverage, although the broader mechanism evidence remains limited.

Table A6 returns to $\ln(\text{real ZHVI})$ and restricts the baseline event study to ZIP codes with non-stable NFIP SFHA-share changes around the LOMR. The key regressors are still the event-time indicators, but the sample is narrower and closer to the subset where the proxy zoning-direction measure is most likely to reflect a real shift in mapped exposure. The late-post-treatment coefficient falls to -0.0192 with a standard error of 0.0124 , which remains statistically insignificant. Therefore, the baseline result does not become sharper when I focus only on the subset most likely to capture meaningful map-status changes under the proxy definition.

Table A7 keeps $\ln(\text{real ZHVI})$ as the dependent variable but replaces the full event-time path with one post-LOMR indicator. The key regressor is therefore a single treated-after-treatment dummy in a pooled two-way fixed-effects difference-in-differences design. The Post-LOMR coefficient is -0.0070 with a standard error of 0.0081 . That estimate is much smaller than the late-period event-study coefficient and statistically indistinguishable from zero, which is exactly what I would expect if a pooled post indicator averages together early near-zero periods and later negative periods.

Table A3: ln(Real ZHVI) on Upzoning vs Downzoning

	(1)	(2)
	Upzoned (Risk ↑)	Downzoned (Risk ↓)
$\tau = -4$ (3+ yrs pre)	-0.00612 (0.0115)	
$\tau = -3$ (2-3 yrs pre)	0.00552 (0.00659)	
$\tau = -2$ (1-2 yrs pre)	-0.00176 (0.00260)	
$\tau = 0$ (0-1 yr post)	0.00441 (0.00355)	
$\tau = +1$ (1-2 yrs post)	0.000825 (0.00633)	
$\tau = +2$ (2-3 yrs post)	0.000427 (0.0102)	
$\tau = +3$ (3-4 yrs post)	-0.000621 (0.0121)	
$\tau = +4$ (4+ yrs post)	-0.0142 (0.0170)	
$\tau = -4$ (3+ yrs pre)		-0.00462 (0.0144)
$\tau = -3$ (2-3 yrs pre)		-0.000508 (0.0102)
$\tau = -2$ (1-2 yrs pre)		-0.00215 (0.00581)
$\tau = 0$ (0-1 yr post)		0.000759 (0.00621)
$\tau = +1$ (1-2 yrs post)		-0.00167 (0.00838)
$\tau = +2$ (2-3 yrs post)		-0.00579 (0.0122)
$\tau = +3$ (3-4 yrs post)		-0.00851 (0.0147)
$\tau = +4$ (4+ yrs post)		-0.0237 (0.0186)
Observations	208,170	205,208
Within R ²	0.0016	0.0018

Standard errors in parentheses

Zip and county×year FE. SE clustered at county level.

Each spec uses its treated subsample + all never-treated controls.

Risk direction classified by SFHA zone share change (± 1 pp) around LOMR.

