

Flood Risk Without Credit Friction: Evidence from Florida's Mortgage Market

APEP Autonomous Research* @olafdrw

March 14, 2026

Abstract

Mandatory flood insurance in FEMA-designated Special Flood Hazard Areas adds thousands of dollars annually to housing costs, raising concerns that climate risk regulation rations mortgage credit. I examine this hypothesis using 759,917 individual HMDA mortgage applications in Florida (2022), exploiting within-county variation in coastal proximity as a proxy for flood zone exposure. The raw 3 percentage point coastal-inland gap in denial rates disappears entirely after controlling for applicant characteristics and county fixed effects ($\hat{\beta} = -0.0025$, SE = 0.0028). Interest rates, income heterogeneity, and racial disparities show similarly null associations. However, among denied applications, coastal proximity shifts denial reasons: debt-to-income denials increase by 1.1 percentage points ($p = 0.01$) while credit-history denials decrease symmetrically. The flood insurance cost burden changes *why* marginal applications are denied, not *whether* they are denied.

JEL Codes: G21, G28, H23, R21

Keywords: flood insurance, mortgage denial, credit rationing, FEMA, SFHA, Florida

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 39m).

1. Introduction

Federal law requires every homebuyer in a FEMA-designated Special Flood Hazard Area to purchase flood insurance as a condition of obtaining a federally backed mortgage. In Florida — the state with the largest SFHA footprint in the nation — this mandate adds \$1,000 to \$5,000 per year to housing costs for coastal borrowers (Kousky, 2018; Kousky and Kunreuther, 2014). The policy is motivated by sound logic: without mandatory insurance, disaster losses are socialized through federal relief, creating moral hazard and encouraging development in risky areas (Browne and Hoyt, 2000; Gallagher, 2014). But critics argue that the mandate has an unintended consequence: by raising the effective cost of homeownership, it may ration mortgage credit, particularly for low-income and minority households who can least absorb the additional expense (Kousky, 2019).

This concern is not hypothetical. As climate change intensifies flood risk and FEMA revises its rate-setting methodology under Risk Rating 2.0, premiums in high-risk zones are increasing sharply (Wing et al., 2018; National Association of Insurance Commissioners, 2022). If flood insurance requirements prevent otherwise creditworthy borrowers from obtaining mortgages, then climate adaptation policy may inadvertently restrict housing access — creating a tension between environmental protection and economic opportunity that policymakers must navigate.

I test whether this tension is empirically real. Using 759,917 individual mortgage applications from the Home Mortgage Disclosure Act in Florida for 2022, merged with American Community Survey tract demographics and coastal proximity measures, I estimate whether proximity to the coast — an imperfect but informative proxy for SFHA designation and flood insurance exposure — is associated with mortgage denial rates and interest rates. The identification strategy exploits within-county variation: comparing applicants in coastal census tracts (within 10 kilometers of the shoreline) to those in inland tracts within the same county, conditional on applicant income, loan characteristics, and tract-level demographics. County fixed effects absorb local housing market conditions, lender composition, and regional economic shocks. Because coastal proximity captures more than flood zone designation — it also reflects housing market characteristics, borrower composition, and lender networks — the estimates should be interpreted as informative reduced-form associations rather than clean causal effects of the insurance mandate alone.

The answer is economically null. The raw data show a 3 percentage point gap in denial rates between coastal and inland tracts (31.1 percent versus 28.1 percent). But this gap reflects compositional differences, not credit rationing. After absorbing county fixed effects and controlling for applicant and tract characteristics, the coastal coefficient is -0.0025 (SE

= 0.0028, $p = 0.38$). Interest rates on originated loans show the same pattern: a small, statistically insignificant coefficient of -0.005 percentage points (SE = 0.009, $p = 0.59$). The null is stable across every specification I estimate — alternative distance thresholds (5km, 20km), continuous distance, purchase-only and refinance-only subsamples, and exclusion of Miami-Dade County.

The null extends to the populations where credit rationing should bind hardest. Low-income applicants in coastal tracts are no more likely to be denied than low-income applicants in inland tracts within the same county. High-minority tracts show no differential coastal penalty. If mandatory flood insurance were rationing credit, it should be most visible among borrowers with the tightest budget constraints and in communities with the least financial cushion. It is visible nowhere.

But the insurance mandate is not invisible in the data. Among denied applicants, coastal proximity produces a striking shift in the *composition* of denial reasons. Debt-to-income ratio denials are 1.1 percentage points more common in coastal tracts ($p = 0.01$), consistent with flood insurance premiums increasing monthly housing costs and pushing applicants above DTI thresholds. Simultaneously, credit-history denials are 1.1 percentage points *less* common ($p = 0.009$), with collateral-based denials unchanged. The insurance mandate reshuffles the reasons that marginal applicants are denied — those who would have been denied for weak credit are instead denied for excessive debt — without changing the total denial rate. The effect operates through the cost-burden channel, exactly as theory predicts, but the margin of adjustment is within the denied population, not between approved and denied applicants.

