Insurance Protection Gap *

Parinitha Sastry† Tess Scharlemann‡ Ishita Sen§ Ana-Maria Tenekedjieva ¶

Current Draft: July 2026
Original Draft: May 2024

Abstract

Homeowners insurance is typically viewed as a key tool protecting household wealth against rising climate damages by facilitating rebuilding after disasters. We argue that this view is incomplete for the majority of U.S. homeowners with mortgages. Using newly available individual-level homeowners insurance data spanning the United States over more than a decade, we show that insurance choices are driven by lender rules and serve to protect their exposures. We show that mortgage lenders impose a de facto rule tying coverage to loan balances, which generate pronounced bunching and kinks in insurance coverage choices. This rule alone accounts for nearly two-thirds of the variation in observed coverage, and our results suggest that coverage would be 18–46 percent lower absent mortgage mandates. Furthermore, because mandates do not adjust with prices, total coverage is highly price inelastic. In contrast, excess coverage above lender-imposed minimums, which reflects households' discretionary demand, is strongly price sensitive, particularly among financially constrained borrowers. Finally, we show that these mandates shape post-disaster outcomes by shifting borrowers from delinquency toward prepayment, thereby protecting lender claims while leaving household equity only partially insured. Taken together, the findings show that mortgage mandates are a central institution in homeowners insurance markets, and has important consequences for how climate risk is allocated between households and lenders.

Keywords: Climate Risk, Insurance Protection Gap, Property Insurers, Banks, Mortgages, Household Finance.
JEL Codes: G21, G22, G51, Q54, R31.

---

*An earlier version of the paper circulated with the title "Climate Risk and the U.S. Insurance Gap: Measurement, Drivers and Implications." We thank Boaz Abramson, Matteo Benetton, Lauren Cohen, Tony Cookson, Cameron Ellis, Shan Ge, Ed Glaeser, Robin Greenwood, Mallick Hossain, Ben Keys, Ralph Koijen, Ulrike Malmendier, Philip Mulder, Borghan Narajabad, Monika Piazzesi, Hyeyoon Jung, Adi Sunderam, Motohiro Yogo, and seminar and conference participants at Berkeley, Imperial College, Boston University, Harvard Business School, Columbia Business School, Arizona State University, Babson College, Yale Junior Finance Seminar, NBER Household Finance, NBER IO of Housing Markets Conference, SITE Climate Finance & Banking, NBER Insurance Working Group, WFA, SFS Cavalcade, AFA, RFS-Oklahoma Climate Conference, the Duke-UNC Corporate Finance Conference, UNC Wealth Inequality Conference, NYU Stern Climate Finance Conference, NBER Conference on Research in Income and Wealth for valuable comments and feedback. Quadrant Information Services own the copyright to their respective data, which we use with permission. The views expressed in this paper are those of the authors only and should not be interpreted as reflecting the views of, or implying any responsibility for, the Federal Reserve System.

†Wharton; psastry@wharton.upenn.edu.
‡Federal Reserve Board; Tess.Scharlemann@frb.gov.
§Harvard Business School & CEPR; isen@hbs.edu.
¶Arizona State University; ana-ateneked@asu.edu.

Over the last fifty years, property damage from natural disasters has risen sharply. In the United States, an estimated 35 million homes—roughly one-third of the housing stock—face high exposure to natural disaster risk.[1] Property insurance is therefore increasingly viewed as a central tool for protecting household wealth against climate risk, since insurance payouts can finance the repair or rebuilding of homes damaged by natural disasters.

In this paper, we argue that this view is incomplete for the roughly two-thirds of U.S. homeowners who carry a mortgage. In the United States, homeowners insurance is not chosen by the household alone–it is jointly chosen with mortgage lenders. Lenders are exposed to losses when natural disasters damage mortgaged properties and precipitate borrower default. Lenders, therefore, require borrowers to maintain insurance throughout the life of the loan, specify minimum coverage amounts, and have direct claims to insurance proceeds. As a result, insurance contracts have two economically distinct claimants: the lender, whose exposure is the outstanding mortgage balance, and the household, whose protection depends on the residual coverage left after the lender's claim is satisfied. This institutional feature raises two central questions: how much do mortgage lenders shape insurance demand, and how are the benefits of insurance allocated between lenders and households?

This paper provides the first comprehensive analysis of how households and lenders jointly determine homeowners insurance coverage. We construct a novel loan-level dataset linking insurance policies, mortgage contracts, borrower, and property characteristics, and local climate risk for U.S. households from 2013 to 2023. We show that mortgage requirements are a first-order determinant of insurance coverage. Lenders enforce a de facto rule requiring coverage to exceed the loan balance, protecting their exposure to catastrophic natural disasters. This mandate is quantitatively important, influencing coverage even when it does not strictly bind, and generating sharp kinks and bunching in coverage choices at underwriting thresholds. Importantly, mortgage mandates render overall coverage to be price inelastic, since these requirements do not vary with insurance prices. In contrast, households' discretionary coverage above the lender minimum is highly price sensitive and reduces sharply as premiums rise, particularly for financially constrained borrowers. Finally, we show that a key function of the mortgage mandates is to shift borrowers away from delinquency and toward mortgage prepayment after disasters. This suggests that the mandates protect lender claims through repayment, while leaving households only partially protected to rebuild. Taken together, our findings show that homeowners insurance is organized around the mortgage: mandates keep lender exposure insured, while rising premiums increasingly leave household equity exposed.

The key institutional feature underlying our analysis is the focus on two types of lender requirements that determine how much insurance coverage each household is required to get. The government-sponsored enterprises (GSEs), which back two-thirds of the U.S. mortgage market, have underwriting guidelines which specify minimum coverage requirements to be eligible for GSE securitization. GSE guidelines tie minimum required coverage to the ratio of loan balance to replacement cost, what it costs to rebuild a property in a total loss, L/RC, in a piecewise linear function.

[1] Based on estimates from Cotality, see CNBC, "Nearly one-third of U.S. homes are at high risk of natural disaster," Jan. 27, 2021.

The required coverage equals 80% of replacement cost when $L/RC \leq 0.8$, rises linearly with the loan balance when $0.8 < L/RC < 1$, and equals replacement cost when $L/RC \geq 1$ (Figure 1). In contrast, other mortgage lenders and the Federal Housing Authority require a simpler "de facto" rule, which requires that coverage exceed the minimum of the loan balance and the replacement cost. Importantly, both rules are completely determined by the loan balance and replacement cost, and neither varies with insurance prices or any other household or property characteristics. This distinction between price-insensitive lender requirements and price-sensitive household demand is central to our empirical design.

We organize the analysis around a simple conceptual framework in which observed coverage equals a lender-required component plus discretionary coverage chosen by the household. Coverage above the lender floor is the component that can protect household equity after the mortgage claim is repaid. The framework yields three broad classes of empirical tests. First, if lender requirements matter, coverage levels should comply with lender rules, exhibit kinks at underwriting thresholds, and bunch at required coverage levels. Second, if lender requirements make up a large share of observed coverage, total coverage should be relatively price inelastic, while discretionary coverage above the lender floor which is chosen by households should be much more price sensitive. Third, if insurance coverage protects the mortgage claim after disasters, borrowers with coverage just above the loan balance should have different outcomes than similar borrowers just below the threshold. In particular, they should be more likely to prepay, and less likely to become seriously delinquent, after a severe disaster.

We start by evaluating the first class of predictions about whether mortgage lenders' requirements shape the *level* of insurance coverage. The first test concerns compliance with mortgage lenders' rules: if mortgage mandates are enforced strictly, we should observe high compliance with these rules (Prediction 1). We compare observed coverage to the formal GSE rule and to the de facto loan-balance rule. Interestingly, we find evidence of partial compliance with the GSE rule. While coverage closely tracks GSE-required levels when L/RC is high (above 80%), borrowers with low L/RC (below 80%) have coverage ratios of roughly 70%, below the GSE's 80% requirement. Yet, these borrowers are almost universally in compliance with the de facto rule requiring coverage at least equal or exceed the loan balance. We find that that the mortgage constraint is relatively slack for borrowers with low L/RC, as their insurance coverage exceeds the outstanding loan balance. In contrast, for borrowers with high L/RC, the constraint is much tighter, with households purchasing coverage that is close to the lender's minimum requirement. More broadly, we find that nearly all borrowers are in compliance with the simpler de facto rule, and that this rule alone explains 67% of the variation in coverage, suggesting that it is a quantitatively important determinant of insurance coverage choices.

We next examine a second prediction implied by the de facto rule: if mortgage mandates actively shape coverage, then the sensitivity of coverage to loan balances and replacement costs should change at the underwriting thresholds. In particular, under the de facto rule, the relevant coverage benchmark switches from being the loan balance to the replacement cost when $L/RC = 1$. Thus,

the sensitivity of coverage to the loan amount should change sharply at the point where loan balance exceeds replacement cost. Similarly, the sensitivity of coverage to replacement cost should also change at this same threshold (Prediction 2). Consistent with this prediction, we find a sharp kink at $L/RC = 1$: below the threshold, coverage is substantially more responsive to the loan balance, whereas above it, coverage becomes more responsive to replacement cost. In contrast, we find no comparable kink at the GSE threshold of $L/RC = 80\%$, consistent with the evidence that the formal GSE schedule is not fully enforced in the low-L/RC region. We show that our kink estimates are robust across flexible specifications, bandwidth choices, and a range of placebo and robustness tests, providing causal evidence that mortgage underwriting rules influence household coverage and its sensitivity to loan balances. The kink at $L/RC = 1$ shows that mandates shape coverage choices, but the magnitudes imply that they are not the only force. Because the pass-through estimates are below the one-for-one relationship implied by strict adherence to the de facto rule, observed coverage appears to combine a lender-imposed floor with discretionary household demand above that floor. Our estimates imply that, all else equal, insurance coverage would be between 18–46% lower absent lender mandates tying coverage to the loan balance.

The distribution of coverage provides a third piece of evidence that the level of coverage is shaped by mortgage mandates. We examine whether households bunch at lender-mandated minimum coverage levels. This prediction follows from the possibility that some households have little or no latent demand for insurance and therefore purchase only the minimum coverage required by the mandate. We focus on the region where the loan balance is the binding constraint under the de facto rule. Consistent with this rule, we find that more than 90% of households in this region obtain coverage that exceeds the outstanding loan balance, producing a sharply censored distribution with very little mass below the threshold where coverage equals the loan balance. In addition, we find pronounced bunching, with nearly 10% of households choosing coverage exactly equal to the loan amount. This pattern is difficult to reconcile with smooth household demand alone, indicating that the loan balance acts as a binding underwriting floor for insurance coverage. Taken together, the compliance, kink, and bunching tests show that mortgage market requirements strongly determine household coverage choices.

Next, we analyze a second class of predictions related to the sensitivity of coverage demand to insurance prices. We exploit the key feature that mortgage market imposed coverage requirements depend only on loan balances and replacement costs, not on insurance prices. In contrast, households' underlying demand for insurance should be price sensitive when prices are above actuarially fair rates (Arrow (1963), Mossin (1968), Einav et al. (2010)). If mortgage mandates are an important driver of coverage choices, observed coverage should combine a price-insensitive component determined by mandates with a price-sensitive discretionary component chosen by households. Moreover, coverage should be less price responsive when mortgage market requirements are tighter (high L/RC) and more price responsive when households have more discretion and mortgage constraints are slacker (low L/RC).

To estimate these demand elasticities, we exploit regulatory shocks to insurance prices, following

Oh et al. (2021)). Because insurance pricing is state-regulated, insurers require regulatory approval for rate changes. In certain "high friction" states, this process involves significant, often unpredictable delays, resulting in approved rates that differ substantially from requested ones, generating persistent insurance underwriting losses. In contrast, in "low friction" states, insurers can adjust prices more responsively to changing risks and demand conditions. We instrument prices in low-friction states using contemporaneous prices in high-friction states along with zip-year-credit score fixed effects. This instrument is relevant because multi-state insurers cross-subsidize risks across states due to their financial constraints.[2] We argue the exclusion restriction holds because variation in the within-year timing of regulatory decisions in high-friction states is unlikely to be driven by shocks to demand in low friction states. Thus, the identifying variation exploits differences in the timing of mortgage origination among observationally similar borrowers in the same ZIP code and credit-score. Borrowers originating their mortgages in different months of the same year face different insurance prices, reflecting the timing of rate regulation decisions in other high-friction states.

Using this IV approach, we estimate a small overall price elasticity of approximately -0.088, implying that a 1% increase in insurance prices reduces coverage by less than one-tenth of one percent. This inelastic response is consistent with lender requirements muting the price sensitivity of overall coverage demand. Moreover, this inelasticity is strongest for borrowers facing tighter mortgage requirements (high L/RC), while borrowers with more discretion over their coverage (low L/RC) exhibit somewhat larger, though still modest, elasticities. We then explore heterogeneity in elasticities across borrower characteristics. Financially vulnerable households, defined as those with high debt-to-income ratios, are somewhat more price sensitive than low DTI households, suggesting a role for borrower characteristics in shaping insurance demand, though on average all groups remain largely inelastic. These findings suggest that a large share of observed coverage is pinned down by mortgage market requirements rather than households' underlying demand.

We next isolate the discretionary component of coverage that reflects household intrinsic demand over coverage. We first examine the price sensitivity of "excess coverage," defined as coverage above the lender-required minimum. We limit to borrowers for whom the loan balance is the relevant mortgage required benchmark. In sharp contrast to the overall elasticity, we find that excess coverage is highly price sensitive: a 1% increase in insurance prices reduces discretionary coverage by nearly 2%. This elasticity is more than an order of magnitude larger than that of total coverage, indicating that households respond strongly to insurance prices where they are free to adjust coverage, but lender mandates substantially mute the response of observed coverage. Furthermore, in a related test, we also find that as insurance prices rise, households increasingly reduce coverage to the lender-imposed minimum, evidenced by a sharp increase in bunching at coverage equal to the outstanding loan

[2] Oh et al. (2021) shows that delays in regulatory approval generate persistent underwriting losses in "high-friction" states, raising insurers' financing costs and, in turn, necessitating higher prices in "low-friction" states. This cross-subsidy occurs because financing costs are a key component of insurers' costs (Froot et al. (1999), Koijen and Yogo (2015)), and in turn these costs are a function of insurers' prices in every state. Because most insurers operate across states, persistent and severe losses in high friction states erode insurer capital, which increases their financing costs. In turn, higher financing costs lead insurers to raise prices in less regulated states, making prices in high friction states relevant for prices in low friction states.

balance. This means that rising premiums not only reduce discretionary coverage, but also push some borrowers all the way to the lender minimum. The discretionary response is strongest among high debt-to-income borrowers: financially constrained households reduce excess coverage more sharply and are more likely to bunch at the loan balance when prices rise. Together, these findings show that households are price sensitive where they have discretion, but mortgage market mandates substantially dampen the price sensitivity of overall coverage. As a result, rising insurance prices would primarily erode the component of coverage that protects household equity, while leaving the component of coverage protecting lenders largely intact.

The final part of the paper studies why mortgage-market mandates are organized around the loan balance. To shed light on this question, we study mortgage outcomes after severe disaster exposure and use variation in insurance coverage to distinguish between two functions of the contract: financing household rebuilding and protecting lender claims. Coverage equal to or exceeding the loan balance is sufficient to repay the lender, but not necessarily to rebuild the home. This distinction allows us to test whether insurance coverage operates mainly as household protection for rebuilding or as lender protection through mortgage repayment. The key threshold is whether insurance coverage exceeds the outstanding loan balance ($C > L$). We therefore compare borrowers just above and just below the loan-sufficiency threshold, testing whether those above the threshold are more likely to prepay and less likely to become delinquent after a disaster.

We exploit the landfall of Hurricanes Harvey and Irma, which struck six states in 2017, as a natural experiment. Our empirical design compares mortgage outcomes of borrowers before and after the storms, across severely and less severely affected counties, and within narrow bandwidths around the threshold $C/L = 1$. We ask whether borrowers just above the loan-balance threshold experience different post-disaster outcomes than borrowers just below it. We find that coverage just above the loan amount increases mortgage prepayment and reduces default after disasters, suggesting that lender insurance requirements shift households away from costly default and toward repayment of the loan. A back-of-the-envelope calculation using our estimates suggests that sufficient coverage enables nearly 25% of households experiencing a total loss to prepay their mortgage rather than default. These results show that a central function of mortgage mandates is to protect lender exposures by facilitating prepayment, challenging the view that the main purpose of homeowners insurance is to help households finance rebuilding.

A key question is whether households are better or worse off under the mandates. We argue that the implications for households are nuanced. Prepayment is likely better for borrowers than default or foreclosure, since it avoids missed-payment costs, credit-score damage, and foreclosure-related losses. Thus households benefit from prepaying rather than becoming delinquent or defaulting. However, this does not necessarily restore households to their pre-disaster financial position. If insurance proceeds are used to repay the mortgage, the household retains only the residual amount, $C - L$, together with any remaining value of the damaged property. This residual may be smaller than the amount required to rebuild or the household's accrued equity prior to the disaster. Thus, lender mandates can improve mortgage outcomes after disasters, while still leaving households substantially

exposed to uninsured climate losses.

The central role of mortgage mandates in shaping insurance outcomes raises broader questions about their design and consequences. Our results imply that mortgage lenders effectively impose an "individual mandate" in homeowners insurance markets, ensuring coverage even when households have limited latent demand for insurance. In effect, lender mandates tie the size of the insurance market to the size of the mortgage market. This institutional arrangement may mitigate adverse selection and protect lenders from default losses, but it also shapes the allocation of climate risk across households, insurers, and mortgage markets. In particular, lender mandates limit households' ability to reduce coverage as prices rise, helping to sustain a large insurance market size despite rising costs. At the same time, by addressing climate-related risks through insurance rather than mortgage pricing, lenders may shift risks between markets. While insurance mandates reduce default risk, reducing lender or GSE risks, they can increase prepayment risk borne by mortgage-backed securities investors. More broadly, our results suggest that the current role of insurance mandates reflects a set of institutional constraints, raising questions about how to optimally design both insurance and mortgage markets in a world of rising climate risk.

Related literature. This paper contributes to three strands of the literature. First, we contribute to the literature on the determinants of households' demand for insurance. Much of the theoretical work on insurance demand builds on models of optimal insurance choice, beginning with Arrow (1963) and Mossin (1968), and extending to settings with financial constraints in Froot et al. (1993) and Rampini and Viswanathan (2013), Ericson and Sydnor (2018). Relative to this work, our setting highlights a distinct but closely related force: insurance demand can be shaped not only by households' preferences and constraints, but also by the contractual requirements of their creditors. In our setting, mortgage lenders can require borrowers to maintain insurance coverage as a condition of the loan. This creates a direct link between debt contracts and insurance contracts, raising questions about their complementarity and substitutability. This connection between credit and insurance links our paper to a rich literature in development economics on the optimal bundling of crop insurance and credit under financial and informational frictions (Braverman and Stiglitz, 1982, Banerjee et al., 2014, Casaburi and Willis, 2018, Karlan et al., 2014, Annan, 2021). We show that similar forces are present in the U.S. mortgage market. In particular, we show that lenders use insurance requirements to address monitoring frictions and moral hazard concerns about borrower default after disasters by requiring sufficient coverage as a condition of mortgage credit.

Empirically, we follow the approach of Einav et al. (2010) and Einav and Finkelstein (2023) in estimating demand elasticities. To our knowledge, we are the first to estimate these parameters for the U.S. homeowners insurance market. We are also the first to document a divergence between lender and household demand for insurance. Lender mandates make overall coverage demand nearly inelastic, while household demand for the component of coverage that protects home equity is highly elastic. We further estimate elasticities for discretionary coverage and document bunching at lender mandates, which helps establish this divergence in demand.[3]

[3]There are existing estimates of demand elasticities on the extensive margin for flood insurance in Wagner (2022),

Second, our paper contributes to the small but growing literature on homeowners insurance markets.[4] Thus far, this literature has focused primarily on the supply side of insurance markets, including the determinants of pricing and insurer rationing. Prior work has studied the role of capital market frictions (Froot and O'Connell, 1999, Jaffee and Russell, 1997), state-level price regulation (Oh et al., 2021), informational asymmetries (Boomhower et al., 2023), trust and claims validity (**?**Gennaioli et al., 2021, Jotikasthira et al., 2025), reinsurance pricing (Mulder and Keys, 2024), disaster risk (Blonz et al., 2024, Solomon, 2025), and mortgage market frictions (Sastry et al., 2024). On the demand side, there is earlier work showing that homeowners insurance contract choices are starkly different than choices in auto insurance (Barseghyan et al., 2011, 2013). Similarly, in contemporaneous work, Cookson et al. (2023) study insurance demand among victims of the Marshall Fire in Boulder, Colorado, documenting the role of behavioral biases and shrouded replacement cost characteristics in household choice. Relative to this literature, our contribution is to emphasize the first-order role of mortgage mandates in shaping homeowners insurance coverage demand and to identify lenders as key beneficiaries of insurance protection, a stark difference between homeowners insurance and other markets.

Third, our findings speak to the literature on climate risk and household financial resilience. Existing work has studied the effects of climate and disaster risk on mortgage default (Issler et al., 2020, Kousky et al., 2020, Sastry, 2025, Ge et al., 2024, Biswas et al., 2023), post-disaster financial recovery (Gallagher and Hartley, 2017, Cookson et al., 2023), and house prices (Baldauf et al., 2020, Murfin and Spiegel, 2020, Bernstein et al., 2019).[5] A key finding of our paper is that mortgage mandates tying coverage to loan balances facilitate the use of insurance proceeds to prepay mortgages, improving both lender and household outcomes after natural disasters. This finding is consistent with evidence of elevated prepayment following natural disasters (Gallagher and Hartley, 2017, Issler et al., 2019, Biswas et al., 2023, Kousky et al., 2020). Our contribution is to show that such prepayment is not merely an incidental household response to disasters; it is an intended consequence of mortgage lender mandates, with insurance shifting households from default toward prepayment.

Lastly, our findings speak to a literature on how lenders manage collateral risk. It is generally well-known in the corporate finance literature that collateral availability and quality impacts firm debt capacity and pricing (Rampini and Viswanathan, 2013, Shleifer and Vishny, 1992, Benmelech and Moskovitz, 2010). In the mortgage context, there is significant evidence that lenders use a range of tools beyond pricing to insulate against losses stemming either from collateral deficiency or credit risk. For example, lenders protect themselves using an imperfect appraisal process (Agarwal et al., 2015, Calem et al., 2021, Ambrose et al., 2026), and limit credit availability for properties with particularly uncertain valuations (Jiang and Zhang, 2025), or to borrowers whose default

Collier et al. (2024), and Liao and Mulder (2021). Flood insurance covers flood damage and is required for mortgages only in flood zones, leading to limited take-up nationwide (Amornsiripanitch et al., 2024). In contrast, homeowners insurance is required for all mortgages, and nearly 90% of homeowners purchase homeowners insurance.

[4] See Sastry and Sen (2025) for a recent review of the literature.

[5] Households may also bear indirect climate exposures through labor markets (Kruttli et al., 2019) and municipal bond prices (Goldsmith-Pinkham et al., 2020).

would expose lenders to legal liability under post-GFC Ability to Repay rules (Defusco et al., 2019). Fannie Mae and Freddie Mac require borrowers to purchase mortgage insurance on high-LTV loans - which are most vulnerable to house price volatilility (Bhutta and Keys, 2022), and private lenders starkly limit credit access to high-LTV borrowers. Additionally, lenders limit refinance opportunities for borrowers who appear risky due to job loss (Defusco and Mondragon, 2020). Our relative contribution is to show that mortgage lenders rely heavily on insurance coverage mandates to address concerns about collateral risk stemming from climate damages, consistent with significant non-price adjustments in the mortgage context.