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

Table A4: $\ln(\text{Real ZHVI})$ on $\text{LOMR} \times \text{Disclosure}$

	(1) Binary	(2) \times Strict Disclosure	(3) \times Broad Disclosure
$\tau = -4$ (3+ yrs pre)	-0.00839 (0.00765)	-0.0196** (0.00846)	-0.0136 (0.0123)
$\tau = -3$ (2-3 yrs pre)	-0.000353 (0.00470)	-0.00280 (0.00438)	-0.00556 (0.00581)
$\tau = -2$ (1-2 yrs pre)	-0.00138 (0.00303)	-0.00172 (0.00238)	-0.00319 (0.00295)
$\tau = 0$ (0-1 yr post)	-0.0000485 (0.00231)	0.000317 (0.00272)	0.00550* (0.00297)
$\tau = +1$ (1-2 yrs post)	-0.00454 (0.00400)	-0.00505 (0.00512)	0.000766 (0.00605)
$\tau = +2$ (2-3 yrs post)	-0.00926 (0.00599)	-0.00864 (0.00837)	0.00157 (0.00835)
$\tau = +3$ (3-4 yrs post)	-0.0139* (0.00714)	-0.0126 (0.0114)	0.00139 (0.0117)
$\tau = +4$ (4+ yrs post)	-0.0287** (0.0113)	-0.0324* (0.0194)	-0.00786 (0.0199)
$\tau = -4 \times \text{Disclosure}$		0.0280* (0.0166)	
$\tau = -3 \times \text{Disclosure}$		0.00468 (0.0107)	
$\tau = -2 \times \text{Disclosure}$		-0.000187 (0.00680)	
$\tau = 0 \times \text{Disclosure}$		-0.000305 (0.00453)	
$\tau = +1 \times \text{Disclosure}$		0.000922 (0.00807)	
$\tau = +2 \times \text{Disclosure}$		-0.00138 (0.0119)	
$\tau = +3 \times \text{Disclosure}$		-0.00269 (0.0139)	
$\tau = +4 \times \text{Disclosure}$		0.00645 (0.0231)	
$\tau = -4 \times \text{Broad Disc.}$			0.00840 (0.0155)
$\tau = -3 \times \text{Broad Disc.}$			0.00768 (0.00855)
$\tau = -2 \times \text{Broad Disc.}$			0.00276 (0.00508)
$\tau = 0 \times \text{Broad Disc.}$			-0.00770* (0.00421)
$\tau = +1 \times \text{Broad Disc.}$			-0.00746 (0.00792)
$\tau = +2 \times \text{Broad Disc.}$			-0.0149 (0.0113)
$\tau = +3 \times \text{Broad Disc.}$			-0.0204 (0.0145)
$\tau = +4 \times \text{Broad Disc.}$			-0.0260 (0.0238)
Observations	228,002	228,002	228,002
Within R ²	0.0029	0.0036	0.0033

Standard errors in parentheses

Sample: all treated zips + never-treated controls (full event study sample).

ebin = baseline LOMR effect (non-disclosure). dbin/dbinb = differential disclosure effect.

Zip and county \times year FE. SE clustered at county level.

Strict: CA, IL, IN, LA, MS, OR, SC, TX, WI (9 states).

Broad adds FL, VA, NC, NY (13 states).

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

Table A5: ln(NFIP Policies) on Upzoning vs Downzoning

	(1)	(2)	(3)	(4)
	Policies (Up)	Policies (Down)	Claims (Up)	Claims (Down)
up_m4	-0.0152 (0.0341)		-0.0613** (0.0298)	
up_m3	-0.000135 (0.0332)		-0.0839* (0.0450)	
up_m2	-0.0174 (0.0252)		-0.0334 (0.0389)	
up_p0	-0.0138 (0.0127)		-0.0664** (0.0313)	
up_p1	-0.00873 (0.0161)		-0.0448 (0.0372)	
up_p2	0.00248 (0.0244)		-0.0816** (0.0405)	
up_p3	-0.00130 (0.0293)		-0.0490 (0.0427)	
up_p4	-0.0102 (0.0341)		-0.0693* (0.0388)	
down_m4		-0.0354 (0.0320)		0.0227 (0.0170)
down_m3		-0.0560* (0.0306)		0.0599*** (0.0225)
down_m2		-0.0198 (0.0253)		0.0577** (0.0243)
down_p0		-0.0291 (0.0323)		0.0965** (0.0387)
down_p1		-0.0648** (0.0311)		0.0486* (0.0290)
down_p2		-0.0924** (0.0429)		0.0513* (0.0276)
down_p3		-0.0482 (0.0411)		0.0291 (0.0239)
down_p4		-0.0649 (0.0698)		0.0275 (0.0174)
Observations	208,170	205,208	208,170	205,208
Within R ²	0.0022	0.0025	0.0002	0.0002

Standard errors in parentheses

Zip and county×year FE. SE clustered at county level.

Each spec uses its treated subsample + all never-treated controls.

Risk direction: SFHA zone share Δ (± 1 pp threshold).

Reference period: $\tau = -1$ (12-0 months before LOMR).