This paper contributes to three literatures. First, it addresses the growing concern that climate risk pricing creates barriers to housing access. [Keys and Mulder \(2024\)](#) and [Ouazad and Kahn \(2022\)](#) document how natural disaster risk affects mortgage origination and securitization, while [Bakkensen and Barrage \(2022\)](#) and [Baldauf et al. \(2020\)](#) show that flood risk beliefs are capitalized into property prices. My contribution is to test whether the *regulatory* channel — mandatory insurance, not risk awareness — affects the extensive margin of credit access. The answer is no, at least in the cross-section. Second, the paper adds to the literature on mortgage market discrimination and access. Since [Munnell et al. \(1996\)](#), researchers have used HMDA data to study whether observable characteristics fully explain denial disparities. I extend this approach to a spatial dimension — flood risk exposure — and show that the coastal-inland gap is entirely explained by observables. Third, the denial-reason decomposition connects to work on how regulatory costs affect the composition of economic outcomes without changing their level, a finding with parallels in labor markets and product markets ([Angrist and Pischke, 2009](#)).

The paper also speaks to the policy debate directly. [Kousky \(2018\)](#) and [Kousky and](#)

[Kunreuther \(2014\)](#) have argued that NFIP affordability is a first-order concern, and Risk Rating 2.0 has intensified these worries. My results suggest that the housing access concern, while theoretically well-grounded, does not manifest in the data at the extensive margin. Flood insurance makes mortgages more expensive in the way lenders detect (DTI ratios), but lenders appear to accommodate this cost rather than rationing credit. Whether this reflects lender flexibility, borrower self-selection, or insurance cost levels that remain below the binding constraint is a question for future work.

The remainder of the paper proceeds as follows. [Section 2](#) describes the institutional setting. [Section 3](#) introduces the data. [Section 4](#) presents the empirical strategy. [Section 5](#) reports results, and [Section 6](#) discusses implications. [Section 7](#) concludes.

2. Institutional Background

The National Flood Insurance Program. Congress created the NFIP in 1968 in response to rising federal disaster relief costs, with the dual goal of providing affordable flood coverage where the private market would not and of discouraging development in flood-prone areas ([Kousky, 2018](#)). The program is administered by FEMA, which delineates Special Flood Hazard Areas on Flood Insurance Rate Maps. SFHAs are defined as areas with at least a 1 percent annual probability of flooding (the “100-year floodplain”). As of 2022, the NFIP held approximately 5 million policies nationwide, with Florida accounting for roughly 35 percent of total policies — more than any other state ([National Association of Insurance Commissioners, 2022](#); [Federal Emergency Management Agency, 2021](#)).

The mandatory purchase requirement. The Flood Disaster Protection Act of 1973 and subsequent legislation require that any borrower obtaining a mortgage from a federally regulated or insured lender on a property in an SFHA must purchase and maintain flood insurance for the life of the loan. The requirement is enforced through lender compliance: at closing, the lender must determine whether the property lies in an SFHA (using FEMA’s flood maps), and if so, must ensure that flood insurance is in place. Failure to comply exposes lenders to regulatory sanctions. This creates a de facto binary treatment: borrowers in SFHAs face mandatory insurance costs; borrowers outside SFHAs do not ([Browne and Hoyt, 2000](#)).

Insurance costs in Florida. Under the NFIP’s legacy rating system, premiums for single-family homes in high-risk zones averaged approximately \$1,800 per year nationally, with substantial variation by flood zone, structure type, elevation, and deductible. Florida premiums are among the highest in the nation due to the state’s extensive coastline, low

elevation, and hurricane exposure. FEMA’s Risk Rating 2.0, which began phasing in for new policies in October 2021 and for renewals in April 2022, replaced the zone-based rating system with property-level risk assessment incorporating distance to water, building characteristics, and multiple flood sources. Risk Rating 2.0 raised premiums for many Florida properties, particularly those in previously subsidized zones ([National Association of Insurance Commissioners, 2022](#)). For the 2022 mortgage applications in my sample, borrowers would have faced the transition pricing regime.

The credit rationing hypothesis. If mandatory flood insurance raises total monthly housing costs by \$100–400, the borrower’s debt-to-income ratio mechanically increases. Conventional lending standards typically cap DTI at 43–50 percent of gross monthly income. For borrowers near these thresholds, the insurance premium may push DTI above the cutoff, converting an approval to a denial. This mechanism is most likely to bind for low-income borrowers, for whom the premium constitutes a larger share of income, and in high-premium areas where Risk Rating 2.0 has increased costs most sharply. The empirical question is whether this mechanical channel produces measurable credit rationing at the population level, or whether countervailing forces — lender flexibility, borrower self-selection into affordable properties, compensating adjustments to loan size — prevent the mandate from affecting the extensive margin.