1. Individual-level Data on Insurance and Mortgages

We construct a novel granular database that integrates a wide number of mortgage and insurance datasets to provide a comprehensive joint picture of mortgage and insurance markets. Our extensive database provides loan-level information on the insurance contract (e.g., coverage, premium, insurer), mortgage contract, borrower characteristics, mortgage performance, property characteristics, and climate risk exposures. The dataset spans from 2013 to 2023, encompassing households across the United States. To our knowledge, we are the first to bring together such wide ranging aspects of mortgages, insurance contracts, and property characteristics covering the entire U.S. for over a decade at a granular level. Specifically, our analysis combines the following six databases to obtain a loan-level dataset that includes data on the borrower, property, mortgage, and insurance contract for single-family home purchase mortgages originated on or after 2013. We describe each database briefly below. Appendix A provides further details on the data construction and sample merge.

1.1. Data Sources

(1) ICE McDash Mortgage Modules: We start by using the ICE McDash loan-level dataset on mortgage origination and performance history that is widely used in the mortgage literature. McDash collects these data from mortgage servicers, with the data representing nearly two-thirds of the U.S. mortgage market. The key variables include loan amount, loan-to-value (LTV) ratio, origination month, interest rate, debt-to-income (DTI) ratio, borrower origination credit score, loan maturity, property value, and the type of mortgage (e.g., FHA, VA, jumbo). We merge this with McDash's mortgage performance modules, which contain data on a loan's performance history from origination to final payment. This includes whether the loan is current, the delinquency severity, and also whether a loan has exited the database via prepayment or foreclosure in a given month.

(2) ICE McDash Property Insurance Module: To jointly study insurance choices, we exploit a newly available granular dataset which contains detailed contract-level information on homeowners insurance for individual mortgage borrowers. The database, introduced in 2024, is sourced from ICE McDash and addresses a key gap in the literature, as historically, insurance data could not be linked to mortgages at a loan-level. The dataset covers mortgages from 2013 to 2023 and provides information on monthly insurance premiums, coverage amounts, and deductibles. ICE

McDash collects these data from a subset of mortgage servicers. Table C.1 discussed in Appendix A shows that the sample of mortgages in the McDash insurance module is broadly representative of all McDash loan originations across a range of borrower characteristics.

(3) Equifax data: We obtain additional borrower characteristics from Credit Risk Insights Servicing McDash (CRISM), which anonymously matches Equifax data at the loan level to ICE-McDash servicing records data.

(4) Cotality Deeds Module: Next, we bring in property characteristics, such as the geographic coordinates, structure age, and size of the property (square feet) using the Cotality Deeds database (formerly CoreLogic), which compiles property characteristics from tax assessment records. We anonymously merge these data to the loan-level data using a precise matching algorithm we have developed based on address (ZIP code), origination date, loan amount, loan purpose, and loan type.

(5) Cotality Climate Module: We also integrate key measures of property-level climate risk from the Cotality climate module. Specifically, we use average annual losses (AAL), which represents the annual expected losses of different types of climate-related hazards for the property. Importantly, Cotality's climate module is also widely used by lenders and insurance companies (Boomhower et al., 2023).[6]

(6) MSCI-First Street Foundation and R.S. Means: We incorporate data on property-level replacement costs estimates from the MSCI-First Street Foundation and obtain city-level local construction cost indices from R.S. Means, a widely used database in the urban and housing literature that is developed based on detailed surveys of home builders across metropolitan cities nationwide (e.g., Glaeser and Gyourko, 2005). We link MSCI-First Street Foundation estimates to Cotality deeds using their coordinates, combined with a nearest-neighbor algorithm. We apply R.S. Means indices to our replacement cost estimates from First Street to estimate the replacement cost at mortgage origination. We discuss below how we benchmark and validate these replacement cost measures in more detail.

1.2. Measurement of Key Variables and Sample

Our study centers around three main variables. The primary variable of interest is a household's homeowners insurance **coverage (C)**, defined as the maximum amount a household can claim in the event of a property damage.[7] The second variable is the first-lien **loan balance (L)** on the mortgage. The third variable is the **replacement cost (RC)**, which is the amount it would cost to

[6]For example, the *Wall Street Journal* writes: "CoreLogic [now Cotality], paid by the insurer to advise it on risks, said in an interview it first warned State Farm four years ago of a heightened danger in the Los Angeles neighborhoods that were hit by the latest fires." See Wall Street Journal, "State Farm Was All In on California", February 6, 2026.

[7]Our coverage variable refers to "Coverage A," which is the coverage for the structure of a property. Some homeowners may choose to add-on extended coverage, but they can only do so if their basic coverage is at least equal to 100% of their replacement cost. Klein (2018) writes that extended coverage cannot be provided unless "the basic coverage is at least as great as the estimated replacement cost of the property." Similarly, the insurance consulting firm Marsh McLennan also writes "When you insure your home to 100% of its replacement cost value, some insurance companies will offer the benefit of extended replacement cost. This provision will pay beyond your policy limit should the amount at the time of loss not be adequate," available here. See Appendix Section A.2 for more detail.

rebuild the property if a total loss event completely destroys the structure. If coverage equals rebuild cost, then a homeowner can expect insurance payments to cover all potential costs associated with repair and rebuilding, whereas if coverage is below the rebuild cost, the homeowner is responsible for any repair costs that exceed the coverage limit. As we will discuss in Section 2, these three variables are linked together by mortgage market mandates.

Replacement cost measurement: Unlike coverage and loan balance, which we observe directly in the McDash data, replacement cost is not directly observed and may be measured with error. This measurement challenge is not unique to our data. In actual insurance and mortgage underwriting, market participants also do not observe a single "true" replacement cost. Instead, insurers, lenders, and vendors estimate rebuilding costs using structure characteristics, such as square footage, age, and quality, together with local construction costs and model assumptions. Many insurers rely on third-party valuation tools, including Verisk 360Value, Cotality Marshall & Swift, and First Street Foundation, which provides the replacement cost data used in this paper. We therefore interpret our measure as a proxy for the replacement cost estimates available to insurers and lenders at mortgage origination. The inability to know the true value of replacement costs with certainty creates fundamental challenges for any study of homeowners insurance markets (and likely also for the smooth functioning of these markets), and distinguishes this market from other settings such as health or life insurance.

We conduct several validation exercises that suggest our measure closely aligns with the replacement cost estimates used in practice on average. First, the average implied replacement cost per square foot in our data is $209.60 over 2013–2023, closely matching the National Association of Home Builders estimate of $208.52. Second, the growth rate in our R.S. Means construction-cost index is 4.5%, similar to both Verisk's estimate for the insurance industry, 4.7%, and the BEA producer price index for construction costs, 4.5%. Third, our coverage-to-replacement-cost ratios align with other estimates from the contemporaneous literature.[8] Fourth, our replacement cost estimates vary with property characteristics in expected ways: replacement costs are higher for larger and older homes and in areas with higher construction costs, consistent with insurer and vendor methodologies (Table C.2).[9] Finally, the empirical patterns in Section 3 provide an internal validation: coverage limits align with replacement costs precisely in the regions where mortgage rules make replacement cost the relevant underwriting benchmark, and the relationships between coverage, loan amount, and replacement cost exhibit kinks at the thresholds predicted by mortgage market rules.[10]

[8] Jung and Jung (2025) obtain replacement costs for 2021 from the National Structure Inventory, while Cookson et al. (2023) estimate replacement costs for Colorado using R.S. Means and structure characteristics.

[9] See, for example, the Insurance Information Institute, a trade group for insurers.

[10] One notable feature of the data is that estimated replacement cost exceeds market value for more than 80% of homes. However, this pattern does not appear to reflect systematic measurement error. Insurance coverage itself exceeds market value for the majority of homes, and a replacement-cost–market-value wedge is consistent with the fact that land depreciates more slowly than structures (Glaeser and Gyourko, 2005), especially in an older housing stock. This mechanism predicts that the replacement cost–market value wedge should be larger for older and larger structures, a prediction we verify in Table C.2. Given that the median home in our sample is 40 years old, it is therefore plausible that replacement cost exceeds market value for a large share of properties.

Taken together, these validation exercises suggest that our replacement cost measure captures the object used in underwriting reasonably well on average. At the same time, we do not treat replacement cost as perfectly measured. Throughout the paper, we explicitly consider how measurement error in replacement cost could affect each empirical exercise. In each subsequent section, we address the specific challenges that replacement cost measurement poses for that section's analysis and describe the steps we take to mitigate, sign, or bound any resulting biases.

Insurance Price: The demand estimation in Section 4 also examines insurance prices. Because premiums vary mechanically with coverage, the relevant price of insurance is the monthly insurance premium per $100,000 of coverage, which we obtain from the McDash Insurance module.

1.2.1. Sample and Summary Statistics

Our analysis focuses on understanding the drivers of coverage choices at mortgage *origination*. This is because most of the variation in coverage choice arises at the time of mortgage origination, with the ratio of coverage to replacement costs remaining stable over the life of the mortgage (Figure C.2).[11] The summary statistics in Table 1 shows the distribution of our three main variables (coverage, loan balance, and replacement cost), as well the distribution of each variable as a ratio of the other ($C/L, C/RC, L/RC$), which will be important as we evaluate the extent to which the mortgage mandates matter.

We refer to the ratio of coverage to replacement cost (C/RC) as **coverage ratios**. This variable captures how much households would have to pay out-of-pocket to completely rebuild their properties. Figure C.1 shows a histogram for the distribution of coverage ratio at origination. The majority of the distribution has coverage below replacement costs: half of American households have less than 70% of their rebuild cost covered by insurance, and 10% of households have less than 50% of their rebuild cost covered. Appendix A.3 provides further analysis of coverage ratios in more detail, along with how it relates to key borrower characteristics such as credit scores and income (Figure C.15).

2. Mortgage Market Requirements and Testable Predictions

Mortgage lenders require that borrowers purchase homeowners insurance and obtain a certain minimum amount of coverage. As a result, homeowners insurance contracts are nearly ubiquitous: nearly 93% of U.S. homeowners have a policy, amounting to over $15 trillion in coverage annually (Jeziorski et al., 2021, Oh et al., 2021). The standard contract is annual and covers damages from most natural disasters, except flooding, with 90% of insurance claims caused by weather events.[12] The insurance contract has two key features: coverage and premiums.[13] Households are entitled to

[11] We keep loans for which the first observation occurs within the first two years of the loan. We do so because insurance is not part of the origination information in McDash, but part of the monthly performance data. Over 85% of first observations fall within the first year.

[12] Insurance Information Institute. "Facts + Statistics: Homeowners and renters insurance". III.org; Munich Re: "Climate change is showing its claws".

[13] The deductible, which is another feature of the contract, refers to amount the borrower must pay before insurers start paying for insured events. In the data, there is little variation in deductibles, with most contracts ranging

claim payouts up to a pre-specified maximum coverage limit in case an insured loss event materializes. The premium is the amount that the household pays to the insurer in exchange for this protection.

2.1. Mortgage Market Minimum Coverage Requirements

Mortgage lenders and the government-sponsored enterprises (GSEs) have a direct interest in whether the home is adequately insured, because the home serves as collateral for the loan. As a result, mortgage origination and servicing guidelines set minimum required coverage for households. Specifically, the required coverage depends on the ratio of two key variables: the outstanding loan balance, L, and the property's replacement cost, RC. The ratio L/RC, captures the extent to which the structure itself collateralizes the loan. When the ratio L/RC is low, the structure collaterizes a small portion of the loan balance, and the required coverage is typically lower. When the ratio L/RC is high, the structure collateralizes a large portion of the loan balance, and the required coverage is typically higher.

We focus on two broad types of mortgage market mandates: (i) a rule coming from the GSEs' formal underwriting guidelines for originators and servicers; and (ii) a rule implied by de facto underwriting norms in lender guidelines, which is also reflected in the guidelines of the Federal Housing Administration (FHA).

GSE rule. For loans to be GSE-eligibile, the minimum required coverage is the lesser of 100% of the replacement cost value or the unpaid loan balance, provided that the balance is no less than 80% of the replacement cost of the property. Specifically, the GSE required coverage is

$$C^{GSE} = min(RC, max(0.8RC, L)) \tag{1}$$

Equation 1 says that when $L < 0.8 \cdot RC$, the GSE rule requires households to obtain at least 80% of the RC so long as this level of coverage still exceeds the loan amount. When $0.8 \cdot RC < L < RC$, the required coverage equals the loan balance. Finally, when $L > RC$ the required coverage equals full replacement cost, ensuring that the structure can be rebuilt and the lender's collateral position is preserved. Panel A of Figure 1 visually depicts these GSE requirements (blue line).[14]

De facto rule. While the rules are made explicit in GSE guidelines, in practice mortgage lenders and servicers are known to follow a simpler de facto coverage rule: that coverage must at least equal the minimum of the outstanding loan balance. For example, several loan underwriting guidelines, including those of the FHA, state that coverage must meet or exceed the loan amount.[15] These

between \$500 and \$2000 dollars (Figure C.8). In case of catastrophic losses, the amount by which households are under-insured would matter a lot more than the deductible amount they would need to pay, even at the higher end of deductible levels.

[14] See the Fannie Mae's and Freddie Mac's guidelines for mortgage origination and servicing. Private lenders may similarly require full replacement cost coverage or establish minimum coverage thresholds.

[15] See, for example: Zebra Insurance Broker, "Lender Requirements for Homeowners Insurance." FHA also requires homeowners insurance coverage to "equal the lesser of the loan's unpaid balance or 100% of the insurable value [replacement cost]."

practices imply a de facto rule in which the required coverage can be defined as

$$C^{Defacto} = min(RC, L) \tag{2}$$

Equation 2 says that when $L < RC$, households must obtain at least the loan balance, and when $L > RC$ they must obtain coverage equal to full replacement cost. Panel A of Figure 1 depicts these requirements (red line). We therefore treat the loan balance as a natural benchmark for the level of coverage lenders seek to protect and, together with information on L/RC, use it to infer how tightly lender constraints bind in practice.[16] Figure 1 shows that the de facto rule and GSE rule only differ in the region where $L/RC < 0.8$, making it strictly looser than the GSE rule. Note that this de-facto rule ensures that the lender's loan exposure remains protected. Panel B of Figure 1 summarizes the expected minimum coverage in the de facto and the GSE rule.

2.1.1. Beneficiaries and Enforcement

There are two additional aspects worth noting. First, lenders and the GSEs are also beneficiaries on the insurance contract, giving them a direct claim to insurance proceeds. In the event of a disaster, when the insurance company issues a payout, the check is often made out jointly to both the homeowner and the servicer (on behalf of the lender).[17] The servicer typically controls how those insurance proceeds are disbursed, ensuring that the funds can be applied appropriately, e.g., to rebuilding and repairs or towards the loan balance.[18] As a result, a portion of the insurance proceeds acts as protection for lenders: if the borrower defaults after a disaster, the funds can be used either for restoring the property (preserving collateral value) or go directly to repaying the lender, thereby shielding the lender from potential losses.

Second, to simplify payments and enforce requirements, mortgage servicers collect and pay taxes and insurance premiums using escrow accounts.[19] The majority of mortgages feature escrow deposit accounts (Anderson and Dokko (2008)), and they are required by FHA, Fannie Mae, and Freddie Mac in most circumstances.[20] When a mortgage is originated, originators estimate the annual tax and insurance outlays associated with the mortgage, establish a monthly payment sufficient to meet those payments, and collect a reserve in the escrow account.[21] Because homeowners insurance is paid annually for the year ahead, the originator collects the entire homeowners insurance payment as part of the initial closing costs, or requires proof of insurance payment, making the insurance burden salient at origination.

[16] While these rules primarily apply to conforming loans (loans eligible to be sold to the GSEs), in practice, many lenders adopt parallel standards, so that even loans outside the conforming segment, must carry enough insurance to satisfy either the GSE style 80% of replacement cost condition or coverage up to the loan amount.

[17] See "*Understanding the insurance claims payment process,*" Insurance Information Institute (report here).

[18] For example, see I. Kuhn, "*Can a Mortgagee/Lender Withhold Property Insurance Proceeds?*", propertyinsurance-coveragelaw.com. Report can be found here.

[19] Mortgage escrow accounts are separate legal arrangements that are used to hold any partial mortgage payments and payments toward homeowners insurance premiums, and property taxes.

[20] The GSEs at times allow lenders to waive the requirement, but they retain the right to enforce the escrow requirement if the borrower fails to pay his or her property taxes.

[21] https://www.consumerfinance.gov/rules-policy/regulations/1024/17/

2.2. Testable implications of lender mandates for coverage

To understand how lenders influence coverage outcomes, we exploit the structure of the GSE and de facto rules. These rules generate several empirical predictions for coverage that are less consistent with how households would likely choose insurance absent lender mandates. We outline these predictions below.

2.2.1. Conceptual framework

To fix ideas, we start with a simple framework. Let $C^M(L, RC)$ denote the minimum coverage benchmark implied by the mortgage market mandates, where $M \in \{GSE, Defacto\}$. We allow the mortgage market mandate to serve as a benchmark for households choosing coverage. Households must satisfy the lender minimum, but may also purchase discretionary coverage above that minimum. We therefore write observed coverage as

$$C^{Actual} = C^M + E,$$

where $E \geq 0$ denotes discretionary excess coverage above the lender benchmark. This excess coverage may depend on insurance prices, household characteristics, property characteristics, and the remaining uninsured gap above the lender requirement.

A useful benchmark is the one in which replacement cost is the natural upper reference point for households, and households insure a fraction θ of the gap between the lender benchmark and replacement cost:

$$C^{Actual} = C^M + \theta(RC - C^M), \quad 0 \leq \theta \leq 1.$$

Equivalently,

$$C^{Actual} = (1 - \theta)C^M + \theta RC.$$

This representation nests several cases. When $\theta = 0$, the household chooses exactly the lender minimum and coverage bunches at C^M. When $0 < \theta < 1$, the household purchases some discretionary coverage above the lender minimum, but coverage remains indexed to the mortgage market benchmark. When $\theta = 1$, the household obtains full replacement cost coverage, and the lender minimum is not binding, playing no role in household choice. A household's θ can reflect its own beliefs, preferences, and characteristics. For example, under the frictionless model of insurance demand (Arrow, 1963, Mossin, 1968), a risk averse household would obtain full coverage when prices are actuarially fair. Under this frictionless null, $\theta = 1$. On the other hand, if households have low willingness to pay for insurance, for example because of financial constraints (Rampini and Viswanathan, 2013, Ericson and Sydnor, 2018) or beliefs (Bakkensen and Ma, 2020), we expect $\theta < 1$.

Our formulation clarifies how mortgage mandates can affect coverage even for households that

do not bunch exactly at the minimum. For example, under the de facto rule, when $L < RC$, the lender benchmark is the loan balance, $C^M = L$. In this region,

$$C^{Actual} = L + \theta(RC - L) = (1 - \theta)L + \theta RC.$$

Thus, the loan balance can influence coverage choices even when households buy coverage above the loan amount (i.e. the constraint does not bind). This is because the mandate may also serve as a reference point for households, shaping coverage even when it is slack; for instance, if households anchor on the required amount as a signal of "adequate" coverage. Bunching at $C = L$ captures only the special case $\theta = 0$, where coverage is fully pinned down by the lender minimum. More generally, any household with $\theta < 1$ has coverage that remains partly indexed to the loan balance. Bunching identifies the share of borrowers whose coverage is exactly equal to the lender minimum, while other empirical estimates targeting θ (such as the pass-through of loan amounts and estimates around the kinks in the rule, discussed below) capture the broader influence of the mortgage market requirement even when it does not bind.[22] This representation delivers several testable implications on both coverage levels, as well as how coverage varies with insurance prices that we take to the data.

An important point to note is that in standard models of Arrow (1963), Mossin (1968), households' demand is not determined by the loan amount. In other models that explicitly consider the role of household leverage, the loan amount enters negatively, as borrowers only have incentive to insure their equity (Mahoney (2015)).[23] Therefore, a positive relationship between loan amount and coverage likely reflects the influence of lenders seeking to protect their exposures.

2.2.2. Predictions on coverage levels

The first three predictions examine what determines the level of coverage, isolating variation that is most naturally explained by the lender mandate rather than households' latent demand. We test these three predictions in Section 3.

Prediction 1: Compliance with rules. First, if lenders enforce coverage rules at origination, we should observe that most households obtain enough coverage to meet or exceed requirements.[24] We can directly assess compliance with the two rules by examining how the coverage ratio (C/RC) and the coverage-to-loan-balance ratio (C/L) varies across the three L/RC regions shown in panels a and b of Figure 1. Under both rules, households with $L/RC > 1$ should have coverage equal to replacement costs, and households with $0.8 < L/RC < 1$ should have coverage exceed the loan amount. However, the two rules differ when $L/RC < 0.8$; there, the GSEs require a coverage ratio (C/RC) of 80%, whereas the de-facto rule requires a coverage-to-loan-balance (C/L) of 100%. We

[22] The standard max formulation, $C^{Actual} = \max\{C^{HH}, C^M\}$, is a special case in which the household either chooses its unconstrained latent demand or is pushed exactly to the lender minimum. Our more general formulation allows the lender benchmark to affect coverage even when households choose coverage above the minimum.

[23] For example, higher LTV borrowers would obtain less coverage, all-else-equal.

[24] If observed coverage were systematically below lender rules, that would indicate rampant non-compliance, and would be inconsistent with lender enforcement of coverage rules.

note that strict enforcement of the GSE rule would lead to a coverage ratio of at least 80% for all borrowers. Additionally, because the de-facto rule is a looser version of the GSE rule, we can distinguish between the two rules by evaluating outcomes in the region where they differ.

Prediction 2: Kinks in the pass-through to coverage. Second, because required coverage is a piecewise linear function of L and RC with kinks at known underwriting thresholds, the pass-through from these variables to coverage should change exactly at these points if the mandate binds. Conditional on replacement cost, the effect of the loan amount on coverage should be larger when $L/RC < 1$ than when $L/RC > 1$. Similarly, conditional on the loan amount, the effect of replacement cost differs below and above the thresholds. We can exploit these implied changes in slope to identify the causal impact of loan balances on coverage and to separate the lender-driven component of coverage from households' own demand. Because lender mandates still impact coverage even for households with $\theta > 0$, we can exploit the kink in underwriting rules to identify the effect of lender mandates more broadly.

Prediction 3: Bunching at lender mandates. Third, if households have a latent demand for coverage below the required level but adjust upward to satisfy the rules, the distribution of coverage should exhibit bunching at the mandated points, and also be censored with limited mass below lender mandates. Specifically, we would expect to see bunching of coverage at loan balance. The extent of bunching indicates the share of borrowers whose coverage is completely determined by the lender mandate. This prediction arises because there will be some households with $\theta = 0$ who obtain the bare minimum required by lenders.

2.2.3. Predictions on coverage elasticities

In the next three predictions, we will examine the price-sensitivity of coverage choices, to exploit another key divergence between lender mandates and households. Crucially, lender mandates C^M do not vary with insurance price; they are determined entirely by the loan amount and replacement cost. In contrast, we expect household latent demand to be price-sensitive, with downward-sloping demand curves (Einav et al., 2010). This divergence yields three additional predictions that can be evaluated in the data. We test these three predictions in Section 4.