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

Table A6: ln(Real ZHVI) on SFHA-Crossing LOMRs

	(1)	(2)
	All LOMRs	SFHA-Crossing Only
$\tau = -4$ (3+ yrs pre)	-0.00839 (0.00765)	-0.00625 (0.00907)
$\tau = -3$ (2-3 yrs pre)	-0.000353 (0.00470)	0.00335 (0.00588)
$\tau = -2$ (1-2 yrs pre)	-0.00138 (0.00303)	-0.00173 (0.00313)
$\tau = 0$ (0-1 yr post)	-0.0000485 (0.00231)	0.00249 (0.00349)
$\tau = +1$ (1-2 yrs post)	-0.00454 (0.00400)	-0.000532 (0.00498)
$\tau = +2$ (2-3 yrs post)	-0.00926 (0.00599)	-0.00295 (0.00732)
$\tau = +3$ (3-4 yrs post)	-0.0139* (0.00714)	-0.00504 (0.00926)
$\tau = +4$ (4+ yrs post)	-0.0287** (0.0113)	-0.0192 (0.0124)
Observations	228,002	215,580
Within R ²	0.0029	0.0020

Standard errors in parentheses

Zip and county×year FE. SE clustered at county level.

SFHA-crossing = zone_risk_direction is 'up' or 'down' (not 'stable').

Reference period: $\tau = -1$ (12-0 months before LOMR).

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

Table A7: ln(Real ZHVI) on LOMR (TWFE DiD)

	(1) TWFE DiD
Post-LOMR	-0.00699 (0.00805)
County Unemp. Rate (%)	-0.00216*** (0.000346)
NFIP Policies (qtr avg)	0.000000941 (0.0000101)
Constant	12.91*** (0.00219)
Observations	228,002
Within R ²	0.0016

Standard errors in parentheses

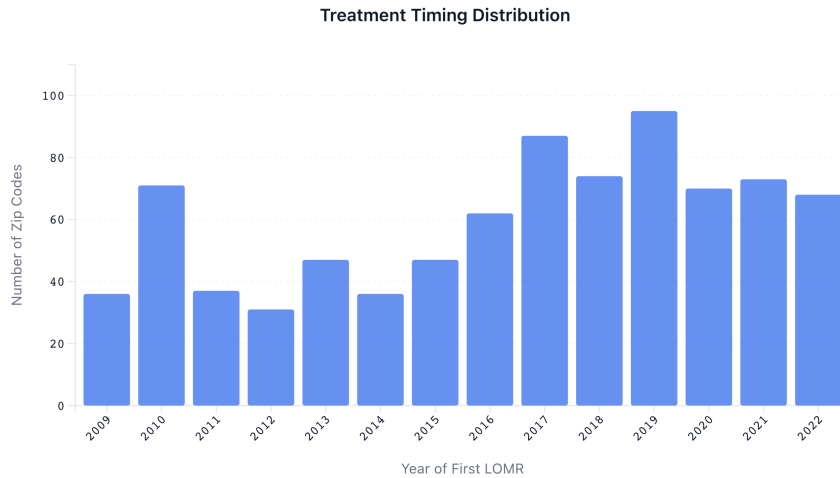
Zip and county×year FE. SE clustered at county level.

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

B Additional Figures

This section collects the event-study plots for the subsample and mechanism specifications reported in the tables above, along with the diagnostic figures referenced in Section 7.2. The main value of these figures is not to reveal a hidden sharp effect. It is to show either that treatment timing is spread across the sample period or that confidence intervals widen once the treated sample is split into narrower groups.