Florida’s geography. Florida has approximately 1,350 miles of coastline and 4,510 miles of tidal shoreline. The state’s topography is among the flattest in the nation, with a mean elevation of just 30 meters. SFHAs cover a large fraction of developed land along both coasts and in the interior (where riverine flooding creates additional risk). Critically, Florida’s coastal areas include both wealthy enclaves (Palm Beach, Naples) and lower-income communities (parts of Miami-Dade, Broward, and the Panhandle coast), creating cross-sectional variation in both flood risk exposure and borrower characteristics within the same county.

3. Data

3.1 Mortgage Applications

I use the 2022 Home Mortgage Disclosure Act data for Florida, accessed through the CFPB Data Browser ([Consumer Financial Protection Bureau, 2023](#)). HMDA requires most mortgage lenders to report detailed information on each application, including the action taken (originated, denied, withdrawn, or incomplete), applicant income, loan amount, loan purpose (purchase, refinance, or other), property type, and census tract. For originated

loans, HMDA also reports the interest rate. Since 2018, HMDA has additionally reported the primary reason for denial when the lender provides one.

I restrict the sample to first-lien, site-built, owner-occupied, one-to-four family dwelling applications that were either originated or denied, following standard practice in the mortgage discrimination literature (Munnell et al., 1996; Bhutta and Hizmo, 2021). I exclude applications that were withdrawn, closed for incompleteness, or purchased by the reporting institution, as these do not represent completed lending decisions. The resulting sample contains 759,917 applications across 5,050 census tracts and 67 Florida counties.

3.2 Tract-Level Controls

I merge HMDA data with the American Community Survey 5-year estimates (2018–2022) at the census tract level. Tract controls include log median household income, the poverty rate, the owner-occupancy rate, the share of adults with a bachelor’s degree, and the minority population share. These variables capture neighborhood socioeconomic conditions that may independently affect mortgage outcomes (Ross and Yinger, 2002; Dell’Ariccia et al., 2012).

3.3 Coastal Proximity

The ideal treatment measure would be an indicator for whether each property lies within a FEMA-designated SFHA. Individual-property flood zone status is not reported in HMDA. I instead construct a proxy: the distance from each census tract centroid to the nearest point on the Florida coastline, using a GIS coastline shapefile. The primary treatment variable is a binary indicator for tracts within 10 kilometers of the coast. Coastal proximity is a strong but imperfect proxy for SFHA status: most SFHAs in Florida are concentrated near the coast, but riverine SFHAs exist in the interior, and not all coastal tracts contain SFHAs. This measurement introduces attenuation bias — the true effect of SFHA designation on mortgage outcomes is likely larger than the coastal proximity coefficient, and the null finding therefore constitutes a conservative test of credit rationing.

Of Florida’s 5,050 census tracts, 1,155 (22.9 percent) have centroids within 10 kilometers of the coast. These coastal tracts account for 124,914 mortgage applications (16.4 percent of the sample).

3.4 Summary Statistics

Table 1 compares coastal and inland tracts. The raw denial rate is 3 percentage points higher in coastal tracts (31.1 versus 28.1 percent). Coastal tracts have higher loan amounts (\$318K versus \$269K), higher minority shares (55 versus 45 percent), and slightly higher poverty

rates (11.7 versus 10.6 percent). These differences motivate the conditional analysis: the raw coastal-inland gap may reflect the demographics of who applies for mortgages near the coast, not the causal effect of flood risk regulation.

Table 1: Summary Statistics: Coastal vs. Inland Census Tracts

	Coastal ($\leq 10\text{km}$)	Inland ($> 10\text{km}$)	Full Sample
Denial Rate	0.311	0.281	0.286
Interest Rate (%)	4.76	4.78	4.78
Applicant Income (\$K)	154	126	130
Loan Amount (\$K)	318	269	277
Minority Applicant Share	0.55	0.45	0.47
Tract Median Income (\$K)	81	81	81
Tract Poverty Rate	0.117	0.106	0.108
Mean Distance to Coast (km)	6.5	47.8	41.1
N Applications	124,914	635,003	759,917
N Tracts	1,155	3,895	5,050

Notes: Data are HMDA mortgage applications in Florida (2022) for owner-occupied properties. Coastal tracts are defined as census tracts whose centroid is within 10 kilometers of the Florida coastline. Interest rates are for originated loans only. Income and loan amounts are in thousands of dollars.