Prediction 4: Overall coverage inelastic. Fourth, if mortgage market rules bind, coverage choices do not vary with price, meaning that the component of insurance demand coming from mortgage markets is insensitive to insurance prices. In contrast, we expect latent household demand to be price sensitive. Therefore, if lender mandates are a quantitatively important driver of coverage, then these requirements would render overall coverage demand to be inelastic.

Prediction 5: Excess coverage elastic. For lender mandates to have an impact on coverage choices, it must also be the case that households have low latent demand for insurance. If not, households would obtain full coverage on their own and lender mandates would not matter. We therefore also expect to see that the part of coverage demand that is in excess of the lender requirements ("excess coverage") is elastic, again reflecting that households' latent demand is

price-sensitive.

Prediction 6: High bunching elasticity. As prices exogenously rise, household latent demand for insurance decreases, making lender mandates even more binding for some households close to the threshold. For example, a rise in insurance prices would result in an increase in bunching at the loan balance, indicating that a growing share of coverage is determined by mortgage requirements rather than household demand.

3. Testing Mortgage Mandates: Coverage Levels

In this section, we examine how lender mandates influence household coverage choices. We consider three tests, as described in Section 2.2.2: the overall compliance with lender rules; kinks in the pass-through to coverage implied by lender rules; and bunching at lender mandates.

3.1. Prediction 1: Compliance with mortgage market rules

We first examine whether households obtain the coverage required by the mortgage market rules, and how compliance varies across the distribution of L/RC where the mandates change. Coverage requirements depend on the loan balance and replacement cost; therefore, we can assess compliance by looking at how the coverage ratio (C/RC) and the coverage-to-loan-balance ratio (C/L) vary across the three different L/RC regions depicted in panel (b) of Figure 1 and described in Section 2.

Panel (a) of Figure 2 shows evidence of partial household compliance with the GSE rule. When $L/RC < 0.8$, the average coverage ratio is approximately 70%, below the 80% required by the GSEs. However, in the other regions, where the GSE and the de facto rules align, the actual coverage ratio complies with the mortgage requirements on average. Although households with L/RC ratio below 0.8 appear out of compliance with the GSE rule, we find near-perfect compliance with the de facto rule. Panel (b) of Figure 2 shows when $L/RC < 0.8$, on average household coverage still exceeds the loan balance, consistent with the de facto rule. For the middle region, where the GSE rule also requires coverage to equal the loan amount, households obtain exactly this level on average. Taken together, we conclude that, on average, there is near-full compliance with the mortgage requirements from the de facto rule.

The high compliance estimates are suggestive that lender mandates are relevant for household coverage choices. To assess their quantitative importance, we conduct an R-squared decomposition to measure the importance of mortgage requirements. Figure 3 shows that the de facto rule alone explains 67% of the variation in coverage. Adding borrower and property characteristics raises the explained variation to over 80%. Furthermore, Table C.7 confirms that the L/RC underwriting variable is strongly predictive of coverage ratios, and the coefficient remains stable as we add property controls, ZIP-by-year fixed effects, insurer-by-year fixed effects, and borrower risk proxies. The stability of the coefficient suggests that the mortgage requirement channel is not merely capturing other observed property quality or borrower risk.[25]

[25] Figure C.10 plots actual coverage against the coverage predicted by the de facto rule, showing that actual coverage

Taken together, these results motivate our focus on mortgage mandates as a quantitatively important driver of household coverage choices. We note that while the compliance with the mortgage rules is a necessary condition for mandates to matter, it is not sufficient. Households may satisfy the required coverage levels simply because their latent demand for coverage already exceeds the lender minimum. Thus, compliance alone cannot distinguish coverage induced by mortgage mandates from coverage households would have chosen absent those mandates. The next two predictions provide sharper tests of whether mortgage mandates actively shape coverage choices: if the mandates bind, coverage should exhibit kinks at the rule-implied thresholds and bunching at mandated coverage levels.

3.2. Prediction 2: Kink in the pass-through to coverage

We next ask whether there are kinks in the pass-through of loan amount and replacement cost to coverage at the underwriting thresholds implied by the mortgage rules? Figure 1 implies that coverage should be sensitive to both loan amount and replacement costs, and that these sensitivities would vary at the L/RC underwriting thresholds. Figure 4 shows the predictions for each coefficient for the two rules under the assumption that coverage is determined entirely by mortgage mandates, and household demand plays no role.

GSE rule. Under the GSE rule, coverage should respond one-for-one to loan amount when L/RC lies between 80% and 100%. Outside this range, coverage should be completely insensitive to loan amount. The pass through of replacement cost, should be 0.8 when $L/RC < 0.8$; 0 when $0.8 < L/RC < 1$; and 1 when $L/RC >= 1$. This implies two kinks: at $L/RC = 0.8$, where the rule transitions from requiring $C = 0.8RC$ to $C = L$; and at $L/RC = 1$, where the rule transitions from requiring $C = L$ to $C = RC$. Correspondingly, the coefficients on loan amount should jump from 0 to 1 at $L/RC = 0.8$ and drop back to 0 at $L/RC = 1$. The coefficients on RC should drop from 0.8 to 0 at $L/RC = 0.8$ and jump back to 1 at $L/RC = 1$.

De facto rule. The de facto rule requires coverage equal to the loan amount when $L/RC < 1$, and coverage equal to RC when $L/RC > 1$. Sensitivity of coverage to loan amount and RC changes precisely at $L/RC = 1$, and nowhere else. Thus, the rule predicts one kink at $L/RC = 1$, where the coefficient on loan amount drops from 1 to 0, and the coefficient on RC jumps from 0 to 1. We note that these pass-throughs occur when coverage is determined entirely by mortgage mandates. More generally, when households purchase discretionary coverage, the coefficients need not be exactly zero or one, but they should still exhibit kinks at the thresholds where the mandate changes. For example, when $L < RC$, coverage can be written as $C = (1 - \theta)L + \theta RC$ (see Section 2), which implies a pass-through of $1 - \theta$ for loan amount, rather than one, and θ for RC, rather than zero.

3.2.1. Evidence from the full distribution of L/RC

We begin by estimating a flexible specification that allows the pass-through of loan amount and replacement cost to vary across the full distribution of L/RC. This provides a non-parametric

choices closely track this requirement.

diagnostic for where the relationship between coverage and the two mortgage market variables changes. Specifically, we estimate:

$$\text{C}_{i,z,t} = \sum_j \beta_j \cdot (\text{LTRC Bin}^j_{i,z,t} \times Loan_{i,z,t}) + \sum_j \gamma_j \cdot (\text{LTRC Bin}^j_{i,z,t} \times RC_{i,z,t}) + \gamma' W_{i,z,t} + \varepsilon_{i,z,t} \quad (3)$$

where $C_{i,z,t}$ is the household's coverage choice, $L_{i,z,t}$ is the loan amount, and $RC_{i,z,t}$ is replacement cost. LTRC Bin^j is a dummy variable for 10 percentage point bins of the loan-to-replacement-cost ratio ($L/RC \in \{.4, .5\}, \{.5, .6\}, ...\{1.1, 1.2\}, \{1.2+\}$). $W_{i,z,t}$ comprises borrower-year fixed effects, structure controls (including house price per square foot), zip-year fixed effects, and insurer-year fixed effects. The subscripts denote loan i in zip z originated in year t. Each β_j captures the pass-through of loan balance to coverage for LTRC bin^j, and each γ_j captures the pass-through of replacement cost to coverage for the same bin.

This specification in Equation 3 allows us to test the pass-through coefficients implied by both rules simultaneously. Under strict enforcement of the GSE rule, the loan coefficient should jump up at $L/RC = 0.8$ and fall at $L/RC = 1$, while the replacement-cost coefficient should move in the opposite direction at the same two thresholds. Under the de facto rule, by contrast, there should be only one kink, at $L/RC = 1$: below this threshold, coverage should be more sensitive to the loan amount, and above it, coverage should be more sensitive to replacement cost. Equation 3 therefore allows us to test whether the data exhibit the kink pattern implied by either rule without imposing the location of the kink in advance. Under the identifying assumptions detailed in Section 3.2.3, we can interpret the kinks at these underwriting thresholds as causal evidence that mortgage requirements influence household coverage choices.

Results. Figure 5 plots the estimated β_j and γ_j coefficients. Panel A shows that coverage is strongly related to loan amount when $L/RC < 1$: across the bins below this threshold, a \$1 increase in loan amount is associated with roughly a 60¢ increase in coverage. At L/RC=1, this relationship shifts sharply. For loans above the threshold, the pass-through of loan amount falls to roughly 40¢. Panel B shows the mirror-image pattern for replacement cost: the replacement cost coefficient rises at $L/RC = 1$, exactly where the loan amount coefficient falls. Thus, coverage becomes less sensitive to loan amount and more sensitive to replacement cost precisely when the de facto rule switches from $C = L$ to $C = RC$. Importantly, we find no comparable kink at $L/RC = 0.8$. This pattern is inconsistent with strict enforcement of the full GSE schedule and aligns with the compliance evidence: borrowers with $L/RC < 0.8$ often fall below the GSE required coverage ratio of 80%, but generally maintain coverage above the loan balance as implied by the de facto rule. This suggests that the loan amount is an important benchmark for household coverage choices.

The magnitudes also show that the lender rule is important but not a complete determinant of coverage. Under strict adherence to the de facto rule, the loan amount coefficient β_j would be 1 below $L/RC = 1$ and 0 above it, while the replacement cost coefficient γ_j would be 0 below $L/RC = 1$ and 1 above it. Instead, the estimates are more muted. We interpret this pattern to be consistent with the framework in Section 2, where observed coverage reflects both household discretionary

demand and mortgage mandates. The estimated coefficients of roughly 0.6 on loan amount and 0.4 on replacement cost are therefore consistent with households covering the loan balance and then insuring part, roughly 40%, but not all, of the remaining gap to replacement cost. In other words, households do not appear to choose coverage solely according to the lender rule, but their choices are strongly influenced by it.

There are two possible interpretations of the coefficient on loan amount ($\beta_j = 0.6$) for borrowers with $L/RC < 1$. The first stronger interpretation is that it reflects mortgage market driven causal pass-through of loan amount to coverage. Several patterns support this view: (i) The kink occurs exactly at $L/RC = 1$, where the de facto rule switches from $C = L$ to $C = RC$, and the loan and RC coefficients move in opposite directions at this threshold, as predicted by the rule. (ii) The coefficients are highly stable within each side of the threshold. (iii) The estimates are robust to detailed borrower, property, insurer-year, and house-price controls. This stability is difficult to reconcile with explanations where the coefficients are driven by smoothly varying household preferences or borrower sophistication across the L/RC distribution.[26] The positive loan amount coefficient above $L/RC = 1$ can also be interpreted in this light: although the de facto rule implies that loan amount should no longer matter once $L > RC$, the remaining pass-through of about 0.4 may reflect some lenders or servicers continuing to use the loan balance as a salient benchmark. Under this interpretation, both the 0.6 coefficient below the threshold and the 0.4 coefficient above it reflect lender-related indexing to loan amount, but this indexing is stronger below $L/RC = 1$, where the de facto rule directly ties coverage to the loan balance.

A second, more conservative interpretation is that this coefficient is not fully causal. Loan amounts are jointly determined with house values, borrower characteristics, and other determinants of insurance demand. A natural concern is house value: loan amounts are mechanically related to house prices through the LTV ratio, and households may use house value as a reference point when choosing coverage. Although our controls absorb much of this variation, including ZIP-year fixed effects and house price per square foot, loan amount may still proxy for unobserved features of the borrower or property. Under this conservative interpretation, the coefficient below $L/RC = 1$ combines both the causal effect of the lender rule and the endogenous relationship between loan amount and coverage, and the coefficient above $L/RC = 1$ captures the pure endogenous component.

This conservative interpretation motivates a kink design, where we seek to estimate the discontinuous change in the loan amount coefficient at the threshold. If the confounding relationship between loan amount and coverage varies smoothly through $L/RC = 1$, then the change in the loan amount pass-through at the threshold differences out this endogenous component. In this case, the drop in the coefficient from below the threshold to above the threshold isolates the mortgage mandate component. Thus, even under the conservative interpretation that the level coefficients are partly endogenous, the kink at $L/RC = 1$ provides causal evidence that mortgage market mandates

[26] For example, if low L/RC borrowers had stronger unobserved preferences for full RC than high L/RC borrowers, we would expect the RC coefficient to vary systematically as L/RC changes. Instead, the main change occurs discretely at the threshold predicted by the de facto lender rule.

affect how households choose coverage. The next section formalizes this argument using a narrow bandwidth kink design around $L/RC = 1$.

3.2.2. Pass-through around narrow bandwidths of $L/RC = 1$

A concern with the estimates from the full distribution is that borrowers with different beliefs or preferences may systematically sort into high or low L/RC regimes. While the stability of the coefficients within each L/RC group suggests this is less likely, there may still be concerns about unobserved selection across the distribution of L/RC. To address this concern, we zoom in narrowly around the $L/RC = 1$ threshold and test whether the sensitivity changes precisely at this threshold. We consider households just above and below the threshold, that are arguably similar on observable and unobservable dimensions, to isolate variation from the mortgage mandate. We focus on the de facto rule because of our earlier results that there are no observed kinks at $L/RC = 0.8$, and weak evidence that the GSE rule is being followed strictly. We consider the following specification:

$$C_{i,z,t} = \beta_1 Loan_{i,z,t} + \beta_2 Loan_{i,z,t} \times \mathbb{I}\Big(LTRC_{i,z,t} > 1\Big) + \gamma_1 RC_{i,z,t} + \gamma_2 RC_{i,z,t} \times \mathbb{I}\Big(LTRC_{i,z,t} > 1\Big) + \eta W_{i,z,t} + \varepsilon_{i,z,t} \tag{4}$$

where $C_{i,z,t}$ is the coverage amount, $Loan_{i,z,t}$ is the loan amount, $RC_{i,z,t}$ is the replacement cost, and $\mathbb{I}\Big(LTRC_{i,z,t} > 1\Big)$ is a dummy indicating that the loan-to-replacement cost ratio exceeds 1. $W_{i,z,t}$ includes borrower-year fixed effects, structure controls (including house price per square foot), zip-year fixed effects, and insurer-year fixed effects, as before. We limit the sample to loans within 5% of the $L/RC = 1$ threshold. In robustness checks, we vary our use of controls and consider different bandwidths.

Under the identifying assumptions described in Section 3.2.3, the coefficient β_2 measures the discontinuous change in the slope of coverage on loan amount precisely at the threshold, identifying the causal effect of the mortgage mandate, i.e. how loan amount impacts coverage. Perfect adherence to the mortgage mandate (de facto rule) would imply that $\beta_1 = 1$, $\beta_2 = $ -1, $\gamma_1 = 0$, and $\gamma_2 = 1$, with both β_2 and γ_2 statistically significant. More generally, if the mortgage mandate matters but does not fully determine coverage, we expect $\beta_2 < 0$ and $\gamma_2 > 0$, i.e. coverage should become less sensitive to loan amount and more sensitive to replacement cost once L/RC exceeds one.

Results. Table 2 presents the results. In the baseline specification with ZIP-year fixed effects (column 1), the pass-through from loan amount to coverage is approximately 0.72 when $L/RC < 1$, meaning that a \$1 increase in loan amount is associated with a 70¢ increase in coverage. This falls by 0.28¢ for loans with $L/RC > 1$, implying that for this group the pass-through of loan amount to coverage is roughly 40¢. The replacement-cost pass-through moves in the opposite direction; it is is small (less than 0.30) when $L/RC < 1$ but increases to 0.57 when $L/RC > 1$. The estimates match the full distribution results in the previous section, and remain stable across all specifications, indicating robustness to borrower, property, and insurer controls and fixed effects. In our most

conservative specification in column 4, we obtain that going from above the kink (where loan amount does not bind) to below the kink (where loan amount binds) leads to a significant change in the sensitivity of coverage to loan amount. We obtain that a $1 exogenous increase in loan amount leads to a 23¢ increase in coverage.[27]

3.2.3. Identifying assumptions and robustness

The fundamental identifying assumption of the kink design is that any confounding unobservables are continuous at the $L/RC = 1$ threshold within narrow bandwidths. In our setting, there could be three potential empirical concerns.

Identifying Assumption: Balanced characteristics and continuity of unobservables. The first identifying assumption for our empirical design is that unobserved borrower, lender, or insurer characteristics vary continuously through the $L/RC = 1$ threshold. This assumption is critical: if borrowers with systematically different preferences for coverage sorted sharply into the high-L/RC versus low-L/RC regions, the change in coefficients around the kink would reflect borrower selection rather than the causal effect of the change in the mandate. To validate this assumption, we conduct the following tests.

First, we examine whether borrowers bunch or discretely sort at the $L/RC = 1$ threshold. Figure C.4 plots a histogram of the L/RC ratio and shows no evidence of mass points or bunching at L/RC =1. This suggests that borrowers are not strategically clustering just below or above the threshold, which would indicate awareness of the rule and intentional sorting to obtain a different requirement. Second, we directly evaluate covariate balance on either side of the threshold. Table C.3 compares borrower and property characteristics for loans just below versus just above $L/RC = 1$, within 5pp of the threshold. Across observables, including loan amount, replacement cost, home price, dollar per square foot, borrower credit score, and other property characteristics, we find that borrowers are similar on observables on either side of the threshold. This balance on observables provides reassurance that the populations on either side of the threshold are similar on observables.

Third, Table C.4 tests whether these characteristics themselves exhibit discontinuous changes in their relationship with loan amount or replacement cost at the threshold of $L/RC = 1$ (smoothness test). We find no significant smoothness violations, consistent with the characteristics varying continuously across the boundary. Fourth, Table 3 demonstrates that our main coefficient estimates are robust across alternative bandwidths (5%, 10%, 15%). Crossing $L/RC = 1$ reduces the sensitivity of coverage to loan amount and increases the sensitivity of coverage to replacement cost. Relatedly, table C.6 also show that the kink results are robust to removing households who directly bunch at the loan amount. Furthermore, Table 2 shows that the coefficients are stable as we add or remove controls for borrower, property, and insurer characteristics. These results suggest that our estimates are not driven by the particular bandwidth choice or confounding from observables. Lastly,

[27] As discussed earlier, if we assume that the effect of the unobserved confounder fully explains the pass-through above the kink, and the effect of this confounder is smooth through the kink, then subtracting the above kink coefficient (0.40), from the below kink coefficient (0.70) identifies the pure causal effect of loan amount.

as discussed earlier, Figure 5 shows that kinks do not appear at other thresholds in the L/RC distribution, further supporting the conclusion that the kink at $L/RC = 1$ reflects the mortgage mandate, not spurious patterns in the data.

Concern 1: Replacement Cost Measurement Error. A potential concern is that replacement cost is not directly observed and may be measured with error. This concern is especially relevant in our setting because replacement cost enters both the mortgage rule and the running variable, L/RC. The key question is whether measurement error in our RC measure could generate the kink patterns we document. We distinguish between two cases. If measurement error is systematic and correlated with the underwriting threshold, it could in principle bias our estimates. In contrast, if measurement error is classical, it should primarily add noise, attenuate the estimated pass-through coefficients, and blur the assignment of borrowers around the $L/RC = 1$ threshold. In that case, the estimated kink would likely be biased toward zero.

Several pieces of evidence suggest that RC measurement error is unlikely to explain our findings. First, a non-classical measurement-error explanation would need to be quite specific. It would have to generate a sharp change in the pass-through of loan amount to coverage exactly at $L/RC = 1$, while also generating an offsetting change in the pass-through of replacement cost at the same point. Moreover, it would have to do so without producing comparable kinks elsewhere in the L/RC distribution. In other words, a systematic measurement error explanation would need to mimic the exact institutional prediction of the mortgage mandate, which is unlikely.

Second, we validate the accuracy of our RC measure by comparing it several to independent benchmarks. Our RC estimates match the average levels reported by insurers and builders (e.g., Verisk, NAHB, R.S. Means), and the growth rates of our RC measure also align with external cost indices (e.g., NAHB, BEA). This matching on both levels and trends suggests we do not have systematic biases in our measurement on average.

Third, the estimated kink is stable when we add detailed structure controls, including square footage, year built, property type, and house price per square foot. This matters because many plausible forms of systematic RC error would be correlated with observable structure characteristics. For example, if our RC measure systematically under- or overestimated rebuilding costs for larger homes, older homes, or particular property types, then adding these controls should materially change the kink estimates. Instead, the coefficient patterns remain stable across specifications with and without these controls. This stability is more consistent with measurement error that adds noise than with systematic error that drives the kink.

Fourth, a further internal validation comes from the high L/RC region. When $L/RC > 1$, the mortgage rules imply that RC is the relevant coverage benchmark. In the data, households in this region obtain coverage close to, and on average above, our measure of RC. Thus, the RC measure appears to align with the underwriting benchmark precisely where replacement cost should bind. For measurement error to explain the kink, it would therefore need to be highly asymmetric: our RC measure would have to be accurate in the high L/RC region, but systematically mis-measured

in the low L/RC region in a way that makes coverage appear to track loan amount and creates a discontinuity exactly at $L/RC = 1$. Taken together, this evidence makes the measurement error alternative much less plausible.

A remaining concern is that our RC measure is also used to construct the running variable, L/RC, and therefore determines whether a borrower is classified as below or above the $L/RC = 1$ threshold. If replacement cost is measured with classical error, some borrowers will be assigned to the wrong side of the threshold: loans with true $L/RC < 1$ may be classified as $L/RC > 1$, and vice versa. This misclassification mixes observations from the two regions and mechanically smooths the change in slopes at the threshold. As a result, classical measurement error in RC should attenuate the estimated kink and would lead us to *underestimate* the true causal effect, rather than generate a kink when it is not present. We confirm this intuition in simulations. When coverage is generated by household demand and not by the mandate, adding classical measurement error to replacement cost does not produce a spurious kink at $L/RC = 1$. When coverage is partly determined by the mandate, the same measurement error reduces the estimated kink by blurring the threshold assignment. Thus, under classical measurement error, our kink estimates should be interpreted as conservative.

Concern 2: House price effect. We interpret the kink at $L/RC = 1$ as reflecting the mortgage mandate. House prices are a natural threat to this interpretation because house prices are correlated with both loan amounts and replacement costs, and households may use market value as a reference point when choosing coverage. We conduct a series of tests to provide evidence that house prices are not driving our results.

First, our baseline specifications include ZIP-year fixed effects and price per square foot, which absorb a substantial amount of local house-price variation. Second, Figure C.5 shows that the distribution of coverage relative to market value at origination is relatively symmetric and there is little evidence of censoring below market value. Indeed, nearly half of households in our sample obtain coverage below market value. This stands in contrast to the coverage-to-loan-balance distribution, where coverage below the loan balance is rare. Furthermore, there is minimal evidence of bunching at market value. If households used market value as a reference point when choosing coverage, we would expect to see a substantial fraction of households directly obtaining coverage equal to market value at origination. The absence of bunching at market value makes it unlikely that borrowers are sorting on house price in a way that would create spurious discontinuities at $L/RC = 1$. Relatedly, Table C.3 shows that average house prices above and below the L/RC thresholds are similar.

Third, Table C.5 reports two placebo tests that examine whether the loan amount kink is capturing households' responses to market value instead. In the first test, we ask whether the pass-through of house price to coverage changes discontinuously at the threshold, $L/RC = 1$. This is a natural concern because loan amount and house price are correlated. If borrowers sorted around the threshold for reasons related to market value, or if coverage choices responded discontinuously to market value at $L/RC = 1$, then the estimated loan amount kink could partly reflect this house price channel rather than the mortgage mandate. Column 1 shows no evidence of such a discontinuity: the

interaction term of house prices and the dummy for the L/RC threshold is small and statistically insignificant. This reduces the concern that the kink in loan amount pass-through is driven by market value anchoring or strategic sorting around the threshold.