Figure A1: Distribution of First LOMR Year

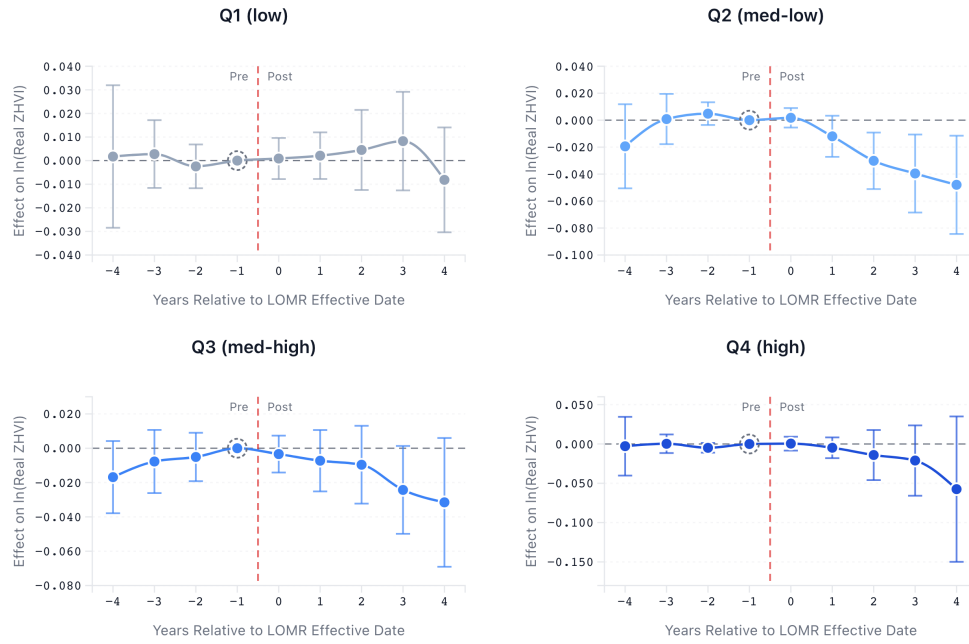


Note: Distribution of the first effective LOMR year across all treated ZIP codes with a first LOMR in the 2009–2022 window, before the multi-LOMR exclusion.

Figure A8 reports the Goodman-Bacon decomposition for the simple TWFE design. The key detail is that 95.4 percent of the identifying weight comes from never-treated versus treated timing comparisons, while only 0.3 percent comes from already-treated comparisons. Therefore, the main concern in this decomposition is not extreme negative weighting from already-treated units. The overall DD estimate from this annual, unweighted, no-controls TWFE is positive (+0.020), which is consistent with the attenuated or reversed sign seen in the unweighted and C&S specifications. The larger issue is still that a single post indicator compresses the dynamic pattern.

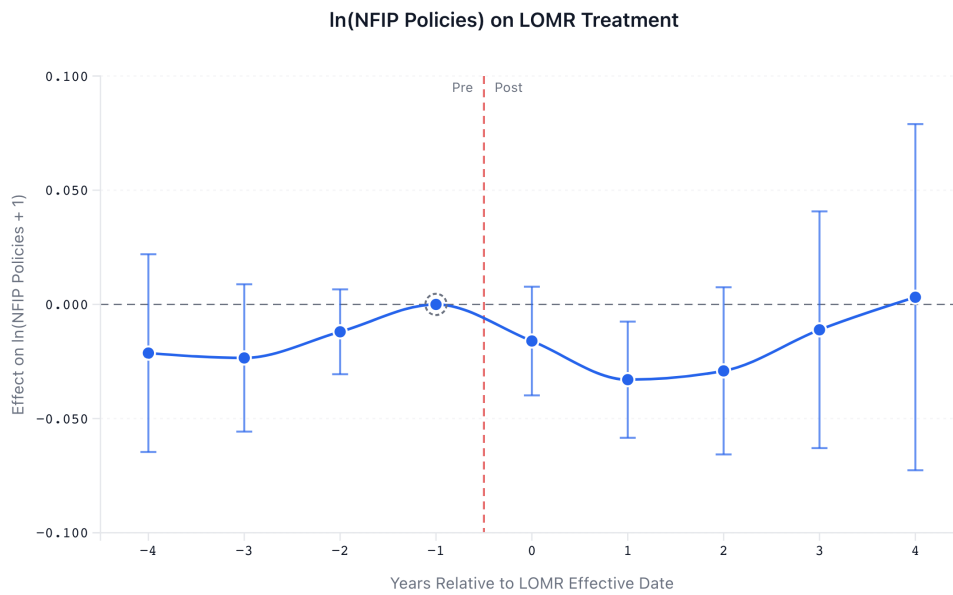
Figure A9 reports the Callaway and Sant’Anna event study. The aggregated pre-treatment average is 0.0059 with a standard error of 0.0020 and a p -value of 0.002. The post-treatment