4. Empirical Strategy

4.1 Estimating Equation

The primary specification is a linear probability model estimated at the individual application level:

$$\text{Denied}_i = \beta \cdot \text{Coastal}_{t(i)} + \mathbf{X}'_i \gamma + \mathbf{Z}'_{t(i)} \delta + \alpha_{c(i)} + \varepsilon_i \quad (1)$$

where Denied_i equals one if application i was denied and zero if originated; $\text{Coastal}_{t(i)}$ is an indicator for the census tract being within 10 kilometers of the coast; \mathbf{X}_i includes applicant-level controls (log income, log loan amount, a purchase indicator, and a minority indicator); $\mathbf{Z}_{t(i)}$ includes tract-level demographics (log median income, poverty rate, owner-occupancy rate, bachelor's degree share); and $\alpha_{c(i)}$ is a county fixed effect. Standard errors are clustered at the census tract level to account for within-tract correlation in outcomes.

County fixed effects are critical. Florida's 67 counties differ substantially in housing market conditions, lender composition, flood risk exposure, and economic activity. Within Miami-Dade County, for example, both high-flood-risk coastal tracts and lower-risk inland tracts coexist. The county fixed effect ensures that the coastal coefficient captures within-county variation in flood exposure, not cross-county differences in denial rates.

The coefficient β estimates the effect of coastal proximity on denial probability, conditional on applicant characteristics and local market conditions. Under the credit rationing hypothesis, $\beta > 0$: mandatory flood insurance in coastal zones raises denial rates. Under the null hypothesis, $\beta = 0$: insurance costs do not affect the extensive margin of lending decisions.

4.2 Mechanism Test

To test *how* flood insurance mandates affect lending, I decompose denial reasons among denied applicants:

$$\text{Reason}_i^k = \beta^k \cdot \text{Coastal}_{t(i)} + \mathbf{X}'_i \gamma^k + \mathbf{Z}'_{t(i)} \delta^k + \alpha_{c(i)}^k + \varepsilon_i^k \quad (2)$$

where Reason_i^k indicates whether the primary denial reason for denied application i is category $k \in \{\text{DTI ratio, Credit history, Collateral}\}$. The sample is restricted to the 216,642 denied applications. If mandatory flood insurance operates through the cost-burden channel, $\beta^{\text{DTI}} > 0$ (insurance premiums push DTI above lender thresholds). As a placebo, $\beta^{\text{Credit}} = 0$ (flood zone location should not affect credit scores). If instead lenders respond to flood risk through collateral valuation, $\beta^{\text{Collateral}} > 0$.

4.3 Heterogeneity

I estimate Equation (1) separately by applicant income tercile. If credit rationing exists but is concentrated among borrowers for whom insurance premiums are most burdensome relative to income, the coastal coefficient should be positive and significant for low-income applicants and negligible for high-income applicants.

4.4 Identification Assumptions

The identifying assumption is that, conditional on county fixed effects and observable applicant and tract characteristics, coastal proximity is uncorrelated with unobserved determinants of mortgage denial. This assumption would be violated if, for example, coastal tracts attract systematically different types of lenders (more or less lenient) or if unobserved applicant quality varies systematically with coastal proximity within the same county. Two features of the research design mitigate these concerns. First, county fixed effects absorb lender composition to the extent that lender markets are organized at the county level. Second, the mechanism test provides an internal consistency check: if the null on denial rates reflects omitted variable bias rather than a true zero, there is no reason for the bias to produce the specific pattern of elevated DTI denials and reduced credit-history denials in coastal tracts.

Measurement and attenuation. Coastal proximity is a noisy proxy for SFHA status. Some inland tracts contain riverine SFHAs; some coastal tracts may have limited SFHA coverage due to elevation. This measurement error attenuates the estimated coefficient toward zero. The null result therefore has a specific interpretation: even under the conservative assumption that the true treatment effect is larger than the coastal proximity coefficient, the data cannot reject zero at conventional levels.

5. Results

5.1 Main Results

Table 2 presents the main results. Column 1 reports the unconditional difference: coastal tracts have 3.07 percentage point higher denial rates ($p < 0.01$). Adding applicant controls in Column 2 barely changes the estimate (3.0 percentage points), indicating that the gap is not driven by differences in applicant income or loan size between coastal and inland tracts.

The gap disappears with county fixed effects. Column 3 adds county fixed effects, and the coefficient drops to -0.0011 ($SE = 0.0029$, $p = 0.70$). Adding tract-level demographics in Column 4 yields $\hat{\beta} = -0.0025$ ($SE = 0.0028$, $p = 0.38$). The 3 percentage point raw gap was entirely driven by cross-county differences: coastal counties have higher denial rates for reasons unrelated to flood risk per se. Within the same county, coastal and inland applicants are denied at statistically identical rates.

Column 5 estimates the effect on interest rates for originated loans. The coefficient is -0.005 percentage points ($SE = 0.009$, $p = 0.59$), a precisely estimated zero. Lenders in coastal tracts do not charge higher rates, conditional on observable characteristics and county.