The second placebo test more directly examines whether households anchor coverage choices to market value (MV) rather than to the loan amount. If households used market value as a benchmark, subject to the requirement that coverage need not exceed replacement cost, coverage would follow a rule analogous to $min(MV, RC)$. In that case, coverage should be sensitive to house price when $MV < RC$, but should become less sensitive to house price and more sensitive to replacement cost once $MV > RC$. This would generate a kink in the pass-through of house price to coverage at $MV/RC = 1$. Column 2 tests this prediction by replacing loan amount with house price and estimating whether the house price and replacement cost pass-throughs change at the $MV/RC = 1$ threshold. We find no such kink: the interaction terms are small and statistically insignificant. This is important because house price is closely related to loan amount through the loan-to-value ratio. The absence of a kink at $MV/RC = 1$ therefore suggests that the kink we estimate at $L/RC = 1$ is not simply capturing household anchoring to market value.

These results strengthen our interpretation that the observed discontinuity in the pass-through of loan amount to coverage is specific to loan amount and generated by the mortgage mandate, rather than a spurious pattern arising from borrowers' responses to home values.

To understand the economic magnitudes of the mortgage mandates, we use our framework to provide a simple back-of-the-envelope calculation, which translates the pass-through estimates into level magnitudes. The mandate-induced component of coverage is $(1 - \theta)L$, so the implied share of observed coverage attributable to the loan-amount floor is approximately $(1 - \theta)L/C$. Using the below-threshold pass-through estimate of 0.6, this implies that $0.6 \times L/C = 46\%$ of observed coverage is tied to the loan-amount floor. The regression-kink estimate is more conservative because it subtracts any loan-related pass-through that is smooth through the $L/RC = 1$ threshold, thereby identifying only the incremental mandate-induced pass-through below the threshold. This yields a lower-bound estimate of roughly $0.23 \times L/C = 18\%$ for the average borrower. Taken together, these estimates imply that lender mandates account for roughly 18–46% of observed coverage.

3.3. Prediction 3: Bunching at lender mandates

The third prediction is that if lender minimums bind for some households because latent demand for coverage is low, the distribution of coverage should exhibit bunching at the mandated coverage level. We test this prediction in the region where the GSE and de facto rules differ most sharply: borrowers with $L/RC < 0.8$. Under the GSE rule, these households would be required to obtain coverage equal to at least 80% of RC. Under the de facto rule, however, the relevant minimum is the loan balance. Thus, if the de facto rule is the binding constraint in this region, the distribution of C/L should exhibit a hard lower bound around one and excess mass at $C = L$.

To evaluate this directly, we plot a histogram of the ratio of coverage to loan balance for borrowers

with $L/RC < 0.8$. Figure 6 reveals a striking pattern: the vast majority (90%) of borrowers hold coverage above their loan balance, with significant bunching directly at the loan amount (roughly 9.5% of households are at the loan amount threshold). Such clustering at the loan amount is inconsistent with smooth house preferences alone. Instead, it is what we would expect if the loan amount acts as a hard underwriting constraint. Consistent with the framework in Section 2, bunching at $C = L$ captures the case where borrowers' coverage is fully pinned down by the mandate. Thus, the roughly 9.5% mass at $C = L$ should be interpreted as the share of borrowers for whom $\theta = 0$. Notably, fewer than 10% of households obtain coverage below their loan amount (i.e. $C/L < 1$), despite this being theoretically possible if lenders permitted it.[28]

This bunching evidence complements the compliance and kink results. Although low L/RC borrowers often fall below the GSE requirement of 80% of replacement cost, they rarely fall below the loan balance and are disproportionately clustered at the loan balance. Thus, the binding threshold in this region appears to be the de facto loan-balance rule rather than the formal GSE rule. The implication is that lender exposures are largely protected even when households remain underinsured relative to replacement cost. In this sense, the protection gap is concentrated in household equity: once coverage protects the lender's claim, any remaining gap between coverage and replacement cost is borne by the household.

4. Testing Mortgage Mandates: Coverage Elasticities

In this section, we examine how mortgage mandates influence coverage demand elasticities. As described in Section 2.2.3, mortgage mandates do not vary with insurance prices or risk, whereas under the standard model household's intrinsic demand does (e.g., Arrow (1963), Mossin (1968), Einav and Finkelstein (2023)). This divergence generates three testable predictions for demand elasticities under the assumption that mandates matter for household choice: the overall coverage demand is inelastic; excess coverage above lender mandates is elastic; and households are more likely to bunch at lender mandates when prices exogenously rise. All three elasticities are necessary to understand the sensitivity of coverage demand to prices; how coverage would evolve in a world of rising prices, and to what extent exposures will be borne by households versus lenders.

4.1. Empirical Design

To estimate demand elasticities, we must address the classic simultaneity bias, whereby the observed price-quantity relationships constitute equilibrium outcomes that reflect both insurers' cost structures (supply) and household and lender demand. To isolate demand elasticities, we require supply shocks that shift the supply curve independently of demand conditions.

We obtain arguably quasi-exogenous changes by exploiting the fact that in some "high friction"

[28]This histogram is consistent with how underwriting constraints operate elsewhere in mortgage markets. Debt-to-income (DTI) ratios, a central underwriting threshold, show analogous patterns: lenders enforce hard cutoffs (e.g., DTI, LTV, and credit score), and borrowers cluster sharply at these boundaries with very few breaching them. These lender thresholds typically appear in the data as bunching along with a censored distribution at hard cutoffs.

U.S. states, insurance prices are heavily regulated, and that insurers cross-subsidize losses across states (Oh et al. (2021)). While there is still variation in prices in high friction states, the timing of price changes is driven by when regulators conclude rate review proceedings and allow insurers to adjust premiums. Furthermore, approved price changes often differ substantially from what is implied by insurers' actuarial assessments. If regulators' price changes in high friction states do not depend on demand conditions in other states, which is likely given the strict nature of rate regulation, then insurance prices in high price regulation states can serve as a valid instrument for insurance prices in low price regulation ("low friction") states. This instrument builds off the intuition of the Hausman instruments that are commonly used in the industrial organization literature (e.g., Nevo (2001), Hausman (1996)), where the prices in a given market are instrumented with the prices in other geographic market segments.[29] However, because our instrument only uses prices in high friction states, we rely on different economic mechanisms. We first explain the instrument, and then discuss the key identifying assumptions below.

We construct the instrument as follows. First, we limit the sample to a subset of high friction states, as defined in Oh et al. (2021).[30] We compute the average insurance price for these states for each month and for each of the four main credit score categories.[31] We call this variable $\bar{P}_{HF,c,m,t}$ where "HF" indicates that the average is based on the high friction states, and c, m, t shows that the variable varies by credit score category c, month m, and year t. We then use this instrument to estimate the effect of premiums shocks in *low friction states*, which we refer to as $P_{i,c,z,m,t}$.

We consider the following two-stage least squares regression:

$$\text{First Stage: } \ln P_{i,c,z,m,t} = \omega + \eta \ln \bar{P}_{HF,c,m,t} + \theta' W_{i,c,z,m,t} + \nu_{i,c,z,m,t} \tag{5}$$

$$\text{Second Stage: } \ln \text{coverage}_{i,c,z,m,t} = \alpha + \beta \ln \widehat{P_{i,c,z,m,t}} + \gamma' W_{i,c,z,m,t} + \varepsilon_{i,c,z,m,t} \tag{6}$$

where *coverage* refers to the insurance coverage chosen and P refers to insurance prices (premiums per dollar of coverage). The subscript i indexes loans, and z indexes ZIP code. W refers to a vector of control variables, which include ZIP $\times$ credit score $\times$ year fixed effects as well as controls for climate risk exposure, replacement cost,[32], building age, borrower debt-to-income ratios, borrower age, and insurer $\times$ year fixed effects. Our measure of climate risk is a dummy variable for whether the property is in the top decile of expected annual climate losses from Cotality.

With ZIP $\times$ credit score $\times$ year fixed effects, we identify our coefficient off of the month-to-month variation in a given year for prices in high friction states. Our estimation compares two borrowers originating their mortgage in the same ZIP of a low friction state, having similar credit scores and

[29] They posit that suppliers, which operate nationally, may need to adjust prices across many markets in response to national supply shocks, and that prices in other markets are exogenous to demand conditions in the current market.

[30] In this paper, they classify states into three groups: high, medium, and low friction. In our context, we combine medium and low friction states together.

[31] The literature has shown that insurance prices vary significantly by credit score groups, separately from location (Blonz et al., 2024).

[32] One may worry about the effect of classical measurement error in replacement cost on our estimates. If our instrument is orthogonal to our estimated replacement cost in low friction states, then including replacement cost as a control will not bias the estimate. This can be verified by showing that the coefficient does not change when excluding this control, which we consider in Table C.12

property-level climate risk. However, they close their mortgages at two different months of the same year, and as a result they face different prices because of price changes in high friction states. Under our identifying assumptions, the coefficient β can be interpreted as the elasticity of demand for coverage with respect to insurance prices. We expect an estimate of $\beta \leq 0$, and an estimate where $\beta \leq -1$ would suggest that demand is price elastic.

4.1.1. Identifying Assumptions

For this instrument to be valid, we require that (i) prices in high friction states are a relevant instrument for prices in low friction states; (ii) prices in high friction states have a monotonic relationship with prices in low friction states; and (iii) prices in high friction states satisfy the exclusion restriction. We consider each of these assumptions in turn below.

Relevance and Monotonicity. Our IV requires that insurer's prices in high friction states are relevant for prices in low friction states. At first glance, this assumption may be surprising, because one may expect insurers to optimize on a state-by-state basis and simply set prices based on local expected losses and demand conditions. However, deviations from state-by-state optimization occur because insurers maximize profits in each state subject to a common across-state capital constraint. Because expected losses are at a state level, but capital buffers and financing costs are at the insurer level, the insurer's optimization problem is no longer separable on a state-by-state basis, creating cross-state linkages.

The insurance literature shows that costly external financing constraints leads insurers to add a surcharge above expected losses even in perfectly competitive markets, and that these financing costs are themselves a function of the insurer's capital buffer (Froot and O'Connell, 1999, Koijen and Yogo, 2016). This link is especially important if insurers expect to operate at a loss in high-friction states (Oh et al., 2021): sustained losses erode insurer capital, which raises their financing costs, leading insurers to charge more elsewhere to maintain regulatory capital buffers.[33] These persistent losses can be seen in the data: Figure C.11b shows that combined ratios in high friction states are usually equal to 1, meaning insurers on average pay out in claims and expenses what they earn in premiums.[34] High combined ratios are consistent with stringent price regulation in these states that prevents insurers from charging fully actuarially fair rates. In this way, persistent losses in high-friction states raise an insurer's marginal cost of operating *everywhere* through higher financing cost. Because financing costs are a function of prices, this incentivizes changing prices even in low friction states, creating a cross-subsidy between states (Oh et al., 2021).[35] As a concrete example,

[33] Regulatory capital constraints do bind for insurers, which suggests that on the margin they would accept marginally higher financing costs and lower capital buffers.

[34] In some years and states, the ratio can be much larger than one, meaning that insurers are operating at a loss.

[35] To see this, note that standard models of insurance supply (Koijen and Yogo (2015)) would imply the following equation for optimal prices in a low friction state s: $P_s = \kappa_s \cdot E[L_s] \cdot \Phi(P_s, P_{s'}, P_{s''}, ...)$, where κ_s is markup, $E[L_s]$ is expected losses, and Φ is shadow cost of capital. A change in price in a high friction state s' affects the insurer's financing constraints, changing Φ, which then changes its optimal price in the low friction state. Consistent with the financing constraints channel, Oh et al. (2021) show that pricing cross-state spillovers are pronounced for insurers that are financially more constrained.

if the price an insurer charges in high-friction California is not high enough to offset its expected losses, that will erode its capital and increase its financing costs, which in turn affects the price it can charge in low-friction Virginia.[36] Consistent with the relevance assumption, we observe a strong first stage (Table C.11).

Monotonicity. We also assume that the first stage relationship between prices in high friction states and low friction states is consistently monotonic. We observe a positive-signed first stage coefficient, which we verify holds consistently in various sub-samples, indicating that price increases in high friction states are consistently associated with price increases in low friction states.

First Stage. The positive sign of our first stage reflects the insurance cross-subsidy discussed earlier, and two key features of insurance markets. First, insurers' expected losses are increasing over time. Second, insurers continue to expect persistent losses in high friction states despite regulators allowing price increases, because these price increases are not high enough to offset rising expected losses. There is significant anecdotal evidence in support of these features.[37] First, public insurer statements and official climate forecasts indicate that insured losses have risen over time, and that they expect losses to continue to rise further.[38] Second, Figure C.11b shows that insurers remain unprofitable in high friction states, despite having received price increases over this period. Third, Figure C.12 shows that the rate increases granted by regulators are lower than what insurers would need to fully offset their expected losses, which they report in their regulatory rate filings.[39] On average, when insurers receive 10%, their models imply they require a 40% rate increase to be actuarially fair. In contrast, when insurers receive 2%, their model implies they require a 10% rate increase. This shows that price increases in high friction states are actually associated with a higher price gap (i.e. higher expected net losses), requiring a larger price increase in low friction states.

Exclusion Restriction: The third identifying assumption is that prices in high friction states only influence coverage in low friction states through their impact on prices in low friction states. In practice, this means that average prices in high friction states do not vary *within the year* for a

[36] We also note that our instrument is valid for single-state insurers that only operate in low-friction states, because they compete with multi-state insurers; such small insurers only have a 10% market share overall, so they play less of a role in the data.

[37] For example, the *Wall Street Journal* writes about California, a high friction state: " [State Farm's] own actuaries repeatedly said the California subsidiary's premiums weren't high enough and its outside consultants warned of the seriously escalating risk of a devastating fire. [...] 'Our California homeowners insurance product line has not been rate adequate since 2007,' the company said on its website. For years, its in-house calculations had been warning that its rate increases were much lower than what was needed to balance its risks. "Yet State Farm repeatedly asked the state for a fraction of the increases that would be needed to bring its rates up to the necessary levels. In 2021, State Farm asked California regulators to approve a 6.9% rise in home-insurance rates, less than a quarter of the 31% increase its own in-house calculations showed was needed to cover the risks of its policies, state filings show. In 2022, it again asked for 6.9%, rather than the 23% rate its own math indicated was needed. "Insurers in general kept rate requests low to navigate California's tough price controls." See Wall Street Journal, "State Farm Was All In on California—Until It Pulled the Plug Before the Fires", February 6, 2026.

[38] For example, SwissRe forecasts that real insured losses will increase on average by 5-7% annually through 2030. See: SwissRe Natural Catastrophe Report. NOAA offers similar forecasts.

[39] "Rate target" refers to the price increase that insurers' models imply would lead to actuarially fair pricing. Rate received refers to the price increase granted by regulators. The y-axis plots the gap between the target and the rate increase received. We note that the "rate target" differs from the rate change that insurers actually request ("rate request"), which is a separate variable in the rate filings.

given credit score group due to demand shocks in low friction states.

We believe this is likely to hold because the variation in our instrument is driven by when state regulators in high friction states make rate decisions. Insurance regulation in high friction states requires that insurance companies receive approval on their pricing models, including any changes in the loadings on their pricing variables,[40] and any aggregate increases in rates for the state. This means that insurance companies cannot simply adjust prices–either raise them or lower them– without obtaining regulatory approval from the relevant state insurance commissioner. Figure C.13 produces a case study for two of the largest insurers in California, showing that prices in California discretely jump around regulatory filing decisions. This supports our interpretation that the variation in average prices in high friction states are driven by regulatory decisions.

The question then becomes what drives the timing and decisions of regulators in high friction states. We argue that demand shocks in low friction states are unlikely to influence the timing within the year of the regulatory process, or the final rate decisions. We provide several pieces of evidence in support of this assumption.

First, there is significant evidence that the timing and decisions of regulators are difficult to forecast. Regulatory decisions can take months or even years to be approved, and final rate approvals can differ substantially from what insurers require. Milliman, an insurance consulting firm, writes:

> "California insurers may face hurdles to changing prices, even using state-prescribed methodologies. Insurers must submit rate proposals for regulatory review as a normal course of business. But in California, the review period can be particularly lengthy, and filings can be subject to costly public intervention and hearings. An insurer's request for a rate increase may lead to their being forced to take a rate decrease, and effective dates may be delayed many months (sometimes years) beyond what insurers originally request."[41]

Second, a concern may be that high friction state regulators consider demand conditions in low friction states in making approval decisions. We provide evidence that prices in high friction states are unresponsive even to risk measures within those same states, both in levels and in terms of the timing of price changes. Table C.8 shows that there is negligible pass-through of average annual losses (Cotality risk measure) to prices in high friction states, though they have a significant pass-through in low friction states. The results are similar using zip-year level realized loss ratios (Table C.9): while loss ratios are priced in low friction states, they are nor priced in high friction states. Crucially, Oh et al. (2021) show that prices in high friction states do not change in the immediate aftermath of natural disasters in these states (e.g., in the following year) as obtaining regulatory approvals takes significant time. They also verify that prices in high friction states do not respond to losses coming from other states. Given that prices do not even reflect actuarial risks and

[40] Any mathematical function that maps any household or property characteristics to a price. If an insurer uses a linear model, even a change in the coefficient on a particular variable requires regulatory approval.

[41] See Milliman, September 2020

losses in high friction states, the timing and magnitude of regulator decisions in high friction states are unlikely to reflect demand or risk factors in low friction states.

These two pieces of evidence help address several potential concerns about the instrument. For example, a standard concern in the IO literature with Hausman insturments is that there can be coordinated nationwide advertising campaigns, that can induce a common demand shock that impact prices in other states and an own-state price. In our setting, even nationwide advertising campaigns are unlikely to be a concern, because insurers must still wait for the state insurance commissioners to deliberate on their rate request and issue a final decision. A second concern is that the experience of a disaster event in a high friction state could also lead to changes in demand in low friction states, for example if households learn about their own risks when climate losses occur elsewhere. This concern is mitigated by the fact that price changes in high friction states do not coincide with incurred losses (Oh et al. (2021)), meaning that *within-year* changes in prices do not coincide with climate events in high friction states.[42] Lastly, in the climate context, it is possible that with increased disaster-related losses, regulators may perhaps become more lenient over time or concerned about climate change, with households similarly concerned with climate change. Such aggregate trends do not threaten the identification because we include zip-credit score-year fixed effects, ensuring that our identification of the coefficients are obtained from within-year variation. Lastly, there may be concerns about common seasonal demand patterns; in robustness exercises we verify that the inclusion of calendar month effects do not change our estimates, assuaging this concern.

Under these three identifying assumptions, the coefficient estimate β from the 2SLS IV regression in Equation 6 identifies a demand elasticity. We next consider three outcome variables to obtain three related demand elasticities.

4.2. Prediction 4: Overall Demand Elasticities

We first estimate the price elasticity of overall insurance coverage, using coverage levels as the outcome variable. As discussed in Section 2, household demand should be price-sensitive, while mortgage mandate required coverage is mechanically insensitive to insurance prices because it is determined by the loan balance and replacement cost. Therefore, if lender mandates are a quantitatively important driver of observed coverage choices, we expect total coverage demand to be relatively inelastic.

Table 4 reports the main results from estimating Equation 6 for overall coverage. Column 1 shows a statistically significant price elasticity of -0.088. This means that the average household is fairly price inelastic: a 1% increase in insurance prices reduces overall coverage by only 0.088%. Thus, overall coverage responds very little to price changes. As discussed above, the instrument has a strong first stage, with the first-stage F-statistic reported at the bottom of the table. While

[42] Relatedly, we note that it is also difficult for households to observe detailed *pricing* or rate decisions in other states, data on which is relatively difficult to obtain. For example, unlike gasoline or mortgage rates, there are no centralized or widely publicized indices of insurance prices by state. This lack of salient, aggregated price information reduces the plausibility that households in low friction states learn about their own risks or pricing by examining pricing or regulatory decisions in high friction states.

the small elasticity estimate does not by itself prove that lender mandates are binding, since other forces could also generate inelastic demand, it is consistent with lender mandates muting the price responsiveness of observed coverage.

To validate that this inelasticity reflects mortgage market mandates, rather than uniformly low price sensitivity among households, we next test whether elasticities vary with the loan-to-replacement-cost ratio. As discussed in Section 3, L/RC determines the relevant mortgage market benchmark and strongly predicts observed coverage choices. This interpretation predicts that coverage should be least price responsive when mortgage market requirements are tighter (high L/RC), and more price responsive when households have more discretion to adjust coverage above the lender floor (low L/RC).

Columns 2-3 show that both groups have relatively inelastic demand, and the point estimates are consistent with the presence of mandates. Borrowers with low L/RC ratios, below 80%, have a larger elasticity of -0.120. These borrowers are subject to the de facto loan balance rule, and on average purchase coverage well above the loan balance (approximately 150% of the loan amount), meaning this constraint is not strictly binding. Because these borrowers have more discretionary coverage above the lender floor, we expect their total coverage to be more responsive to price changes. In contrast, borrowers with higher L/RC ratios have an elasticity that is statistically insignificant and close to zero. For these borrowers, mortgage requirements are tighter: when $0.8 \leq L/RC < 1$, required coverage tracks the loan balance, and when $L/RC \geq 1$, required coverage tracks replacement cost. Consistent with tighter constraints, observed coverage is closer to the requirement, leaving less room for households to adjust coverage as prices change. As earlier, our first stage F-statistic remains strong across each subsample. Overall, the heterogeneity by L/RC supports the interpretation that mandates are an important reason why overall coverage demand is highly inelastic.

In robustness exercises, we verify in Table Table C.12 that the results are not driven by measurement error in replacement costs by showing the results are robust to excluding this control. Furthermore, in Table C.13, we show that the elasticity estimates are robust to the inclusion of calendar month fixed effects, showing that our results are not driven by common seasonal trends in pricing.

Heterogeneity by DTI ratio. We next examine whether the elasticity within the low L/RC group varies with borrower financial constraints. This test helps distinguish the role of lender mandates from household latent demand. Among low-L/RC borrowers, the loan-balance requirement is the relevant lender floor, but many borrowers purchase coverage above this amount. As a result, this group contains a meaningful discretionary margin on which household characteristics may matter. If price responsiveness partly reflects households' willingness or ability to pay for coverage above the lender floor, we would expect, all else equal, more financially constrained borrowers to have lower latent demand for coverage and to be more elastic, when lender mandates do not bind tightly (e.g., Rampini and Viswanathan (2013), Ericson and Sydnor (2018)). Table 5 shows that this is the case. The elasticity for low-DTI borrowers is -0.073, compared with -0.133 for mid-DTI

borrowers and −0.160 for high-DTI borrowers.[43] Thus, even within the low-L/RC group, where total coverage remains fairly inelastic on average, price responsiveness increases monotonically with DTI. This pattern suggests that household latent demand matters at the margin: borrowers with tighter financial constraints reduce coverage more when prices rise. At the same time, the magnitudes overall remain much smaller, i.e. even high DTI borrowers are quite price inelastic, indicating that lender mandates remain the dominant force behind the overall inelasticity of homeowners insurance demand.