Figure A2: Event Study by Insurance Penetration Quartile



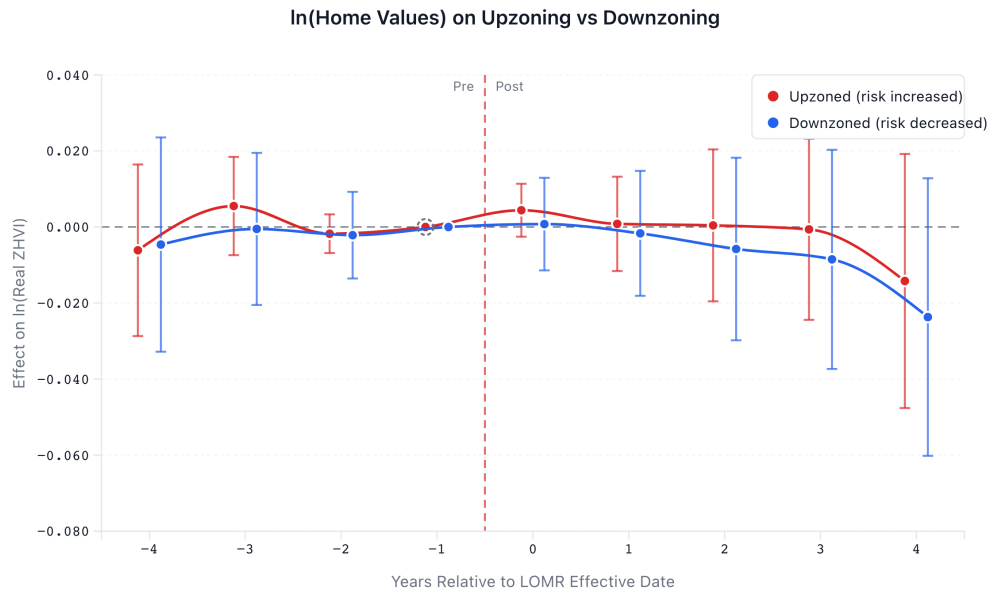
Note: Separate event-study coefficients by pre-LOMR policy intensity quartile. Q1 = lowest penetration, Q4 = highest. 95% confidence intervals shown.

Figure A3: Insurance Policies Event Study




Note: Event-study coefficients for $\ln(\text{NFIP policies} + 1)$ as the dependent variable. Same specification as Equation 1.

Figure A4: Event Study: SFHA-Share Increase vs. Decrease Proxy



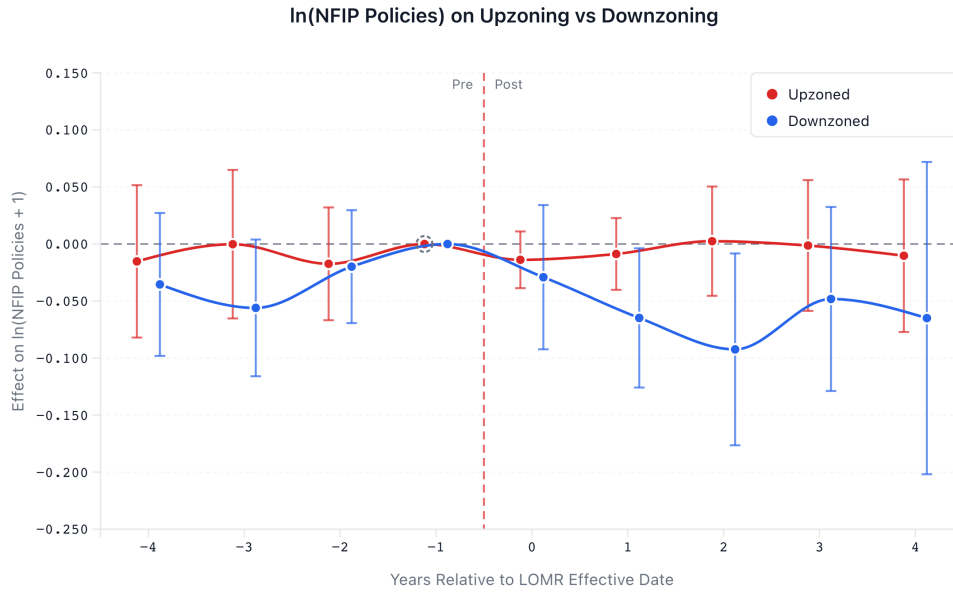
Note: Separate event-study coefficients for treated ZIP codes classified using increases or decreases in NFIP SFHA share around the LOMR.