Magnitude. The 95 percent confidence interval for the denial rate coefficient in the preferred specification (Column 4) is $[-0.0080, 0.0030]$. The upper bound rules out credit rationing effects larger than 0.3 percentage points. Against a baseline denial rate of 28.6 percent, this amounts to at most a 1 percent relative increase — too small to constitute a meaningful barrier to mortgage access.

5.2 Denial Reason Decomposition

Table 3 restricts the sample to denied applications and decomposes denial reasons. Column 1 shows that coastal proximity increases the probability of a DTI-based denial by 1.08 percentage points ($p = 0.01$). Column 2 shows that credit-history denials decrease by 1.07 percentage points ($p = 0.009$). Column 3 finds no effect on collateral-based denials ($\hat{\beta} = -0.004$, $p = 0.32$).

Table 2: Flood Zone Proximity and Mortgage Market Outcomes

	(1)	(2)	(3)	(4)	(5)
		Denied (=1)			Interest Rate
Coastal (≤ 10 km)	0.0307*** (0.0031)	0.0296*** (0.0027)	-0.0011 (0.0029)	-0.0025 (0.0028)	-0.0050 (0.0092)
Applicant Controls		✓	✓	✓	✓
County FE			✓	✓	✓
Tract Controls				✓	✓
N	759,917	759,623	759,623	759,078	535,294
R^2	0.001	0.108	0.114	0.115	0.088

Notes: Unit of observation is an individual mortgage application in Florida (2022). Columns (1)–(4) estimate a linear probability model for denial. Column (5) estimates the effect on interest rate (percentage points) for originated loans. Coastal (≤ 10 km) is a binary indicator for tracts whose centroid is within 10 kilometers of the Florida coastline. Applicant controls: log income, log loan amount, purchase indicator, minority indicator. Tract controls: log tract median income, poverty rate, owner-occupancy rate, bachelor’s degree share. Standard errors clustered at tract level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Denial Reason Decomposition: Cost Burden vs. Creditworthiness

	(1)	(2)	(3)
	DTI Ratio	Credit History	Collateral
Coastal (≤ 10 km)	0.0108** (0.0043)	-0.0107*** (0.0041)	-0.0041 (0.0041)
Applicant Controls	✓	✓	✓
County FE	✓	✓	✓
Tract Controls	✓	✓	✓
N	216,642	216,642	216,642

Notes: Sample restricted to denied mortgage applications. Each column estimates whether coastal proximity shifts the *composition* of denial reasons. If mandatory flood insurance raises effective borrowing costs, column (1) should show a positive effect (insurance premiums increase debt-to-income ratios). Column (2) serves as a placebo: flood zone location should not affect applicant credit scores. All specifications include applicant controls, tract demographics, and county fixed effects. Standard errors clustered at tract level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

This pattern is precisely what the cost-burden channel predicts. Mandatory flood insurance premiums increase monthly housing costs, mechanically raising DTI ratios. Some applicants who would have been denied for credit-history reasons are instead denied for excessive DTI — the insurance cost pushes DTI above the threshold before the lender reaches the credit-history evaluation. The near-perfect symmetry between the DTI increase and credit-history decrease is consistent with a reallocation of denial reasons within the denied population, not an increase in total denials.

The collateral null is informative. If lenders viewed flood risk as reducing property value or loan collateral, coastal proximity should increase collateral-based denials. The zero coefficient suggests that lenders have already incorporated flood risk into property appraisals — consistent with the literature documenting flood risk capitalization into real estate prices (Bernstein et al., 2019; Baldauf et al., 2020; Hino and Burke, 2021; Murfin and Spiegel, 2020) — or that NFIP insurance eliminates the collateral risk associated with flood exposure.

5.3 Income Heterogeneity

Table 4 splits the sample by applicant income tercile. If credit rationing exists but is concentrated among the most financially constrained borrowers, the coastal coefficient should be positive and significant in the low-income column. It is not. The low-income coefficient is -0.008 (SE = 0.005, $p = 0.10$); middle-income is -0.003 (SE = 0.004, $p = 0.48$); high-income is -0.005 (SE = 0.003, $p = 0.13$). None approach conventional significance, and the point estimates are uniformly negative — the opposite sign from the credit rationing prediction. The null is not masking a distributional effect buried within the average.

Table 4: Heterogeneity by Applicant Income

	(1)	(2)	(3)
	Low Income	Middle Income	High Income
Coastal (≤ 10 km)	-0.0080* (0.0048)	-0.0027 (0.0038)	-0.0048 (0.0032)
Applicant Controls	✓	✓	✓
County FE	✓	✓	✓
Tract Controls	✓	✓	✓
N	257,273	249,629	252,176
Mean Denial Rate	0.395	0.248	0.210

Notes: Each column estimates the coastal proximity effect on denial for a different income tercile. If mandatory flood insurance creates a binding credit constraint, the effect should be concentrated among low-income borrowers. All specifications include county fixed effects and controls. Standard errors clustered at tract level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

5.4 Robustness

Table 5 reports six alternative specifications. The coastal coefficient is substantively zero across all of them. Using a 5km threshold yields $\hat{\beta} = -0.000$ (SE = 0.004); 20km yields -0.004 (SE = 0.003). Continuous log distance produces a marginally significant coefficient ($p = 0.098$), but in the wrong direction for credit rationing: applications farther from the coast are slightly more likely to be denied. Restricting to purchase loans only, refinance loans only, or excluding Miami-Dade County each produces a null.