4.3. Prediction 5: Excess Coverage Elasticities

The previous results show that total coverage demand is highly inelastic. However, total coverage combines two conceptually distinct components: the lender-imposed coverage floor and the household's discretionary coverage above that floor. To isolate the component of coverage that more directly reflects household demand, we next study "excess coverage," defined as coverage above the loan balance, $C - L$. We focus on borrowers with low L/RC ratios, for whom the de facto lender rule requires coverage to at least equal the loan balance. In this region, the loan amount is the relevant lender floor, while coverage above the loan balance represents the residual protection potentially available to the household after the lender's claim is satisfied. This margin therefore more closely captures households' willingness or ability to insure their home equity. If the inelasticity documented above is driven by mortgage mandates, rather than by uniformly low household price sensitivity, then excess coverage should be substantially more elastic than overall coverage.

We obtain the excess coverage elasticity using our same instrumental variable specification in Equation 6, using *ln(excess coverage)* as the dependent variable. Table 6 reports the results. Column 1 repeats the overall coverage elasticity for the low L/RC sample: a 1% increase in insurance prices reduces total coverage by only 0.120%. Column 2 uses log excess coverage as the outcome. The estimated elasticity is much larger, at −1.981: a 1% increase in prices reduces excess coverage by nearly 2%. Thus, while overall coverage is nearly price inelastic, the component of coverage above the lender floor is highly elastic. As expected, Column 4 shows that high L/RC households do not display large or statistically significant excess coverage elasticities, which is intuitive since the loan balance is not the binding constraint for such households.

This sharp contrast in elasticities is central to the interpretation of the results. The same price variation that generates only small changes in overall coverage generates large changes in discretionary coverage. This pattern is difficult to reconcile with alternative explanations of low demand elasticities, e.g., households are simply inattentive to insurance prices. Instead, it suggests that households are price sensitive where they have discretion, but that lender requirements mute the response of overall observed coverage. These results indicate that as prices grow, household equity would be significantly more exposed as households continue to reduce their excess coverage,

[43] A DTI threshold of 40% represents a key underwriting cutoff in the mortgage market, above which lenders become substantially more reluctant to extend credit. This institutional feature makes DTI a useful proxy for the tightness of borrowing constraints. We note, however, that DTI heterogeneity may also reflect other differences between high- and low-DTI borrowers, including beliefs, sophistication, or risk preferences.

while lender exposures will be largely covered.

Heterogeneity by DTI ratio. We next examine whether the excess coverage elasticity varies with borrower financial constraints among borrowers in the low-L/RC group. If excess coverage reflects household ability to pay for protection above the lender floor, then financially constrained borrowers should reduce this coverage more aggressively when prices rise. Columns 4–6 of Table 8 shows that this is the case. Among low-DTI borrowers, the excess-coverage elasticity is -1.558. Among high-DTI borrowers, the elasticity is substantially larger, at -2.781. Thus, high-DTI borrowers are nearly twice as responsive as low-DTI borrowers on the discretionary coverage margin. This implies that, as insurance prices rise, and high-DTI households reduce excess coverage, it would leave a larger share of their home equity exposed. A simple back-of-the-envelope calculation illustrates the magnitude: a 10% increase in insurance prices would reduce excess coverage by 28% for high-DTI households, equal to 33% of their average equity.

4.4. Prediction 6: Bunching Elasticities

The excess-coverage estimates show that households reduce discretionary coverage when insurance prices rise. We now test a sharper implication of the mortgage mandate mechanism: higher prices should push some borrowers all the way down to the lender minimum. We therefore evaluate how bunching at the loan balance responds to quasi-exogenous changes in insurance prices for low L/RC households, where the de facto lender requirement is that coverage at least equal the loan balance. Therefore, if rising prices reduce households' latent coverage below the lender minimum, the probability that coverage lies at or near the loan balance should increase. We note that this bunching margin is distinct from the intensive margin response in excess coverage. The excess coverage elasticity measures how much coverage above the mandate declines when prices rise. The bunching response instead measures whether borrowers reduce all the way to the lender floor. An increase in bunching therefore indicates that a larger share of observed coverage is entirely pinned down by mortgage requirements.

We estimate the same IV specification as in Equation 6 except we replace our outcome variable with an indicator that coverage is within narrow bands of the loan amount. In the baseline specification, we define bunching as $C/L \in [0.99, 1.05]$, allowing for rounding and small measurement differences around the threshold.[44] Because the outcome is an indicator and the regressor is log insurance price, the coefficient is a semi-elasticity: a coefficient of β_1 implies that a 1% increase in insurance prices raises the probability of bunching by β_1 percentage points.

Table 7 reports the results. Column 1 shows that higher prices increase bunching in the full sample: a 1% increase in prices raises the probability of bunching at the loan balance by 0.361 percentage points. Column 2 shows that this response is driven by borrowers with low L/RC ratios, for whom the loan balance is the relevant de facto lender floor. In this group, a 1% increase in

[44] In our main table, we consider coverage choices between 99% and 105% of the loan amount to be "bunching." In Table C.14, we consider different bandwidths (within 1% and within 5% of the loan amount, respectively), and show the results are robust.

prices raises the probability of bunching by 0.405 percentage points. By contrast, the estimate for high L/RC borrowers is close to zero and statistically insignificant. The magnitude is economically meaningful. Figure 6 shows that roughly 10% of low-L/RC borrowers bunch at the loan balance at baseline. The estimate in Column 2 implies that a 10% increase in insurance prices would raise this share by about 4 percentage points, from roughly 10% to 14%. Thus, rising prices do not merely reduce discretionary coverage above the lender floor; they also push a meaningful share of borrowers onto the floor itself. This result reinforces the interpretation of the previous elasticities. Overall coverage is highly inelastic because the mortgage mandate component of coverage is difficult to adjust. But households are price sensitive over the discretionary component, and when prices rise enough, some households reduce this discretionary component to zero. For borrowers bunching at the loan balance, insurance coverage primarily protects the lender's claim in a total-loss event, leaving little or no insured residual for household equity after the loan is repaid.

Heterogeneity by DTI ratio. We next examine whether the bunching response is stronger for financially constrained borrowers. Columns 1-3 of Table 8 report estimates within the low L/RC sample. The point estimates are larger for high-DTI borrowers than for low-DTI borrowers: a 1% increase in insurance prices raises bunching by 0.450 percentage points for high-DTI borrowers, compared with 0.366 percentage points for low DTI borrowers. This pattern is consistent with financially constrained households being more likely to reduce discretionary coverage and move down to the lender floor when prices rise. Together with the excess-coverage estimates, the DTI heterogeneity suggests that rising insurance prices disproportionately erode the protection of financially constrained households. High-DTI borrowers reduce coverage above the lender floor more strongly, and they are also more likely to end up at the loan-balance threshold. As a result, lender exposure remains protected, while household equity becomes increasingly exposed.

4.5. Discussion and Robustness

Overall, we find that mortgage market mandates render overall household demand inelastic, while households on their own exhibit elastic demand, evidenced by their high excess coverage elasticities and high bunching elasticities. Furthermore, across all three elasticities, high-DTI households are consistently more responsive to price than low-DTI households. Together, the results suggest that as insurance prices rise, the component of coverage protecting lenders stays largely intact, while the component protecting household equity erodes quickly, leaving homeowners increasingly exposed, especially for the most financially constrained (high-DTI) borrowers.

Taken together, the three elasticities along with the heterogeneity patterns also help address alternative explanations for the overall small elasticity estimate that we obtain. A uniformly low degree of attention, inertia, beliefs about climate risk, or search effort would not naturally predict that elasticities vary systematically with the tightness of lender requirements, nor that excess coverage is much more price sensitive, or increased bunching at the loan-balance threshold, as prices rise. A related concern is that the inelastic overall coverage response may reflect actuarially favorable pricing rather than lender mandates. We show that this is unlikely by comparing high- and low-friction

states. Contracts in high-friction states are more likely to be priced near, or even below, actuarially fair levels, while contracts in low-friction states feature higher markups.[45] If below actuarially favorable pricing were the main reason demand is inelastic, households in high friction states should hold substantially more complete coverage. Instead, they remain substantially underinsured, with average coverage ratios around 70%, similar to low friction states (Table 1). Moreover, Table C.10 shows that coverage elasticities are similar in high and low friction states, despite large differences in contract generosity.[46] These results suggest that actuarial fairness is unlikely to be the primary driver of the observed inelasticity. Instead, the evidence is more consistent with inelastic demand arising from mortgage mandates influencing a large component of overall coverage.

In the next section, we show that insurance proceeds protect lenders by facilitating mortgage prepayment after total-loss events. If households use all insurance proceeds to prepay the mortgage, they receive only the residual amount, $C - L$. This residual effectively determines the extent to which insurance coverage protects the household's equity.

5. Rationalizing Mortgage Mandates

A common prior in the public discourse is that households primarily use insurance proceeds to finance repair and rebuilding after natural disasters. However, this view is difficult to reconcile with the de facto mortgage market mandate that forces coverage to exceed the loan amount for most households, rather than the replacement cost. In this section, we show that requiring coverage above the loan balance helps protect the exposures of lenders, GSEs, and servicers ("lenders") by shifting households from defaulting to prepaying after natural disasters.[47] As a result, these mandates protect lenders, while uninsured losses are borne by households through lost equity.

To understand how insurance coverage choices affect household outcomes after a large natural disaster, we exploit the landfalls of Hurricanes Harvey and Irma, two major storms that struck six states during August-September 2017. These storms offer a natural setting to examine how insurance coverage affects mortgage performance after severe disaster exposure. We ask whether the loan-balance threshold matters after a disaster by examining whether borrowers with coverage just

[45] This follows from combined-ratio data, a measure of insurer profitability. Loss ratios in high-friction states are persistently close to one, while loss ratios in low-friction states are persistently below one, consistent with higher markups in low-friction states (Figure C.11b).

[46] Because the timing and magnitude of price changes in high-friction states are driven by regulators and are less directly tied to contemporaneous expected losses or demand conditions, an OLS specification in high-friction states provides a useful comparison for demand elasticities.

[47] Though mortgage servicers are not directly exposed to the credit losses on loans they service on behalf of third parties, their incentives generally align with lenders' in preferring prepayment to mortgage delinquency or forbearance. Mortgage nonpayment is costly for servicers: GSE and FHA servicers must continue to advance principal, interest, taxes, and insurance payments to investors while a loan is delinquent, which can stress liquidity-constrained nonbank servicers, in particular (Cherry et al. (2022), Kim et al. (2024)). During the Harvey and Irma hurricanes, servicers for Fannie Mae, Freddie Mac, and the FHA were required to automatically offer short-term forbearance for non-paying loans in the disaster declaration area, and curing those nonpaying loans often required evaluation for mortgage modification, or other high-touch interventions. Federally-regulated banks were also mandated to offer forbearance where possible. (Kousky et al. (2020), del Valle et al. (2024)))

above the loan balance are more likely to prepay, and less likely to become delinquent, after a major natural disaster than otherwise similar borrowers with coverage just below the loan balance.

Specifically, our empirical design exploits three sources of variation: time, geography, and coverage relative to the loan balance. First, we compare outcomes before and after the hurricanes. Second, we compare counties that were severely affected by the storms to less affected counties, defining severely hit counties as those in the top quartile of SHELDUS property damage per capita.[48] Third, within narrow bandwidths around the $C/L = 1$ threshold, we compare borrowers whose coverage is just sufficient to repay the mortgage ($C/L \geq 1$) to borrowers whose coverage falls just short of the loan balance ($C/L < 1$). We refer to policies with $C/L \geq 1$ as "loan-sufficient": the coverage amount is high enough to repay the outstanding mortgage in a total-loss event, though not necessarily high enough to rebuild the home. We combine these three sources of variation in the following triple-difference specification:

$$\begin{aligned} Y_{i,c,t} = {} & \beta_1(\text{Post}_t \times \text{SeverelyHit}_c \times \mathbb{I}[C/L \geq 1]_i) + \beta_2(\text{Post}_t \times \text{SeverelyHit}_c) \\ & + \beta_3\left(\text{Post}_t \times \mathbb{I}[C/L \geq 1]_i\right) + \beta_4(\text{SeverelyHit}_c \times \mathbb{I}[C/L \geq 1]_i) + \beta_5(\mathbb{I}[C/L \geq 1]_i) \\ & + \delta_c + X'_{i,t}\theta + \varepsilon_{i,c,t} \end{aligned} \tag{7}$$

where the variable $\mathbb{I}[C/L \geq 1]_i$ indicates that a household's coverage exceeds its loan balance, SeverelyHit_c is a dummy indicating that the county is in the top quartile of property damage, and $Post_t$ takes a value of 1 for each month after the hurricanes, for borrower i in county c in month-year t. The variable X refers to borrower characteristics interacted with month-year fixed effects, which include contemporaneous LTV ratio, credit score, and interest rate bins, which we argue capture unobservable dimensions of credit risk known to the lender at origination. We also include origination cohort fixed effects. We consider two outcome variables: whether the mortgage prepays or becomes delinquent. We limit to the six states that were hit by the storm, and consider outcomes 6 months before and after the storm. We estimate the design within narrow bandwidths around $C/L = 1$ and report results using both $\pm 20\%$ and $\pm 10\%$ bandwidths. The coefficient on the triple interaction term β_1 is the estimate of interest, and captures the additional effect of the disaster on mortgage outcomes in severely hit areas specifically for household with insurance coverage exceeding the loan balance. The other coefficient of interest, β_2, captures the effect of the disaster in severely hit counties for households with coverage below the loan balance.

To implement this design, we restrict attention to households with $L/RC > 0.8$. We do so to obtain variation in the C/L ratio that is driven by the change in the mortgage mandate. Specifically, households with $0.8 < L/RC < 1$ have required coverage tracking their loan balance, whereas households with $L/RC > 1$, have required coverage below loan balance as the binding limit is replacement cost. This allows us to isolate the effect of being just above versus just below the threshold, comparing households whose insurance proceeds are larger than the loan balance.

Our empirical design requires two key identifying assumptions. First, absent hurricanes, mortgage

[48] Our definition of property damage excludes flood damage (which is not covered by homeowners' insurance).

outcomes in severely hit counties would have followed similar trends to those in less affected counties for each coverage group (parallel-trends assumption). Second, we also assume that other borrower and property characteristics vary smoothly through the coverage-to-loan-balance ($C/L = 1$) threshold, so that within narrow bandwidths, borrowers just above and below the threshold are comparable, differing primarily by whether their coverage is sufficient to repay the mortgage. Under these identifying assumptions, the triple-difference coefficient identifies the causal effect of sufficient insurance coverage on prepayment and default behavior in the aftermath of a major disaster.

Table 9 presents the results. Columns 1–2 show that, among severely hit properties, having coverage just above the loan balance substantially increases the probability of prepayment after the hurricanes. Our triple-difference coefficient indicates that, among borrowers in severely hit counties, those just above the $C/L = 1$ threshold are 50-60 basis points more likely to prepay their mortgage than households just below the threshold. The effect is economically meaningful. Contemporary reports suggest that roughly 2% of homes in severely hit counties experienced a total loss. If the prepayment response is concentrated among properties experiencing total losses, a back-of-the-envelope calculation using our prepayment estimates suggests that roughly 25% of households that faced a total loss prepaid their mortgage when their insurance coverage was sufficient.

In contrast, Columns 3–4 show the mirror-image result for mortgage nonpayment: having coverage just below the loan balance substantially increases the probability of default.[49] Borrowers just above the threshold are 90-120 basis points less likely to default than borrowers just below the threshold after the hurricanes. The effects are concentrated in severely hit areas and emerge only after the storm, consistent with the proposed mechanism related to the storm, rather than a pre-existing trend or endogenous sorting of borrowers into the two C/L groups.

Furthermore, other coefficients help support our identifying assumptions. Households with $C/L \geq 1$ do not differentially prepay or default more in the pre-period in general (i.e. $\beta_4, \beta_5 \approx 0$), casting doubt on the concern that such households could generally have a differential propensity to prepay or default.

Taken together, these results suggest that requiring insurance coverage that exceeds the loan balance pushes households away from delinquency and towards prepayment. These results help rationalize the loan-balance component of mortgage market mandates. From the perspective of mortgage creditors and servicers, prepayment is substantially preferable to delinquency or foreclosure after a disaster, particularly given concerns about limited recovery in foreclosure following natural disasters when properties have been damaged and land values decline. Even if a loan is only temporarily in default, for GSE loans, the servicer must advance payments to MBS investors and absorb the costs of delinquency which can be substantial (e.g., contacting borrowers, negotiating modifications, covering property taxes and insurance premiums). These costs are not fully offset since households do not pay interest on these advances, leaving servicers with a time-value loss. By

[49] During this set of storms, most non-paying mortgage loans were automatically granted forbearance. Although many borrowers end up curing eventually, mortgage forbearance and subsequent borrower contact is still costly for servicers, making prepayment preferred for servicers even when borrowers eventually cure.

contrast, prepayment using insurance proceeds allows the mortgage claim to be repaid immediately and lets the servicer fully recoup the loan exposure immediately.[50] Our results suggest that coverage at or above the loan balance facilitates this repayment channel.

The household implications are more nuanced. Prepayment is likely better for borrowers than default or foreclosure, since it avoids the loss of accrued equity, credit-score damage, and the loss of property through foreclosure. Even if households eventually recover from the mortgage default, they still bear the costs associated with their missed payments. By contrast, in a prepayment, households do not bear any costs of default, and also retain full ownership of the (damaged) property. Households may recoup additional proceeds by selling the damaged home; this may offset some of the equity loss they would otherwise face. However, prepayment does not mean the household is fully protected. In particular, the comparison to their pre-storm position is less clear, and depends on the remaining value of the damaged property, particularly recovery proceeds from selling the land and costs they incur in doing so, meaning some households may still end up worse off than they were before the disaster.

We emphasize that our results suggest that the commonly held view–which is that households use insurance proceeds to rebuild–seems vastly incomplete. Lender mandates seem geared around inducing households to prepay rather than default, thereby protecting lender exposures in natural disaster events.[51] Once the mortgage is repaid, the household retains only the residual insurance proceeds, $C - L$, along with any remaining value of the damaged property. If coverage is well below replacement cost, this residual may be far below the amount needed to rebuild. Thus, lender mandates can improve mortgage outcomes while still leaving households substantially underinsured relative to the cost of reconstruction.

6. Conclusion and Broader Implications

This paper provides some of the first evidence on the importance of mortgage lender mandates in the homeowners insurance market using state-of-the-art mortgage, climate, and insurance databases at a property level. Rather than reflecting pure household preferences, as is commonly assumed, we show that coverage choices are driven in a first-order way by a defacto lender rule requiring the minimum of the loan balance and replacement cost. We find significant compliance with this rule, and coverage moves in tandem with loan balances and replacement cost depending on which is required. For 10% of households, they obtain the loan amount directly, the bare minimum that lenders require. Adherance to lender mandates renders overall coverage demand to be inelastic, since mandates do not depend on price. In contrast, households exhibit highly elastic demand for the component of coverage that

[50] The only loss is the foregone interest stream. Notably, from the MBS investor's perspective, default and prepayment look similar — in both cases the investor gets their principal back early and loses the future income stream — so the investor is largely indifferent between the two outcomes.

[51] Our results align with several papers that broadly document enhanced mortgage prepayment and less rebuilding after natural disasters (e.g. Gallagher and Hartley (2017), Issler et al. (2019), Biswas et al. (2023)). Relative to these papers, we show that this increase in prepayment an intended consequence of mortgage mandates, which are designed to protect lender and servicer exposures.

exceeds what lenders require and protects their equity. Households are also significantly more likely to bunch at the lender mandates when insurance prices rise. These results imply that as insurance prices continue to rise, lender exposures will remain largely protected, while homeowner equity — particularly among the most financially constrained borrowers — will bear a disproportionate and growing share of uninsured risk. An event study following Hurricane Irma/Harvey explains the rationale behind lender mandates: requiring households to obtain coverage in excess of the loan balance shifts households from defaulting after natural disasters to prepaying, because households can use insurance proceeds to prepay the mortgage. This unambiguously improves lender recovery, with households also being better off in a prepayment rather than a default.

We conclude with a discussion of policy implications. Our results raise a broader question: if a central function of homeowners insurance mandates is to protect *mortgage lenders* in total loss events, is the current institutional arrangement the most effective way to provide that protection? One alternative would be for lenders to forgo insurance mandates and instead price additional default risk directly into interest rates. In practice, however, several frictions may make this difficult. GSEs face political constraints in pricing climate risk (Hurst et al., 2016, Sastry et al., 2024). Similar constraints may limit lenders' ability to purchase insurance directly if costs cannot be passed through to borrowers, even if the premiums lenders pay may bake in some benefits from geographic diversification. Moreover, when default and foreclosure generate deadweight losses, both lenders and households may prefer insurance arrangements that prevent default rather than contracts that merely compensate lenders for default risk ex-ante (Ganong and Noel, 2020). Alternative arrangements, such as selective mandates based on climate risk or prices, could also create principal-agent problems between lenders and insurers. For example, insurers may interpret selective lender enforcement as a signal of private information about borrowers or properties, leading to market unraveling. We therefore view the current system not as an idiosyncratic contractual feature, but likely as an institutional response to constraints on risk-based pricing, the costs of default, and informational frictions between mortgage and insurance markets.

We also consider the broader implications of our results for insurance and mortgage markets. Mortgage lenders effectively impose an "individual mandate" in insurance markets, leaving households with limited discretion over coverage choices and potentially reducing adverse selection concerns that arise in other insurance markets (Finkelstein and Poterba (2014)). This mandate creates a lower bound on the size of insurance markets, as households have limited ability to reduce coverage before reaching binding lender thresholds.[52] The resulting price insensitivity helps explain how insurers sustain cross-state subsidies (Oh et al., 2021) and provides a rationale for rate regulation and policymaker concerns about markups when insurer competition is limited.

Our findings also shed light on the design of mortgage contracts, and help explain the limited evidence of pricing of climate risk in mortgage interest rates (e.g., Sastry (2025), Santos and Blickle (2022)). Lenders appear to address climate risks primarily through insurance mandates, potentially

[52]This constraint holds only until rising premiums bind debt-to-income (DTI) limits, at which point higher insurance costs may begin to restrict mortgage availability and, in turn, insurance demand. See Meisenzahl et al. (2026).

reflecting political constraints on pricing, insurers' expertise in risk assessment, or lenders' desire to avoid the deadweight losses associated with default and foreclosure. However, while these mandates mitigate climate-related default risk, they simultaneously increase prepayment risk. Our results suggest that climate risk is likely to manifest in the form of growing prepayment risk in mortgage markets, if insurance markets continue to function well, which has important implications for mortgage-backed securities (MBS) investors.

Finally, whether this institutional structure is *optimal* remains an open question, particularly from the perspective of household welfare. Our results nevertheless demonstrate that mortgage mandates play a central role in shaping homeowners insurance demand and determining how climate risk is allocated across households, insurers, and mortgage lenders.

Tables and Figures

Figures

Figure 1: Mortgage Market Mandates

This figure in panel a visually depicts the GSE rule (blue) and defacto rule (red) for required coverage as a function of the loan-to-replacement-cost (L/RC) ratio. The y-axis refers to the coverage ratio. The two rules align when $L/RC > 0.8$. Panel b shows the required coverage in table form as a function of two key ratios which we will evaluate directly: coverage ratio (C/RC) and the coverage-to-loan-balance ratio (C/L).