Figure A5: Disclosure Law Heterogeneity Event Study



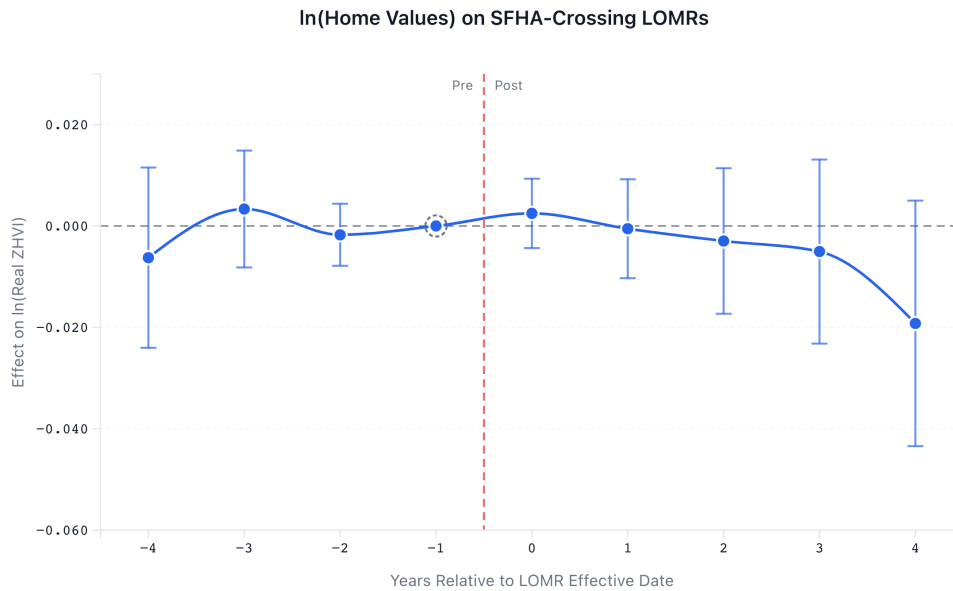
Note: Event-study coefficients interacted with state mandatory flood disclosure indicator (9 strict states).

Figure A6: Insurance Policies Event Study by SFHA-Share-Change Proxy




Note: Event-study coefficients for $\ln(\text{NFIP policies} + 1)$ estimated separately for treated ZIP codes classified by changes in NFIP SFHA share around the LOMR.

Figure A7: SFHA-Share-Change Proxy Subsample Event Study



Note: Event-study coefficients restricted to LOMRs with non-stable NFIP SFHA-share changes around the LOMR. Same specification as Equation 1.

Figure A8: Goodman-Bacon Decomposition

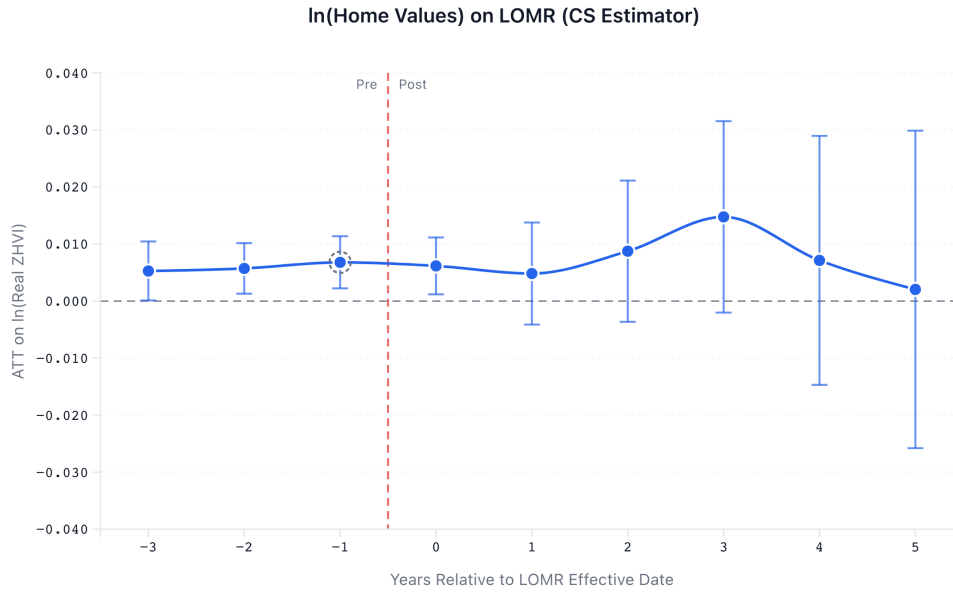


Note: Decomposition of the TWFE DiD coefficient into timing-group comparisons following [Goodman-Bacon \(2021\)](#). 95.4% of weight comes from never-treated vs. timing comparisons.

average is 0.0073 with a standard error of 0.0070 and a p -value of 0.299. That positive pre-treatment average is in tension with the main event-study design. Because this implementation uses annual data, a balanced panel, no weights, and no controls, I treat it as a cautionary check rather than the preferred estimate.