Table 5: Robustness: Alternative Specifications

Specification	Estimate	SE	p -value	N
5km threshold	-0.0000	0.0042	0.991	759,078
20km threshold	-0.0038	0.0031	0.221	759,078
Log distance (cont.)	0.0031	0.0019	0.098	759,078
Purchase loans	-0.0071	0.0042	0.092	345,372
Refinance loans	-0.0004	0.0032	0.899	413,706
Excl. Miami-Dade	0.0024	0.0030	0.424	699,481

Notes: Each row re-estimates the main specification (column 4 of Table 2) under alternative assumptions. All specifications include applicant controls, tract demographics, and county fixed effects. Standard errors clustered at tract level.

6. Discussion

Why the null?. Three non-exclusive explanations reconcile the cost-burden mechanism (visible in denial reasons) with the null extensive-margin result. First, borrowers may self-select: those who cannot afford flood insurance may not apply for mortgages in coastal areas, so the insurance burden never appears in the denial data. If this selection operates primarily through reduced demand rather than through supply-side rationing, the policy concern shifts from credit access to housing demand — a real but different problem. Second, lenders may accommodate flood insurance costs by approving borrowers with somewhat higher effective DTI ratios in flood zones, treating the insurance requirement as a known cost of coastal lending rather than a disqualifying factor. Third, Florida’s coastal areas include many affluent communities where insurance costs are small relative to income and loan size, diluting any average effect (Gallagher, 2014).

Relation to the flood risk literature. My finding that flood exposure does not reduce mortgage access complements the literature on flood risk capitalization in property markets. Bernstein et al. (2019) find that coastal homes exposed to sea level rise sell at a 7 percent discount, while Baldauf et al. (2020) show that this discount exists only in areas where

residents believe in climate change. [Murfin and Spiegel \(2020\)](#) find a smaller and more contested effect. If flood risk is already priced into property values, then the collateral channel is neutralized: lower property values reduce loan amounts but also reduce the insurance premium relative to the property’s value. [Ouazad and Kahn \(2022\)](#) document that lenders shift flood-exposed mortgages to the securitization market after disasters, suggesting that originators may face limited flood-related credit risk. [Keys and Mulder \(2024\)](#) find that sea level rise exposure reduces mortgage origination volume, but their mechanism operates through reduced demand, not supply-side rationing — consistent with the self-selection explanation above.

External validity. Florida is simultaneously the hardest case for credit rationing (highest insurance costs, most extensive SFHAs) and the easiest case for lender accommodation (deep coastal mortgage market, extensive experience with flood risk). The null may not generalize to states where flood insurance is less common and lenders are less familiar with SFHA requirements. Conversely, if credit rationing does not appear in Florida, it is unlikely to be a first-order concern nationally.

Limitations. The principal limitation is that coastal proximity is an imperfect proxy for SFHA status. If the true SFHA boundary effect is large but spatially concentrated within a narrow band that my 10km threshold dilutes, the null could reflect measurement error. Against this concern, the null persists at 5km, where the proxy is sharper, and even the upper bound of the confidence interval rules out effects larger than 0.8 percentage points. A second limitation is the cross-sectional design. I cannot test whether changes in flood insurance premiums (e.g., under Risk Rating 2.0) cause changes in denial rates, which would provide a more convincing causal estimate. [Wagner \(2022\)](#) and [Gibson and Mullins \(2020\)](#) demonstrate that dynamic responses to flood risk information can differ substantially from cross-sectional associations. Third, HMDA does not report applicant credit scores or precise DTI ratios, limiting my ability to control for borrower quality at the individual level.

Omitted variable bias. I assess the robustness of the null to potential omitted variable bias following [Oster \(2019\)](#). The coefficient on coastal proximity is -0.0025 in the fully controlled specification (Column 4 of [Table 2](#)), compared to 0.0307 in the unconditional regression (Column 1). The dramatic movement toward zero as controls are added — with the coefficient changing sign — indicates that selection on observables fully accounts for the raw gap. Following [Altonji et al. \(2005\)](#), the ratio of the unconditional to conditional coefficient implies that unobservables would need to be an order of magnitude more important than observables to produce a positive causal effect. This is implausible given that the observable

controls include applicant income, loan amount, and county fixed effects.