Panel A: Mortgage Market Mandates

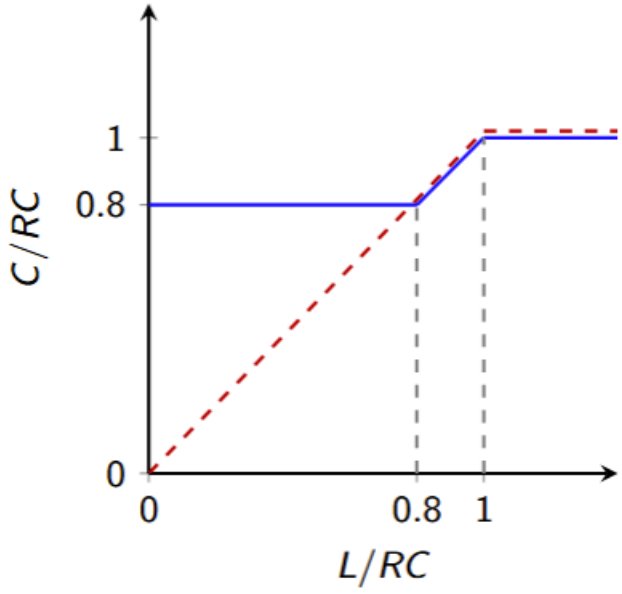

Panel B: Minimum required coverage under the mortgage market mandates

	GSE Rule	De Facto Rule
$L/RC < 0.8$	$C/RC = 0.8$	$C/L = 1$
$0.8 < L/RC < 1$	$C/L = 1$	$C/L = 1$
$L/RC > 1$	$C/RC = 1$	$C/RC = 1$

Figure 2: Compliance with Mortgage Market Rules

This figure shows on average whether observed coverage choices lines up with the mortgage market rules discussed in Figure 1 by plotting coverage ratio C/RC and coverage-to-loan-balance C/L for the three regions of L/RC defined by the GSE rule. C refers to coverage; RC refers to replacement cost; and L refers to loan balance.

(a) Coverage Ratio

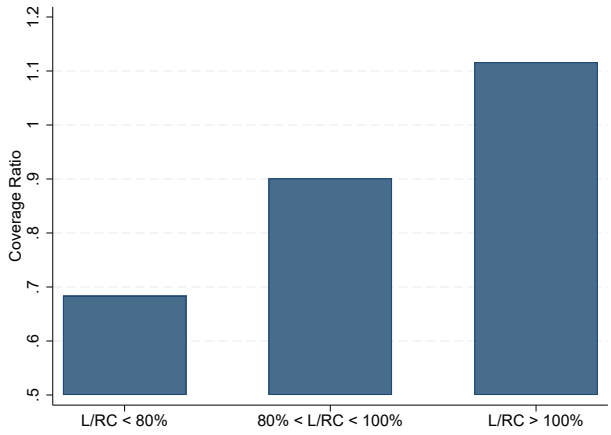

(b) Ratio of Coverage to Loan Balance

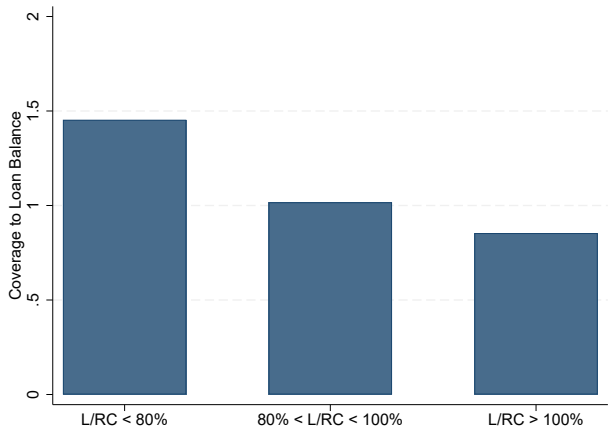

Figure 3: Explaining the Variation in Insurance Coverage

The bar chart shows the adjusted R^2 from a series of regressions with thematic subsets of control variables explaining insurance coverage. The lender rule is coded as the minimum of the loan balance and replacement cost. Borrower and structure controls refer to debt-to-income, credit score, loan-to-value, climate risk, borrower age, house price per square foot, structure size, and building age. Other fixed effects are included as indicated. The figure is based on the main analysis sample, outlined in Section 1.

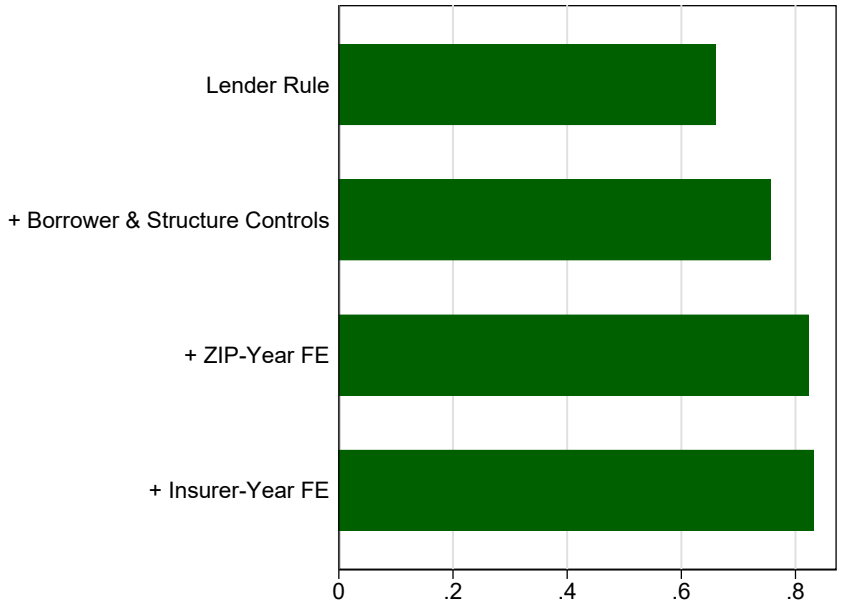

Figure 4: Kink Design: Predictions

This figure shows the implied regression coefficients that would be estimated in Equation 3 on loan balance (panel a) and replacement costs (panel b) if coverage perfectly followed the GSE rule (blue line) or the de facto rule (red line).

(a) Loan amount

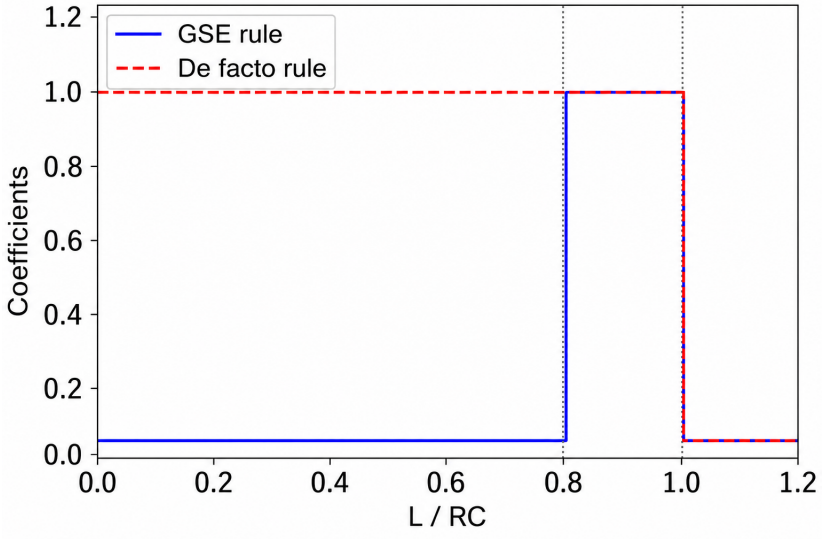

(b) Replacement cost

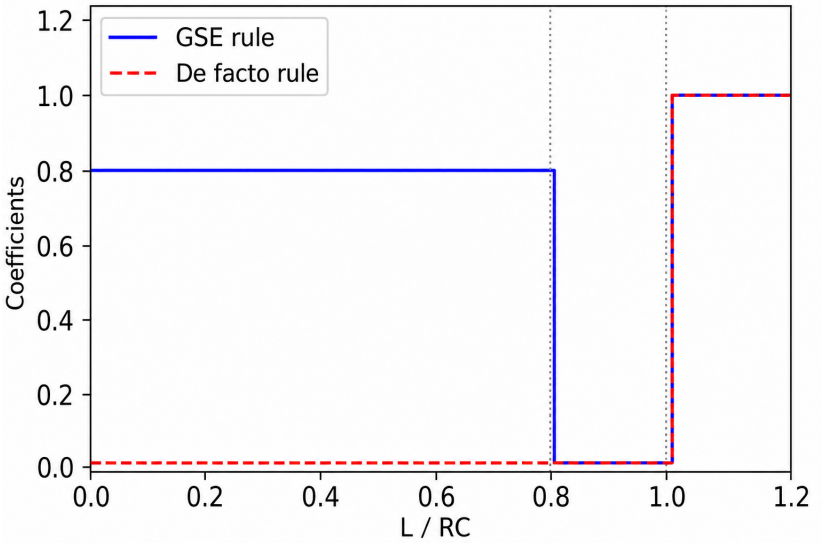

Figure 5: Kink Design: Coverage sensitivity to loan amount and replacement cost

This figure shows a coefficient plot and 95% confidence intervals obtained from estimating Equation 3 in the text. Controls include borrower characteristics (including LTV, DTI, and credit score) × year fixed effects, price per square foot, ZIP × Year, and Insurer × Year fixed effects. The data are at the time of mortgage origination, and spans the period from 2013 to 2023. Standard errors are clustered at the county-year level.

(a) Loan amount

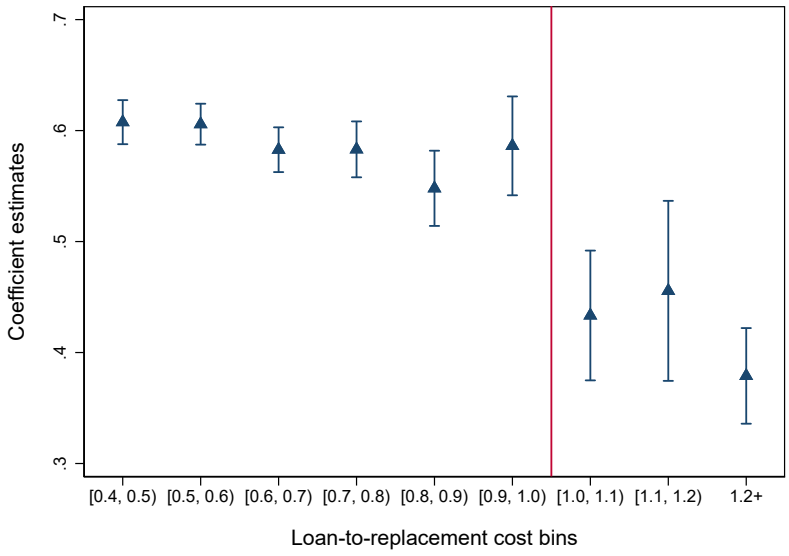

(b) Replacement cost

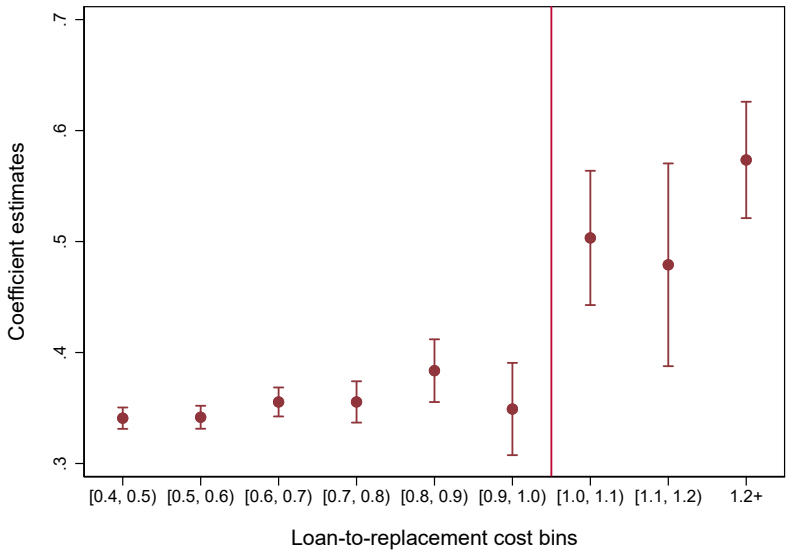

Figure 6: Bunching of Coverage at Loan Amount

This figure shows a histogram of the ratio of coverage amounts to loan amounts at origination for borrowers with the ratio of loan-to-replacement-costs less than 80%. The data are at the time of mortgage origination, and spans the period from 2013 to 2023.

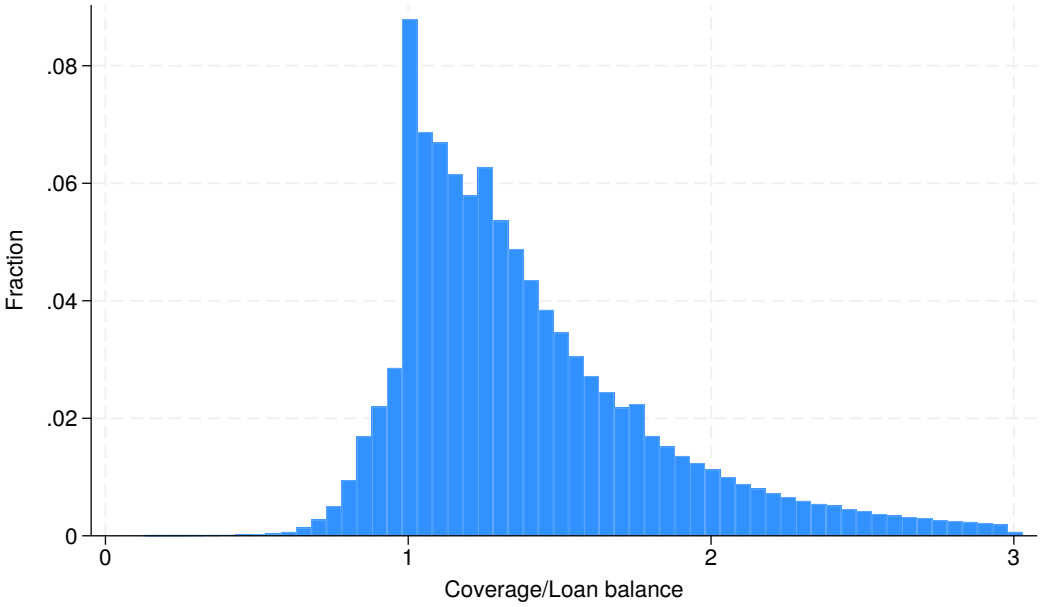

Tables

Table 1: Summary Statistics

This table presents the summary statistics for the main analysis sample, outlined in Section 1. High climate risk refers to properties in the top decile of climate risk by CoreLogic's AAL measure. Coverage amount is from ICE-McDash. Coverage ratio is defined as insurance coverage amount divided by the estimated replacement costs. Insurance price is the monthly insurance premium per $100,000 of coverage. We group states into two types: high and low friction, following the classification in Oh et al. (2021). High friction are states in the top tercile of regulatory friction and low friction are all the remaining states.

	All States		Low Friction		High Friction	
	(1) Mean	(2) SD	(3) Mean	(4) SD	(5) Mean	(6) SD
Coverage	277770.2	141507.5	275422.1	142332.3	280602.3	140454.0
Replacement cost	379187.9	207757.9	376988.2	215509.2	381841.0	197972.8
Loan amount	232933.8	137792.3	234605.5	139407.8	230917.5	135790.8
Coverage Ratio (C/RC)	0.75	0.28	0.75	0.28	0.76	0.27
Loan-to-replacement cost Ratio (L/RC)	0.64	0.33	0.64	0.33	0.63	0.34
Coverage-to-loan balance Ratio (C/L)	1.34	0.57	1.32	0.57	1.37	0.57
Share with coverage near loan balance	0.11	0.32	0.11	0.31	0.12	0.32
Premium	1127.8	609.0	1147.8	640.9	1103.6	567.1
Insurance price	37.9	20.7	38.9	21.9	36.6	19.0
Share of High Climate Risk	0.096	0.30	0.092	0.29	0.10	0.30
Building (Sq ft)	1979.8	825.4	1996.6	832.4	1959.3	816.2
Building age	38.1	28.6	37.4	28.6	39.0	28.5
Borrower age	40.9	12.7	41.4	13.0	40.3	12.4
DTI Ratio	35.8	9.55	35.5	9.64	36.1	9.43
LTV Ratio	86.0	12.7	85.7	13.0	86.5	12.4
Credit score	734.4	51.4	735.4	51.5	733.2	51.2
House price	280999.6	186395.2	284273.2	189211.4	277051.1	182862.9
Observations	3215348		1757908		1457440	

Table 2: Changes in Pass-through to Coverage at Mortgage Rule Kink: 5% Bandwidth

This table shows the results from estimating the regression design as shown in Equation 4 where we estimate how the sensitivity of coverage amount to loan amount and replacement cost changes at the loan to replacement cost (LTRC) threshold of 1. $LTRC > 1$ is a dummy variable that takes the value of 1 if the borrower's loan to replacement cost ratio exceeds 1 and otherwise it is 0. We limit the sample to loans within 5% of the $LTRC = 1$ threshold. We include various combinations of controls and fixed effects as denoted at the bottom of the table, including borrower-year fixed effects, structure controls (including the house price per square foot), zip-year fixed effects, and insurer-year fixed effects. Borrower controls refer to credit score-year fixed effects, DTI-year fixed effects, ltv-year fixed effects, and borrower age. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Coverage$_{i,z,t}$			
	(1)	(2)	(3)	(4)
Loan amount	0.718***	0.646***	0.759***	0.722***
	(0.080)	(0.081)	(0.080)	(0.077)
$LTRC > 1 \times$ Loan amount	-0.281**	-0.272**	-0.255**	-0.232**
	(0.112)	(0.113)	(0.112)	(0.109)
Replacement Cost	0.288***	0.355***	0.234***	0.255***
	(0.078)	(0.080)	(0.078)	(0.076)
$LTRC > 1 \times$ Replacement cost	0.280**	0.270**	0.254**	0.232**
	(0.112)	(0.113)	(0.112)	(0.110)
Borrower-Year FE	N	N	Y	Y
Structure controls	N	Y	Y	Y
ZIP-Year FE	Y	Y	Y	Y
Insurer-Year FE	N	N	N	Y
Observations	98693	95156	95153	94870
Adjusted R^2	0.817	0.823	0.830	0.838

Table 3: Changes in Pass-through to Coverage at Mortgage Rule Kink: Robustness to Bandwidths

This table shows the results from estimating a regression kink deign as shown in Equation 4 where we consider how the sensitivity of coverage amount to loan amount and replacement cost changes at the loan to replacement cost (LTRC) threshold of 1. We consider the most stringent specification from Table 2 with the most controls and fixed effects but vary the bandwidth limit to include loans within 15% (column 1), 10% (column 2), and 5% (column 3) of the $LTRC = 1$ threshold. All variables and fixed effects are as defined in Table 2. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Coverage$_{i,z,t}$		
	+/- 15%	+/- 10%	+/- 5%
	(1)	(2)	(3)
Loan amount	0.619***	0.633***	0.722***
	(0.013)	(0.024)	(0.077)
$LTRC > 1 \times$ Loan amount	-0.131***	-0.123***	-0.232**
	(0.020)	(0.038)	(0.109)
Replacement Cost	0.349***	0.337***	0.255***
	(0.012)	(0.023)	(0.076)
$LTRC > 1 \times$ Replacement Cost	0.133***	0.123***	0.232**
	(0.020)	(0.039)	(0.110)
Borrower-Year FE	Y	Y	Y
Structure Controls	Y	Y	Y
ZIP-Year FE	Y	Y	Y
Insurer-Year FE	Y	Y	Y
Observations	328286	206880	94870
Adjusted R^2	0.836	0.837	0.838

Table 4: Overall Coverage Demand Elasticities

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 6. Column 1 shows the baseline elasticities on the full sample. Columns 2-3 shows sub-samples by whether loan to replacement cost ratios are above or below 80%. The analysis uses the sample of low friction states. The dependent variable is ln coverage amount. The independent variable is ln insurance price (monthly premiums per $100,000 of coverage), instrumented using average prices in high friction states for each credit score group in each month. High and low friction states are as defined following Oh et al. (2021). All specifications include ZIP-credit score-year fixed effects, and additional fixed effects and controls for property and borrower characteristics, as indicated. Property-level controls include a dummy for whether the property is classified high climate risk (average annual loss is in the top decile), building age fixed effects, and the property's replacement cost. Borrower controls include house price, debt-to-income ratios, and borrower age fixed effects. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln coverage$_{i,c,z,m,t}$		
	(1) Overall	(2) Low L/RC	(3) High L/RC
ln $\widehat{P_{i,c,z,m,t}}$	-0.088**	-0.120***	0.039
	(0.032)	(0.023)	(0.088)
ZIP × Year × FICO FE	Y	Y	Y
Insurer × Year FE	Y	Y	Y
Property Controls	Y	Y	Y
Borrower Controls	Y	Y	Y
First Stage F Stat	164	207	32
Observations	1,610,605	1,237,437	307,734
Adjusted R^2	0.697	0.725	0.590

Table 5: Coverage Demand Elasticities: Heterogeneity by DTI Ratio

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 6 for low loan-to-replacement cost households. The first column represents the overall results for this group from Column (2) in Table 4. Columns 2-4 then present the results for three sub-samples: those with low, medium, and high debt to income ratios (Columns 1 and 2), and those with above versus below median house price (Columns 3 and 4). The analysis uses the sample of low friction states. The dependent variable is log coverage amount. The independent variable is log insurance price (monthly premiums per $100,000 of coverage), instrumented using average prices in high friction states for each credit score group in each month. High and low friction states are as defined in Section 1. All specifications include ZIP-credit score-year fixed effects, and additional fixed effects and controls for property and borrower characteristics, as indicated. Control variables are as defined in Table 4. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln coverage$_{i,c,z,m,t}$			
	(1)	(2)	(3)	(4)
	Overall	Low DTI (< 30)	Mid DTI	High DTI (> 40)
$\ln \widehat{P_{i,c,z,m,t}}$	-0.120***	-0.073**	-0.133***	-0.160***
	(0.023)	(0.029)	(0.023)	(0.022)
ZIP × Year × FICO FE	Y	Y	Y	Y
Insurer × Year FE	Y	Y	Y	Y
Property Controls	Y	Y	Y	Y
Borrower Controls	Y	Y	Y	Y
Observations	1237437	305466	365688	375612
Adjusted R^2	0.725	0.736	0.729	0.719
First Stage F Stat	207.125	446.022	550.788	466.620