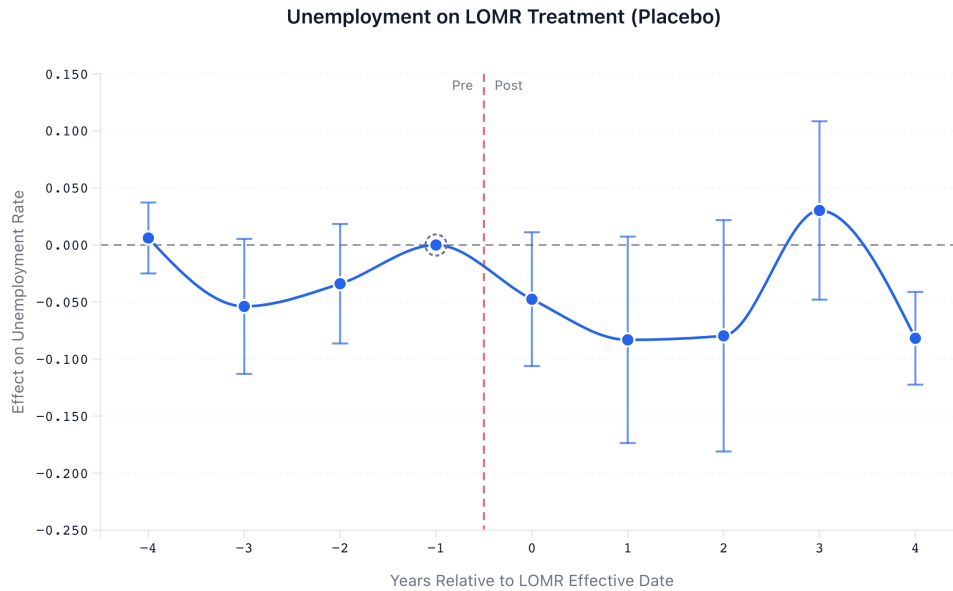
[Figure A10](#) reports the unemployment placebo. The pre-trend test fails, with $F(3, 338) = 7.22$ and $p = 0.0001$, so this figure should be read as a warning sign, not as supportive evidence. I keep it in the appendix because it is informative about the possibility of residual local dynamics around treatment timing.

Figure A9: Callaway & Sant'Anna Event Study



Note: Callaway and Sant'Anna DiD estimator following [Callaway and Sant'Anna \(2021\)](#). Annual frequency, balanced panel, not-yet-treated and never-treated as controls. Cohorts restricted to 2009–2022.

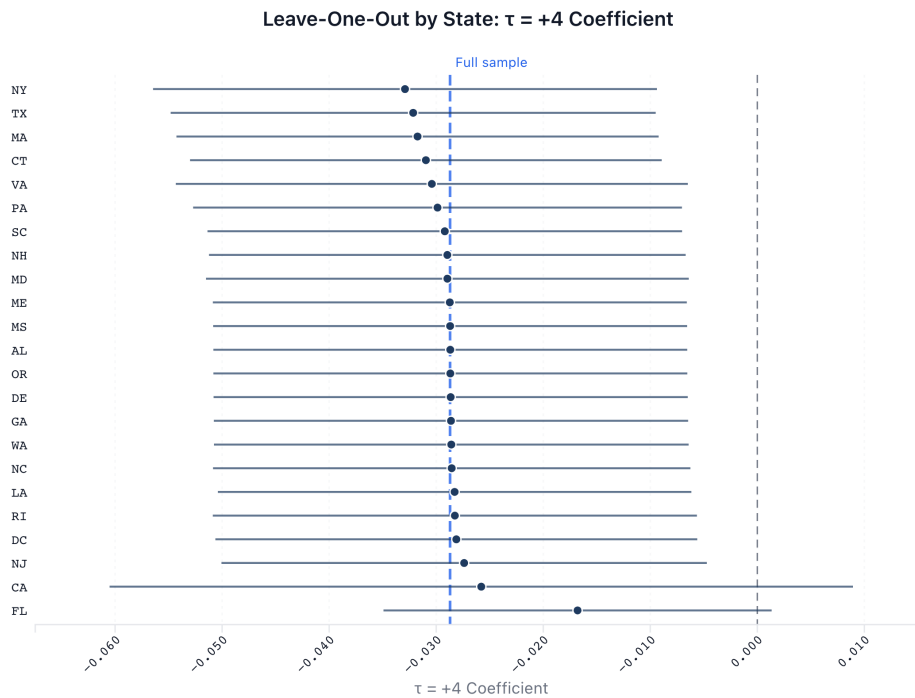
Figure A10: Placebo Event Study: County Unemployment Rate