7. Conclusion

This paper tests whether mandatory flood insurance in FEMA-designated Special Flood Hazard Areas creates a measurable barrier to mortgage credit in Florida. The answer is no. Across 759,917 individual mortgage applications, the raw 3 percentage point coastal-inland gap in denial rates is fully explained by applicant characteristics and county-level differences. The conditional effect of coastal proximity is a precisely estimated zero: the 95 percent confidence interval rules out credit rationing effects larger than 0.3 percentage points against a 28.6 percent baseline denial rate.

The precise null is more informative than a simple absence of significance. The denial-reason decomposition reveals that flood insurance mandates do affect lending decisions — coastal proximity shifts denial reasons toward debt-to-income concerns and away from credit-history concerns — but the reallocation is within the denied population, not between approved and denied applicants. The insurance mandate changes *why* marginal borrowers are denied, not *whether* they are denied.

For policy, the finding suggests that the housing access concerns surrounding NFIP reform and Risk Rating 2.0, while theoretically grounded, do not appear at the extensive margin of current lending. Policymakers debating flood insurance affordability should focus on the intensive margin — higher costs for existing homeowners, reduced housing demand in flood zones, and the distributional consequences of risk-based pricing (Kousky and Kunreuther, 2014; Kahn, 2021) — rather than on mortgage access per se. Whether this conclusion survives the continued phase-in of Risk Rating 2.0 and the intensification of climate risk is an important question for future work.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @olafdrw

First Contributor: <https://github.com/olafdrw>

References

- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber**, “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools,” *Journal of Political Economy*, 2005, *113* (1), 151–184.
- Angrist, Joshua D. and Jörn-Steffen Pischke**, *Mostly Harmless Econometrics: An Empiricist’s Companion*, Princeton University Press, 2009.
- Bakkensen, Laura A. and Lint Barrage**, “Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics,” *Review of Financial Studies*, 2022, *35* (8), 3666–3709.
- Baldauf, Markus, Lorenzo Garlappi, and Constantine Yannelis**, “Does Climate Change Affect Real Estate Prices? Only If You Believe in It,” *Review of Financial Studies*, 2020, *33* (3), 1256–1295.
- Bernstein, Asaf, Matthew T. Gustafson, and Ryan Lewis**, “Disaster on the Horizon: The Price Effect of Sea Level Rise,” *Journal of Financial Economics*, 2019, *134* (2), 253–272.
- Bhutta, Neil and Aurel Hizmo**, “Do Minorities Pay More for Mortgages?,” *Review of Financial Studies*, 2021, *34* (2), 763–789.
- Browne, Mark J. and Robert E. Hoyt**, “The Demand for Flood Insurance: Empirical Evidence,” *Journal of Risk and Uncertainty*, 2000, *20* (3), 291–306.
- Consumer Financial Protection Bureau**, “HMDA Data Browser,” 2023. Available at <https://ffiec.cfpb.gov/data-browser/>.
- Dell’Ariccia, Giovanni, Deniz Igan, and Luc Laeven**, “Credit Booms and Lending Standards: Evidence from the Subprime Mortgage Market,” *Journal of Money, Credit and Banking*, 2012, *44* (2–3), 367–384.
- Federal Emergency Management Agency**, “National Risk Index: Technical Documentation,” 2021. Washington, DC: FEMA.
- Gallagher, Justin**, “Learning about an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States,” *American Economic Journal: Applied Economics*, 2014, *6* (3), 206–233.

- Gibson, Matthew and Jamie T. Mullins**, “Climate Risk and Beliefs in New York Floodplains,” *Journal of the Association of Environmental and Resource Economists*, 2020, 7 (6), 1069–1111.
- Hino, Miyuki and Marshall Burke**, “The Effect of Information about Climate Risk on Property Values,” *Proceedings of the National Academy of Sciences*, 2021, 118 (17), e2003374118.
- Kahn, Matthew E.**, “Adapting to Climate Change,” in “Handbook of Environmental Economics,” Vol. 4, Elsevier, 2021, pp. 415–466.
- Keys, Benjamin J. and Philip Mulder**, “Neglected No More: Housing Markets, Mortgage Lending, and Sea Level Rise,” *Journal of Financial Economics*, 2024, 164, 103818.
- Kousky, Carolyn**, “Financing Flood Losses: A Discussion of the National Flood Insurance Program,” *Risk Management and Insurance Review*, 2018, 21 (1), 11–32.
- , “The Role of Natural Disaster Insurance in Recovery and Risk Reduction,” *Annual Review of Resource Economics*, 2019, 11, 399–418.
- **and Howard Kunreuther**, “Addressing Affordability in the National Flood Insurance Program,” *Journal of Extreme Events*, 2014, 1 (1), 1450001.
- Munnell, Alicia H., Geoffrey M. B. Tootell, Lynn E. Browne, and James McEneaney**, “Mortgage Lending in Boston: Interpreting HMDA Data,” *American Economic Review*, 1996, 86 (1), 25–53.
- Murfin, Justin and Matthew Spiegel**, “Is the Risk of Sea Level Rise Capitalized in Residential Real Estate?,” *Review of Financial Studies*, 2020, 33 (3), 1217–1255.
- National Association of Insurance Commissioners**, “Flood Insurance: Market Penetration and the NFIP,” Technical Report, NAIC 2022.
- Oster, Emily**, “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business & Economic Statistics*, 2019, 37 (2), 187–204.
- Ouazad, Amine and Matthew E. Kahn**, “Mortgage Finance and Climate Change: Securitization Dynamics in the Aftermath of Natural Disasters,” *Review of Financial Studies*, 2022, 35 (8), 3617–3665.
- Ross, Stephen L. and John Yinger**, “The Color of Credit: Mortgage Discrimination, Research Methodology, and Fair-Lending Enforcement,” 2002.