Table 6: Excess Coverage Demand Elasticities

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 6, except we consider in Column 2 the component of coverage that exceeds lender mandates (loan amount). In columns 1 and 2, we limit the sample to households with loan-to-replacement ratios below 80%, where the loan amount is the binding rule. In columns 3 and 4, the sample is limited to households with high loan-to-replacement cost ratios. In Columns 1 and 3, the dependent variable is log of coverage. In Columns 2 and 4, the dependent variable is log of the excess coverage amount, defined as coverage in excess of the loan amount. The independent variable is log insurance price (monthly premiums per $100,000 of coverage), instrumented using average prices in high friction states for each credit score group in each month. High and low friction states are as defined in Section 1. All specifications include ZIP-credit score-year fixed effects, and additional fixed effects and controls for property and borrower characteristics, as indicated. Control variables are as defined in Table 4. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Low L/RC		High L/RC	
	ln coverage$_{i,c,z,m,t}$	ln excess coverage$_{i,c,z,m,t}$	ln coverage$_{i,c,z,m,t}$	ln excess coverage$_{i,c,z,m,t}$
	(1)	(2)	(3)	(4)
ln $\widehat{P_{i,c,z,m,t}}$	-0.120***	-1.981***	0.039	-0.388
	(0.023)	(0.136)	(0.088)	(0.473)
ZIP × Year × FICO FE	Y	Y	Y	Y
Insurer × Year FE	Y	Y	Y	Y
Property Controls	Y	Y	Y	Y
Borrower Controls	Y	Y	Y	Y
First Stage F Stat	207	195	32	31
Observations	1,237,437	1,237,437	307,734	307,734
Adjusted R^2	0.725	0.890	0.590	0.948

Table 7: Bunching elasticities: how bunching at loan balance varies with prices

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 6. The dependent variable is an indicator equal to one if coverage is close to the loan balance, defined as lying within the interval [0.99, 1.05]. The independent variable is log insurance price (monthly premiums per $100,000 of coverage), instrumented using average prices in high friction states for each credit score group in each month. Column 1 shows the baseline elasticities using the same sample, controls, and fixed effects, as described in Table 4. Column 2 and 3 shows sub-samples by whether loan to replacement cost ratios are above or below 80%. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Bunching at loan amount (Y/N)		
	(1) Overall	(2) Low L/RC	(3) High L/RC
$\ln \widehat{P_{i,c,z,m,t}}$	0.361***	0.405***	-0.013
	(0.033)	(0.031)	(0.110)
ZIP × Year × FICO FE	Y	Y	Y
Insurer × Year FE	Y	Y	Y
Property Controls	Y	Y	Y
Borrower Controls	Y	Y	Y
First Stage F Stat	164	207	32
Observations	1,610,605	1,237,437	307,734

Table 8: Bunching and Excess Coverage Elasticities: Heterogeneity by DTI Ratio

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 6, limited to the sample of households with low loan-to-replacement costs. The independent variable is log insurance price (monthly premiums per $100,000 of coverage), instrumented using average prices in high friction states for each credit score group in each month. The dependent variable in columns 1-3 is an indicator equal to one if coverage is close to the loan balance, defined as lying within the interval [0.99, 1.05]. The dependent variable in columns 4-6 is log of the excess coverage amount, defined as coverage in excess of the loan amount. Columns 1 and 4 re-eat the values from Columns 2 of Table Table 6 and Table 7, respectively. The subsequent columns show heteregeneity by DTI ratois, as indicated. Control variables are as defined in Table 4. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Bunching (Y/N)			ln excess coverage		
	(1) Overall	(2) Low DTI (< 30)	(3) High DTI (> 40)	(4) Overall	(5) Low DTI (< 30)	(6) High DTI (> 40)
$\ln \widehat{P_{i,c,z,m,t}}$	0.405***	0.366***	0.450***	-1.981***	-1.558***	-2.781***
	(0.031)	(0.049)	(0.050)	(0.136)	(0.186)	(0.201)
ZIP × Year × FICO FE	Y	Y	Y	Y	Y	Y
Insurer × Year FE	Y	Y	Y	Y	Y	Y
Property Controls	Y	Y	Y	Y	Y	Y
Borrower Controls	Y	Y	Y	Y	Y	Y
First Stage F Stat	207	446	467	195	429	422
Observations	1,237,437	305,466	375,612	1,237,437	305,466	375,612
Adjusted R^2	-0.095	-0.088	-0.113	0.890	0.893	0.873

Table 9: Prepayment and Delinquency after a Disaster

This table reports the estimates of the triple difference-in-differences specification from Equation 7 for mortgage prepayment (Columns 1-2), and delinquency of more than one missed payment (Columns 3-4). We limit the sample to narrow bandwidths of 20% (Columns 1, 3) or 10% (Columns 2, 4) around the the threshold where coverage equals loan balance ($C/L = 1$), and also consider households with loan-to-replacement-cost ratios above 80%. We further limit the sample period to six months before and after the landfall of Hurricanes Irma/Harvey in 2017, and consider the six states that were hit by the hurricanes. SevereleyHit indicates counties in the top quartile of property damage per capita in SHELDUS (excluding flood damage). The specification includes controls and fixed effects as indicated, where borrower characteristics are interacted with month-year fixed effects. Characteristics refer to origination interest rate, loan-to-value ratio measured in December 2016, and credit score measured in December 2016. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Prepayment (Y/N)		Delinq. 30D+ (Y/N)	
	(1) ± 20%	(2) ± 10%	(3) ± 20%	(4) ± 10%
$C/L \geq 1$	0.001	0.001	0.001	0.000
	(0.001)	(0.002)	(0.001)	(0.001)
$PostHurricane \times C/L \geq 1$	-0.002	-0.001	-0.000	0.001
	(0.002)	(0.002)	(0.001)	(0.002)
$PostHurricane \times SeverelyHit$	-0.002	-0.004**	0.014***	0.012***
	(0.001)	(0.002)	(0.003)	(0.003)
$C/L \geq 1 \times SeverelyHit$	-0.003*	-0.004	0.000	-0.001
	(0.002)	(0.002)	(0.001)	(0.002)
$PostHurricane \times C/L \geq 1 \times SeverelyHit$	0.005**	0.006**	-0.011***	-0.008**
	(0.002)	(0.003)	(0.003)	(0.003)
ZIP-Year FE	Y	Y	Y	Y
Borrower-Year FE	Y	Y	Y	Y
Origination Cohort FE	Y	Y	Y	Y
Observations	93,788	49,568	90,870	48,090
Adjusted R^2	0.003	0.003	0.039	0.040

References

Agarwal, S., I. Ben-David, and V. Yao (2015). Collateral valuation and borrower financial constraints: Evidence from the residential real estate market. Management Science 61(9), 2220–2240.

Ambrose, B. W., J. N. Conklin, N. E. Coulson, M. Diop, and L. A. Lopez (2026, 03). Do appraiser and borrower race affect mortgage collateral valuation? Review of Finance 30(2), 639–679.

Amornsiripanitch, N., S. Biswas, J. Orellana, and D. Zink (2024). Flood underinsurance (december 04, 2024). SSRN Working Paper Number 5043931.

Anderson, N. B. and J. Dokko (2008). Mortgage delinquency and property taxes. Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association.

Annan, F. (2021). Does bundling induce adverse selection in insurance? Economics Letters 199, 109665.

Arrow, K. (1963, Dec). Uncertainty and the welfare economics of medical care. The American Economic Review 53(5), 941–973.

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

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

Banerjee, A., E. Duflo, and R. Hornbeck (2014). Bundling health insurance and microfinance in india: There cannot be adverse selection if there is no demand. American Economic Review 104(5), 291–296.

Barseghyan, L., F. Molinari, T. O'Donoghue, and J. C. Teitelbaum (2013). The nature of risk preferences: Evidence from insurance choices. American Economic Review 103(6), 2499–2529.

Barseghyan, L., J. Prince, and J. C. Teitelbaum (2011). Are risk preferences stable across contexts? evidence from insurance data. American Economic Review 101(2), 591–631.

Benmelech, E. and T. Moskovitz (2010). The political economy of financial regulation: Evidence from u.s. state usury laws in the 19th century. The Journal of Finance 65, 1029–1073.

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

Bhutta, N. and B. J. Keys (2022, 02). Moral hazard during the housing boom: Evidence from private mortgage insurance. The Review of Financial Studies 35(2), 771–813.

Biswas, S., M. Hossain, and D. Zink (2023). California wildfires, property damage, and mortgage repayments. Federal Reserve Bank of Philadelphia Working Paper No. 23-05.

Blonz, J., M. Hossain, B. Keys, P. Mulder, and J. Weill (2024). Pricing protection: Credit scores, disaster risk, and home insurance affordability. SSRN Working Paper Number 5026600.

Boomhower, J., M. Fowlie, and A. J. Plantinga (2023, May). Wildfire insurance, information, and self-protection. AEA Papers and Proceedings 113, 310–15.

Braverman, A. and J. E. Stiglitz (1982). Sharecropping and the interlinking of agrarian markets. The American Economic Review 72(4), 695–715.

Calem, P., J. Kenney, L. Lambie-Hanson, and L. Nakamura (2021). Appraising home purchase appraisals. Real Estate Economics 49(S1), 134–168.

Casaburi, L. and J. Willis (2018, December). Time versus state in insurance: Experimental evidence from contract farming in kenya. American Economic Review 108(12), 3778–3813.

Cherry, S., E. Jiang, G. Matvos, T. Piskorski, and A. Seru. (2022). Shadow bank distress and household debt relief: Evidence from the cares act. AEA Papers and Proceedings 112.

Collier, B. L., T. Huber, Jaspersen, and J. G. (2024). Homeowners' willingness to hedge flood risks as prices increase. SSRN Working paper N 4635177.

Cookson, J. A., E. Gallagher, and P. Mulder (2023). Money to burn: Crowdfunding wildfire recovery. SSRN Working Paper N 4535190.

Defusco, A. A., S. Johnson, and J. Mondragon (2019, August). Regulating Household Leverage. The Review of Economic Studies, rdz040.

Defusco, A. A. and J. Mondragon (2020). No job, no money, no refi: Frictions to refinancing in a recession. The Journal of Finance 75(5), 2327–2376.

del Valle, A., T. Scharlemann, and S. Shore (2024). Household financial decision-making after natural disasters: Evidence from hurricane harvey. Journal of Financial and Quantitative Analysis 59(5), 2459–2485.

Einav, L. and A. Finkelstein (2023). Empirical analyses of selection and welfare in insurance markets: a self-indulgent survey. The Geneva Risk and Insurance Review.

Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating welfare in insurance markets using variation in prices. The Quarterly Journal of Economics 125.3.

Ericson, K. M. and J. R. Sydnor (2018). Liquidity constraints and the value of insurance. NBER Working Paper No. w24993.

Finkelstein, A. and J. Poterba (2014). Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market. Journal of Political Economy 112(1), 183–208.

Froot, K. A. and P. G. O'Connell (1999). The pricing of U.S. catastrophe reinsurance. In K. A. Froot (Ed.), The Financing of Catastrophe Risk, Chapter 5, pp. 195–232. University of Chicago Press.

Froot, K. A., P. G. O'connell, et al. (1999). The pricing of us catastrophe reinsurance. The financing of catastrophe risk, 195–232.

Froot, K. A., D. S. Scharfstein, and J. C. Stein (1993, December). Risk management: Coordinating corporate investment and financing policies. Journal of Finance 48(5), 1629–1658.

Gallagher, J. and D. Hartley (2017, August). Household Finance after a Natural Disaster: The Case of Hurricane Katrina. American Economic Journal: Economic Policy 9(3), 199–228.

Ganong, P. and P. Noel (2020, October). Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession. American Economic Review 110(10), 3100–3138.

Ge, S., S. Johnson, and N. Tzur-Ilan (2024). Climate risk, insurance premiums, and the effects on mortgages.

Gennaioli, N., R. L. Porta, F. Lopez-De-Silanes, and A. Shleifer (2021). Trust and insurance contracts. The Review of Financial Studies [Forthcoming].

Glaeser, E. and J. Gyourko (2005, 04). Urban Decline and Durable Housing. Journal of Political Economy.

Goldsmith-Pinkham, P., M. T. Gustafson, R. C. Lewis, and M. Schwert (2020). Sea level rise exposure and municipal bond yields.

Hausman, J. A. (1996). Valuation of new goods under perfect and imperfect competition. Economics of New Goods, ed. by T. F. Bresnahan and R. J. Gordon.

Hurst, E., B. J. Keys, A. Seru, and J. Vavra (2016). Regional redistribution through the US mortgage market. American Economic Review 106(10), 2982–3028.

Issler, P., R. H. Stanton, C. Vergara-Alert, and N. E. Wallace (2019). Mortgage Markets with Climate-Change Risk: Evidence from Wildfires in California. SSRN Working Paper No. 3511843.

Issler, P., R. H. Stanton, C. Vergara-Alert, and N. E. Wallace (2020). Mortgage markets with climate-change risk: evidence from wildfires in California.

Jaffee, D. M. and T. Russell (1997). Catastrophe insurance, capital markets, and uninsurable risks. The Journal of Risk and Insurance 64(2), 205–230.

Jeziorski, W., S. Ramnath, et al. (2021). Homeowners insurance and climate change. Chicago Fed Letter (460), 1–6.

Jiang, E. X. and A. L. Zhang (2025). Collateral value uncertainty and mortgage credit provision. Journal of Financial Economics 169, 104054.

Jotikasthira, C., A. Kartasheva, C. T. Lundblad, and T. Ramadorai (2025, January). Strategic claim payment delays? evidence from property and casualty insurance. Swiss Finance Institute Research Paper No. 25-14, SMU Cox School of Business Research Paper No. 25-4.

Jung, H. and J. K. Jung (2025). Economics of property insurance.

Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014). Agricultural decisions after relaxing credit and risk constraints. Quarterly Journal of Economics 129(2), 597–652.

Kim, Y. S., D. Lee, T. Scharlemann, and J. Vickery (2024). Intermediation frictions in debt relief: Evidence from cares act forbearance. Journal of Financial Economics 158, 103873.

Klein, K. (2018). Minding the protection gap: Resoling unintended, pervasive, profound homeowner underinsurance. University of Connecticut School of Law 25, 35–116.

Koijen, R. S. and M. Yogo (2015). The cost of financial frictions for life insurers. American Economic Review 105(1), 445–475.

Koijen, R. S. and M. Yogo (2016). Shadow Insurance. Econometrica 84(3), 1265–1287.

Kousky, C., M. Palim, and Y. Pan (2020). Flood Damage and Mortgage Credit Risk: A Case Study of Hurricane Harvey. Journal of Housing Research 29(sup1), S86–S120.

Kruttli, M., B. Tran, and S. Watugala (2019). Pricing Poseidon: extreme weather uncertainty and firm return dynamics.

Liao, Y. and P. Mulder (2021, August). What's at Stake? Understanding the Role of Home Equity in Flood Insurance Demand. Working Paper, 51.

Mahoney, N. (2015). Bankruptcy as implicit health insurance. American Economic Review 105(2), 710–46.

Meisenzahl, R. R., A. Polacek, S. Ramnath, and Z. Yudhishthu (2026, January). Measuring the impact of property insurance premiums on the mortgage market. Working Paper WP 2026-01, Federal Reserve Bank of Chicago.

Mossin, J. (1968). Aspects of rational insurance purchasing. Journal of Political Economy 76(4,Part 1), 553–568.

Mulder, P. and B. Keys (2024). Property insurance and disaster risk: New evidence from mortgage escrow data. NBER Working Paper 32579.

Murfin, J. and M. Spiegel (2020). Is the risk of sea level rise capitalized in residential real estate? The Review of Financial Studies 33(3), 1217–1255.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. Econometrica, 307–42.

Oh, S. S., I. Sen, and A.-M. Tenekedjieva (2021). Pricing of climate risk insurance: Regulation and cross-subsidies. Journal of Finance.

Rampini, A. A. and S. Viswanathan (2013, August). Collateral and capital structure. Journal of Financial Economics 109(2), 466–492.

Santos, J. and K. Blickle (2022, April). Unintended consequences of "mandatory" flood insurance. FRBNY Staff Report No. 1012.

Sastry, P. (2025, December). Who bears flood risk? evidence from mortgage markets in florida. Review of Finance (forthcoming).

Sastry, P. and I. Sen (2025). Perspectives on insurance and climate risk. Journal of Finance: Insights and Perspectives, Forthcoming.

Sastry, P., I. Sen, and A.-M. Tenekedjieva (2024). When insurers exit: Climate losses, fragile insurers, and mortgage markets.

Shleifer, A. and R. W. Vishny (1992). Liquidation values and debt capacity: A market equilibrium approach. The Journal of Finance 47(4), 1343–1366.

Solomon, A. (2025, August). Catastrophic climate risk and the limits of private markets. Working Paper, New York University Stern School of Business.

Wagner, K. R. (2022). Adaptation and adverse selection in markets for natural disaster insurance. American Economic Journal: Economic Policy 14(3), 380–421.

A. Additional Details on Data

This appendix provides additional details on the data sources, the process to construct the final sample, and high-level analysis and descriptions of key variables.

A.1. Data Construction

Our main dataset is the ICE-McDash mortgage and insurance modules, which combine data on both mortgages and homeowners insurance for mortgage borrowers whose insurance payments are escrowed. While the ICE-McDash (hereafter "McDash") mortgage data are commonly used in academic research, the insurance module is a newly available dataset that was only recently made available starting in 2024.

Sample representativeness. McDash collects homeowners insurance information for a subset of mortgages, based on whether specific servicers have agreed to report insurance contract details to ICE-McDash and whether borrowers escrow insurance. To ensure that our sample is representative of the broader McDash universe of mortgages, we compare borrower incomes, loan amounts, back-end debt-to-income ratios, loan-to-value ratios, and credit scores for loans in the insurance module with loans in the overall McDash samples. Table C.1 shows that the loans in the McDash insurance module are broadly representative of the loans in McDash's mortgage origination module across a range of borrower characteristics. For example, average credit score is 735 in our insurance module vs. 737 in the overall McDash; LTV ratio is 86 (insurance) vs. 85 (overall); DTI ratio is 35.8 (insurance) vs. 35.1 (overall), loan amount is $243K (insurance) vs. $257K (overall), and annual income is $55K (insurance) vs. $51K (overall). This makes us comfortable that the insurance module is broadly representative of the McDash mortgage sample.

Key details of database merges. Below we provide further details. We start with merging the McDash insurance and mortgage modules. Next, we merge in Credit Risk Insights McDash (CRISM), which includes Equifax's Personal Income Model. We retain the credit profile of the primary borrower, only for matches with over a 80% confidence rate, as determined by McDash.

Next, we merge CoreLogic Deeds data using a fuzzy matching algorithm based on common variables available in both datasets (namely the zip code, origination date, loan amount, loan purpose (refinance/purchase), and loan type (i.e. FHA, VA, etc)). These merge fields generate about a 50% match rate and have a duplication rate of under 5%. We drop duplicate matches. We additionally validated this match algorithm using the Federal Reserve's FR Y-14 data, which allows for an exact address match. We find that unique matches using our match algorithm were accurate in over 95% of cases. A match with the mortgage deeds using Y14 data also has about the same match rate as the McDash data, suggesting that mortgage deeds records-keeping - and not our match algorithm - is the main driver of sample attrition.

Using the CoreLogic unique property identifier, we next merge in property-level tax and climate risk data. The tax data cover property characteristics used for tax assessments, including the

structure size, among other things. With these core merges, we have, for each loan, the mortgage and insurance contract, borrower characteristics, and property characteristics. Because for years 2013 to 2023, insurance data are only updated in December, we use the insurance information as of the first December after origination date to provide us a picture of the key insurance contract terms at origination. We also retain mortgage performance data only for December for subsequent observations on the same loan, to remain consistent with insurance data updates.

Finally, we bring in data on construction costs at the location and year levels. The final merged sample consists of roughly 75 million observations, which include both information as of mortgage origination and subsequent information on mortgage performance at a monthly frequency. We limit the merged sample to single-family home *purchase mortgages* originated on or after 2013, dropping loans associated with re-finances. As described in Section 1, we also limit the main sample to the time of mortgage origination.

A.2. Measurement of Coverage

The coverage variable we observe in our data refers to basic policy limits for the structure. Some homeowners may choose to add-on extended coverage, but they can only do so if their basic coverage is at least equal to 100% of their replacement cost. For example, Klein (2018) writes that extended coverage cannot be provided unless "the basic coverage is at least as great as the estimated replacement cost of the property." Similarly, the insurance consulting firm *Marsh McLennan* also writes "When you insure your home to 100% of its replacement cost value, some insurance companies will offer the benefit of extended replacement cost. This provision will pay beyond your policy limit should the amount at the time of loss not be adequate," available here. Extended benefits are not intended to substitute for basic coverage choices. Therefore, we argue that our measure of coverage accurately captures underinsurance.

To support our claim, we independently verify our estimates with external sources. A number of independent sources also find that the average household is under-insured by nearly 20-30%, matching our estimates. For example, Klein (2018) claims that while prior to 1990 most insurance contracts guaranteed full coverage by default, in the last 30 years the market has moved towards wide-spread under-insurance. Similarly, Nationwide Insurance estimated that households were on average under-insured by 22%, and that two-thirds of Americans are under-insured (see Nationwide's estimates quoted in PropertyCasualty360 (2021)). Similarly, CoreLogic also estimates that nearly two-thirds of Americans are under-insured, and that they are under-insured by an average of 27% (see CoreLogic quoted in Newbrook Insurance Agency (2022)).

A.3. Broad Data Description

Given the novelty of our data, we provide a broad description of the sample, looking at trends over time, correlations across space, and explanatory power of different variables and characteristics.

A.3.1. Summary Statistics

We start by showing significant variation in coverage ratios in the cross-section. Figure C.1 plots a histogram of coverage ratios over the sample. The vast majority of the households have coverage ratios below 1, i.e. their coverage limits are below the rebuilding cost of their homes. Table 1 shows that the average coverage ratio is around 70%, and that households with one standard deviation less coverage have a coverage ratio of about 50%.

Next, we examine the spatial distribution of coverage ratios in Figure C.14, along with the distribution of income and climate risk. Two patters emerge. First, the states in the top quartile by income (between $56,000 and 85,000 per year in Figure C.14c) are also states with the highest coverage ratios (above 73% from Figure C.14a). Second, states with the highest climate risk are *not* the ones with the most coverage. For example, Texas is one of the riskiest states, with 24% of its properties classified as high risk. Yet its average coverage ratio is only 62%, putting it in the bottom quartile. Similarly, Maine is very low climate risk state, with less than 2% of its properties being high risk – yet it has among the highest average coverage ratios in the country, at 82%. At the state level, income and coverage ratios have a 77% correlation. In contrast, climate risk and coverage ratio have a -18% correlation. These broad patterns suggest that household income plays an important role in coverage choices, with higher income states having on average higher coverage ratios.

A.3.2. Explaining the Variation in Insurance Prices

In Figure C.7 we examine the drivers of variation in insurance prices. The most predictive variable category is zip-year fixed effects, explaining close to 65% of the variation in prices. Adding housing, climate, mortgage, or insurance variables increases R^2 by less than 10%.

A.3.3. Coverage Ratio and Loan Age

We next examine how coverage ratios evolve as a loan ages over time. In Figure C.2 we plot a binned scatter plot of coverage ratios by loan age, after controlling for a loan fixed effect. We find limited evidence that a borrower significantly adjusts coverage ratios over the life of the loan. There are two reasons why this result is important. First, as coverage ratios are flat throughout the life of the loan implies that household under-insurance begins at origination and is persistent. Therefore, in our main analysis, we focus on insurance choices at mortgage origination. Second, a flat coverage ratio also points to limited borrower strategic behavior. A natural model of strategic behavior would suggest that households' incentives to insure *increases* as the loan ages, all else equal (Liao and Mulder, 2021), as households build up equity in the house over the life of the loan. Therefore, this pattern suggests that under-insurance is not likely to be driven by strategic behavior by indebted households having limited incentives to hedge.