Note: Event-study coefficients with county unemployment rate as the dependent variable. Pre-trend F-test: $F(3, 338) = 7.22$, $p = 0.0001$.

Figure A11 reports the leave-one-out exercise. Excluding Florida reduces the late-post-treatment estimate from about -0.029 to -0.017 and makes it statistically insignificant, while exclusions of other states leave the estimate close to the baseline value. Under state-level clustering, the $\tau+4$ coefficient remains -0.0287 , and the standard error rises to 0.0132 , which yields a 95 percent confidence interval of $[-0.0545, -0.0029]$. Those checks do not reverse the sign of the main result, and state-level clustering preserves statistical significance.

Figure A11: Leave-One-Out by State: $\tau = +4$ Coefficient



Note: Each point shows the $\tau = +4$ coefficient from the main specification (Equation 1) after dropping the indicated state. Horizontal bars are 95% confidence intervals. Blue dashed line marks the full-sample estimate. Florida is the only exclusion that materially shifts the estimate.

C Data Construction Details

I begin the geographic construction with 33,782 ZIP codes from the SimpleMaps ZIP database and restrict the sample to 5,450 ZIP codes located in NOAA coastal counties. Then, I overlay FEMA’s effective S_LOMR polygons onto Census ZCTA boundaries to identify exposure to official map revisions. At the ZIP level, this procedure records the number of matched LOMRs, the first and last effective dates, and a geographic treatment-intensity measure equal to the overlap area between the LOMR polygon and the ZIP proxy geography divided by total ZIP area. In the full 5,450-ZIP coastal universe (before the Zillow coverage

screen), the overlay yields 2,936 LOMR–ZIP pairs, 1,162 ever-treated ZIP codes, and 886 ZIP codes whose first observed LOMR falls inside the 2009–2022 window. After restricting to the 4,836 ZIP codes with usable Zillow data, those counts fall to 1,125 ever-treated and 855 first-in-window, as reported in [Table 1](#).

I merge those treatment measures to monthly Zillow ZIP-level home values, county unemployment rates from BLS LAUS, and OpenFEMA NFIP policies and claims aggregated to the ZIP-month level. The NFIP aggregation yields policy counts, total premiums, average premiums, SFHA zone share, claims counts, and total claims paid. I deflate Zillow values to December 2022 dollars with the Consumer Price Index for All Urban Consumers. I then export the merged monthly panel for estimation and collapse it to quarters (averaging monthly log values within each quarter) because the ZHVI is already a three-month smoothed index. At the estimation stage, I winsorize the quarterly log home values at the 1st and 99th percentiles and impose the sample restrictions: I drop ZIP codes treated before 2009 and ZIP codes with multiple LOMRs. The resulting quarterly estimation sample contains 4,272 ZIP codes.

The heterogeneity variables require additional construction choices. I define policy intensity as the ZIP’s pre-treatment average NFIP policy count divided by population, and I drop ZIP codes with values above one because those cases imply implausible denominators. I classify proxy upzoning and downzoning by comparing mean NFIP SFHA zone share in the 12 months before and after the LOMR effective date; I classify ZIP codes with a change exceeding ± 1 percentage point as upzoned or downzoned, and the remainder as stable. This measure should be read as an exposure proxy rather than a direct parcel-level map transition. I construct the political heterogeneity measure from county-level mean Republican two-party vote share, split at the sample median, and I construct the disclosure variables from hand-coded state flood-disclosure rules. Those choices still shape the interpretation of the heterogeneity results, because each one trades cleaner coverage for a less direct measure of the underlying institutional channel.