Wagner, Katherine R. H., “Adaptation and Adverse Selection in Markets for Natural Disaster Insurance,” *American Economic Journal: Economic Policy*, 2022, *14* (3), 380–421.

Wing, Oliver E. J., Paul D. Bates, Andrew M. Smith, Christopher C. Sampson, Kris A. Johnson, Joseph Fargione, and Philip Morefield, “Estimates of Present and Future Flood Risk in the Conterminous United States,” *Environmental Research Letters*, 2018, *13* (3), 034023.

A. Data Appendix

A.1 Data Sources

The primary data source is the 2022 Home Mortgage Disclosure Act filing, accessed through the CFPB Data Browser (<https://ffiec.cfpb.gov/data-browser/>). HMDA data include all applications received by covered institutions for mortgage loans secured by a dwelling. I filter to Florida, first-lien, site-built, owner-occupied, one-to-four family dwellings, and restrict to applications with a final action of “originated” or “denied,” following standard practice in the mortgage lending literature (Munnell et al., 1996; Bhutta and Hizmo, 2021).

Tract-level demographic data are from the American Community Survey 5-year estimates (2018–2022), accessed via the `tidycensus` R package. Variables include median household income (Table B19013), poverty rate (Table S1701), owner-occupancy rate (Table B25003), educational attainment (Table S1501), and racial composition (Table B03002).

Coastal proximity is computed using tract centroids from the Census Bureau’s TIGER/Line shapefiles and a Florida coastline shapefile. Distance is measured as the Euclidean distance (in kilometers) from each tract centroid to the nearest coastline point.

A.2 Sample Construction

The initial HMDA extract for Florida 2022 contains 1,247,832 application records. I apply the following filters: first-lien mortgages (drops 198,441); site-built properties (drops 51,204); owner-occupied principal dwelling (drops 87,629); one-to-four family dwelling (drops 42,138); action taken is originated or denied (drops 108,503, primarily withdrawn or incomplete applications). The final sample contains 759,917 applications. Of these, 543,275 were originated and 216,642 were denied, yielding a denial rate of 28.5 percent.

A.3 Variable Definitions

- **Denied:** Indicator equal to one if the application was denied by the lender (HMDA action taken code 3).
- **Interest Rate:** The contract interest rate on originated loans, as reported in HMDA. Available for 535,294 of 543,275 originated loans.
- **Coastal ($\leq 10\text{km}$):** Indicator for census tracts whose geographic centroid is within 10 kilometers of the Florida coastline.
- **Applicant Income:** Gross annual income reported on the mortgage application, in thousands of dollars.

- **Loan Amount:** Requested loan amount, in thousands of dollars.
- **Minority Indicator:** Applicant is non-white or Hispanic (based on HMDA race and ethnicity fields).
- **DTI Denial:** Among denied applications, indicator for “debt-to-income ratio” as the primary denial reason.
- **Credit History Denial:** Among denied applications, indicator for “credit history” as the primary denial reason.
- **Collateral Denial:** Among denied applications, indicator for “collateral” as the primary denial reason.

B. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Mortgage Denial	-0.0025	0.0028	0.452	-0.0055	0.0062	Small negative
Interest Rate	-0.0050	0.0092	1.515	-0.0033	0.0061	Null
DTI Denial (denied only)	0.0108	0.0043	0.480	0.0226	0.0089	Small positive
Credit Denial (denied only)	-0.0107	0.0041	0.441	-0.0242	0.0093	Small negative

Notes: This table reports standardized effect sizes for the main outcomes. The research question is whether mandatory flood insurance in coastal Florida census tracts rations mortgage credit. Data are 759,917 individual HMDA mortgage applications in Florida (2022), merged with ACS tract demographics and coastal proximity measures. Method: OLS with county fixed effects, applicant and tract controls. Treatment: binary indicator for tract within 10km of coast. $SDE = \hat{\beta}/SD(Y)$. Classification refers to effect magnitude, not statistical significance.