A.4. Supplementary Data

FIO Data on Homeowners Insurance: The U.S. Treasury Department's Federal Insurance Office (FIO) in conjunction with the NAIC conducted a data call where they collect homeowners insurance data on premiums and claims from the largest 400 insurance companies, who represent

over 80% of the property insurance market by premium volume. The data was collected at a ZIP code - year level, spanning 2018-2022. While the insurer submissions remain confidential, but the FIO released averages at the zip-level with information on claim probabilities, claim severity, loss ratios, and premiums. We use these estimates to validate the instrument in Section 4

B. Coverage and Household Financial Vulnerability

B.1. Descriptive Facts

The protection gap is strongly positively correlated with measures of financial vulnerability. At the state level, income and coverage ratios have a 77% correlation (Figure C.14c). For example, California and the Northeast exhibit high coverage ratios, whereas the Southeast and Midwest have substantially lower coverage. Figure C.15 shows similar patterns hold at the borrower level using measures of income and credit scores. Going from the lowest ventile of credit scores (around 650) to the highest ventile of credit scores (around 800) is associated with a 10 percentage point increase in coverage ratios, from 64% to 74%. The relationship is quite linear. Figure C.15b shows a similar pattern for primary borrower income. Here, the cross-sectional differences are even starker: going from the lowest ventile of income (around $20,000) to the highest ventile (around $120,000) is associated with a 30 percentage point increase in coverage ratios, from 60% to 90%. These broad patterns suggest that household income plays an important role in coverage choices, with higher income states having, on average, higher coverage ratios.

In contrast, the protection gap has a weaker relationship with measures of climate risk, and seems to vary with the income level of the household. At the state level, climate risk and coverage have a correlation of –18%. At the borrower level, households facing higher climate risk tend to obtain more coverage and are less likely to be severely under-insured (Figure C.15). In contrast, lower income and lower credit score households exhibit different behavior: climate risk appears to have little influence on coverage choices for this group. For example, households with annual primary borrower income below $40,000 show similar rates of severe under-insurance regardless of climate risk exposure; at the very lowest incomes (around $20,000), households in high-risk areas are more likely to be severely under-insured than those in low-risk areas. Similarly, borrowers with low credit scores (around 650) in high-risk areas tend to have less coverage than their counterparts in low-risk areas. In summary, lower-income households in high climate risk areas are particularly vulnerable and most likely to be severely under-insured. These patterns suggest that financial considerations can outweigh risk considerations in determining insurance coverage, consistent with the broad state-level patterns in Figure C.14.

These results suggest that under-insurance is particularly high among the lowest-income households. Such households are also far more likely to be vulnerable to large shocks, lacking the financial resources of their own to smooth such shocks.

B.2. Within Person Income Changes

Within-Person Income Changes. We now consider the dynamics of household demand elasticities, and how they change as the characteristics of households evolve. In particular, we focus on within-household changes in income. We start with the sample of low income people at origination, and define an income transition as a switch from low income to high income within the first 3 years (column 1) and first 5 years (column 2). We then consider how their demand elasticities

change when income increases shortly after mortgage origination for the same borrower. The key identifying assumption is that other drivers of demand (e.g., preferences and beliefs) remain fixed over short horizons. We argue that a benefit of this specification is that other largely fundamental characteristics of a household, such as risk preferences, behavioral biases or beliefs, are unlikely to change so quickly after mortgage origination, suggesting the impact of income.

Table C.15 shows the results. We find that the same individual becomes 20% more inelastic when they transition to high income, even controlling for other contemporaneous changes in household characteristics (such as contemporaneous credit scores). We acknowledge these are not exogenous income shocks, but we find this highly suggestive that broader wealth/income plays a role to play in insurance elasticities. We also argue that lender incentives to monitor is not likely to *increase* with positive income shocks for the same borrower. If anything, monitoring incentives are likely to reduce. Therefore, this test plausibly holds fixed lender constraints, allowing us to uncover changes in household demand.

C. Additional Tables and Figures

Figure C.1: Histogram of Coverage Ratio

This figure shows a histogram of coverage ratios, defined as insurance coverage amount divided by the estimated replacement cost. The insurance coverage data are from ICE-McDash. The figure is based on the main analysis sample, outlined in Section 1.

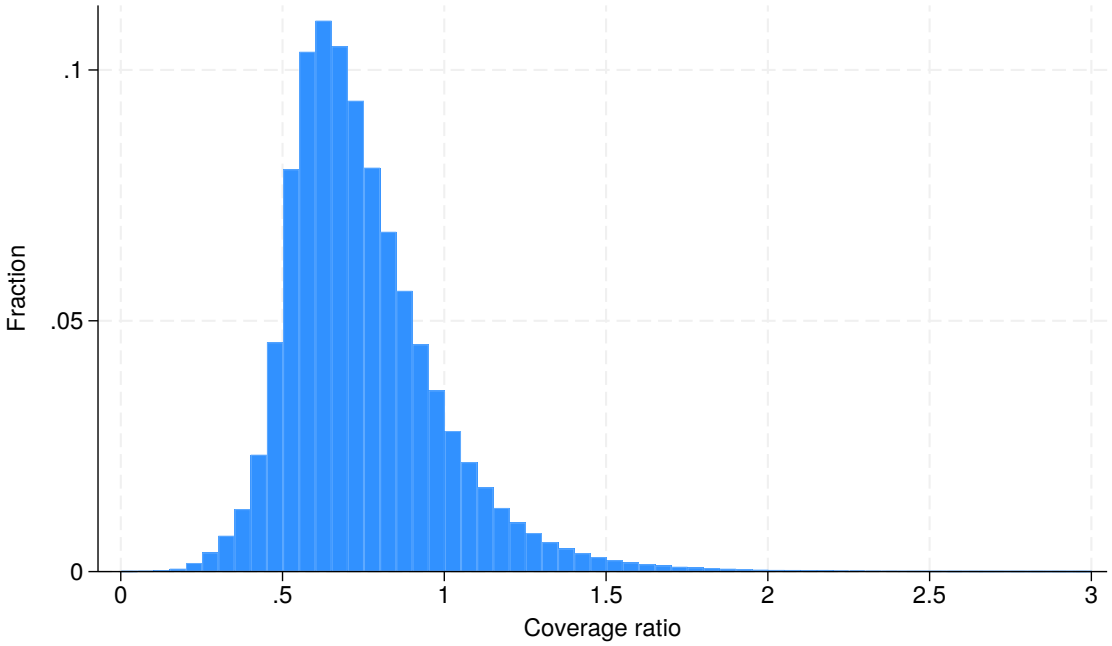

Figure C.2: Coverage Ratios by Loan Age

This figure shows a binned scatter plot of how coverage ratios vary by the age of the mortgage, after controlling for a loan fixed effect. Insurance coverage ratios are defined as defined as insurance coverage amount divided by the estimated rebuild cost. The sample of loans included are based on the main analysis sample, outlined in Section 1, which we observe through time to assess changes in coverage ratio.

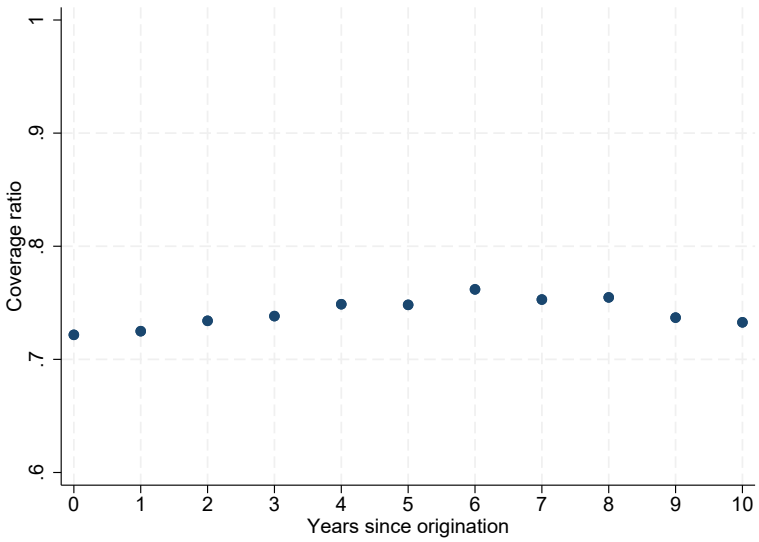

Figure C.3: Rebuild cost and loan amount gradients by loan to replacement cost ratio

This figure shows a binscatter plot of coverage amount against rebuild cost (panel (a)) and loan amount (panel (b)). We report plots for different values of loan to replacement cost ratio: < 1 (red line) and $>= 1$ (blue line). The data are at the time of mortgage origination, and spans the period from 2013 to 2023.

(a) Rebuild cost

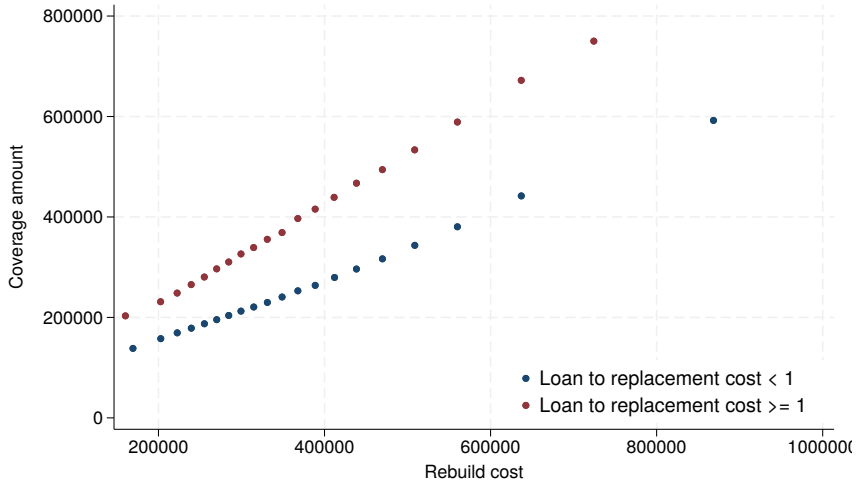

(b) Loan amount

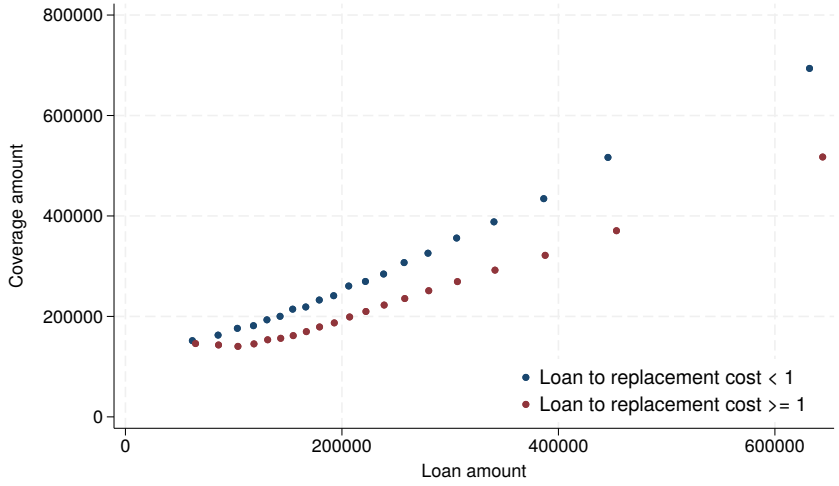

Figure C.4: Histogram of Loan-to-Replacement Cost

This figure shows a histogram of the ratio of loan amounts to replacement costs at origination. The data are at the time of mortgage origination, and spans the period from 2013 to 2023.

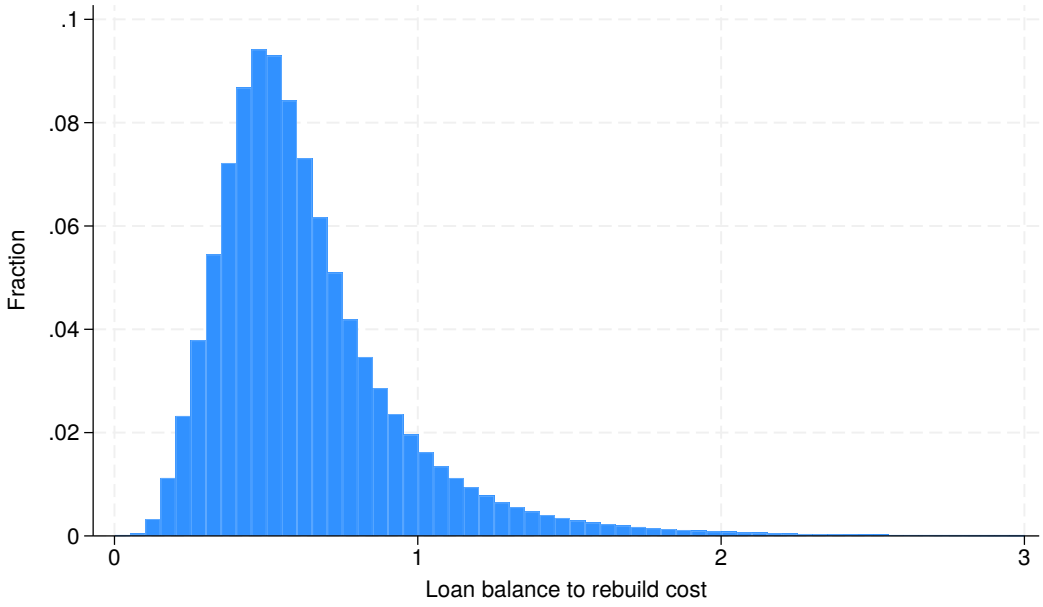

Figure C.5: Histogram of Coverage to market values

This figure shows a histogram of the ratio of coverage to house price at origination. The data are at the time of mortgage origination, and spans the period from 2013 to 2023.

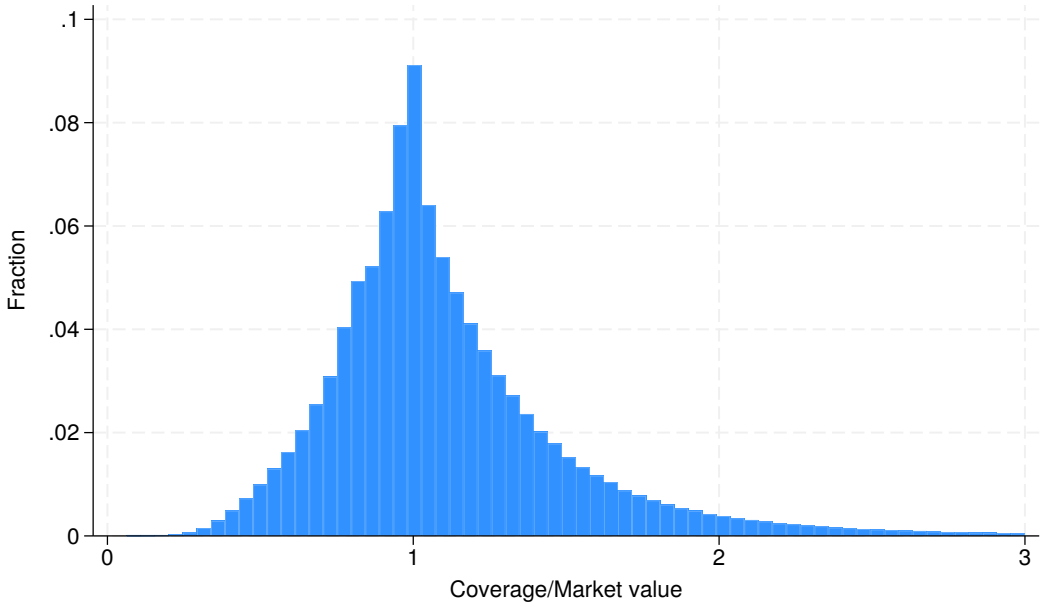

Figure C.6: Coverage to loan balance vs. Loan-to-Replacement Cost

This figure shows a binscatter of how the ratio of coverage to loan balance varies by loan to replacement cost. The data are at the time of mortgage origination, and spans the period from 2013 to 2023.

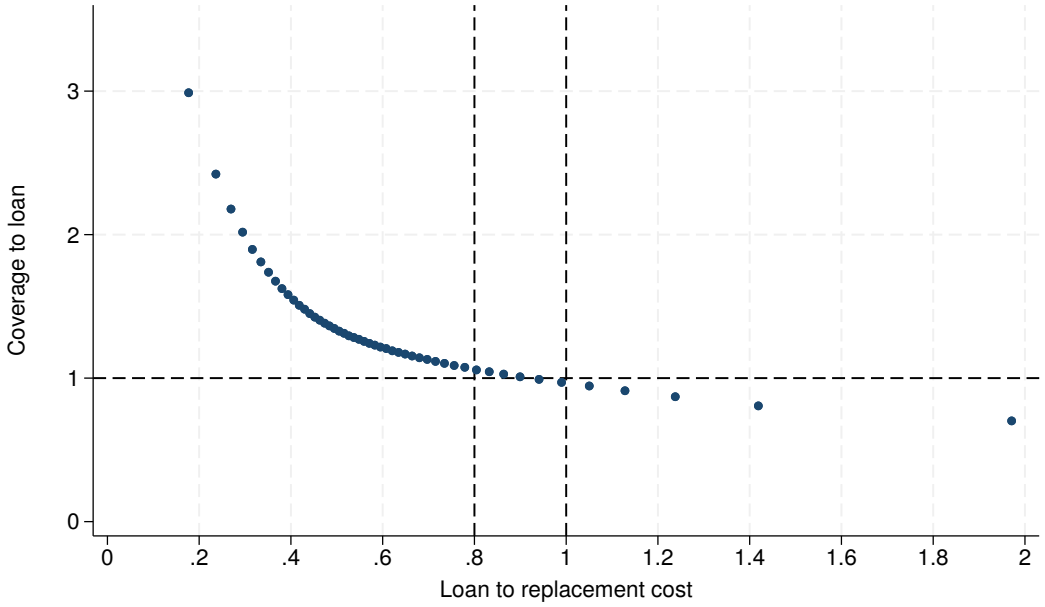

Figure C.7: Explaining the Variation in Insurance Prices

The bar chart shows the R^2 from a series of regressions with thematic subsets of control variables explaining log insurance coverage levels (panel a) and log insurance prices (monthly premium per $100,000 of coverage) (panel b). The variables are (i) ZIP-year fixed effects; (ii) "Housing", includes rebuild cost; (iii) "Climate", includes indicators for properties in top decile of climate risk, and if the house is in a flood plain; (iv) "Mortgage", includes LTV, origination credit score (from ICE-McDash), original principal balance, investor type (FHA, GSE, other) etc.; (v) "Insurance" includes log insurance price and deductible amount for panel (a) and includes log coverage and deductible amount for panel (b). The figure is based on the main analysis sample, outlined in Section 1.

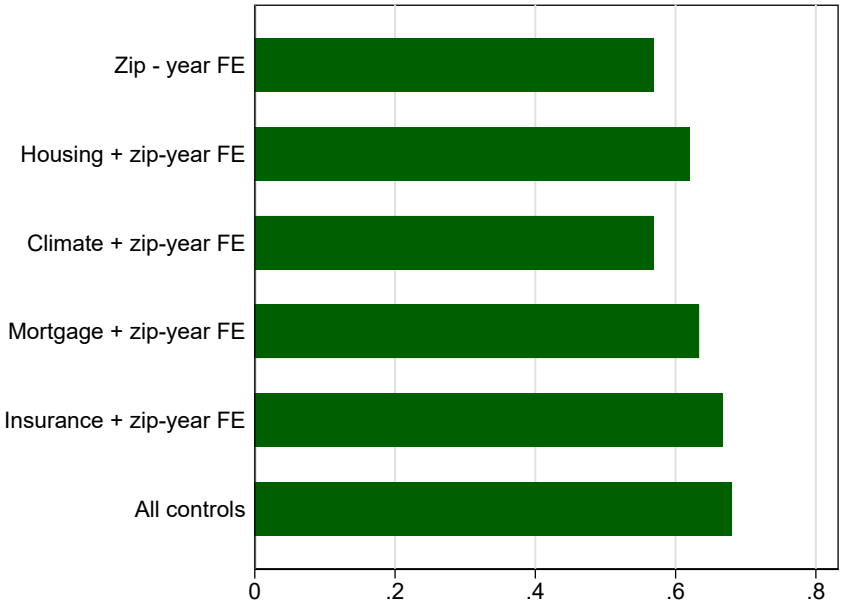

Figure C.8: Deductible Choices

This figure shows deductible choices. Panel (a) reports a histogram of deductible amounts. Panel (b) reports how deductible amounts per dollar of coverage has evolved over time. We report mean, 10th, and 90th percentile values.

(a) Deductible amount

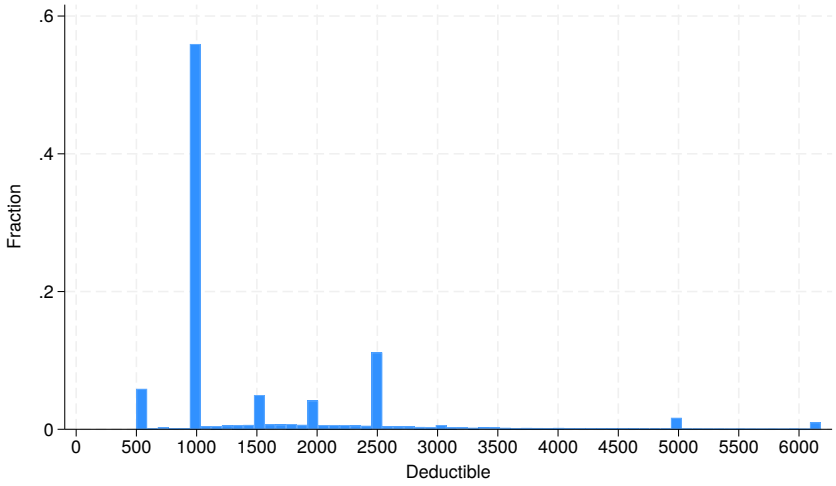

(b) Deductible amount per dollar of coverage

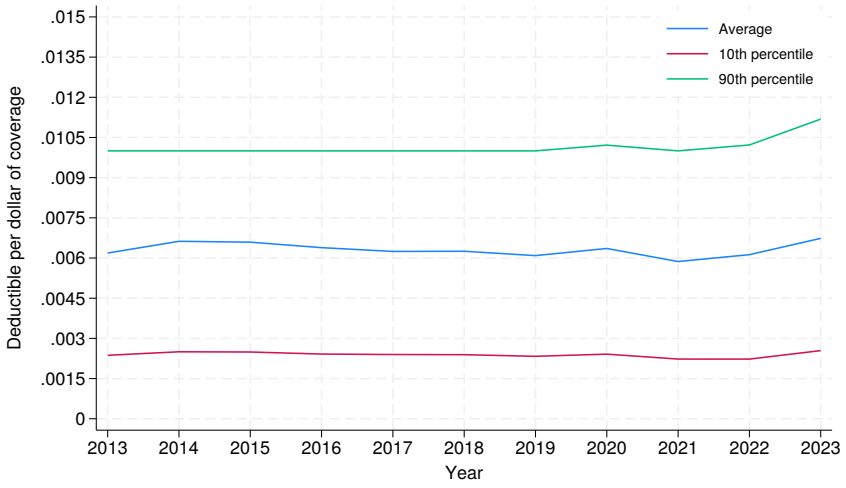

Figure C.9: Kink Design: Coverage sensitivity to loan amount and replacement cost (Robustness)

This figure shows a coefficient plot and 95% confidence intervals obtained from estimating Equation 3 in the text, but limiting controls to ZIP × Year fixed effects only. The data are at the time of mortgage origination, and spans the period from 2013 to 2023. Standard errors are clustered at the county-year level.

(a) Loan amount

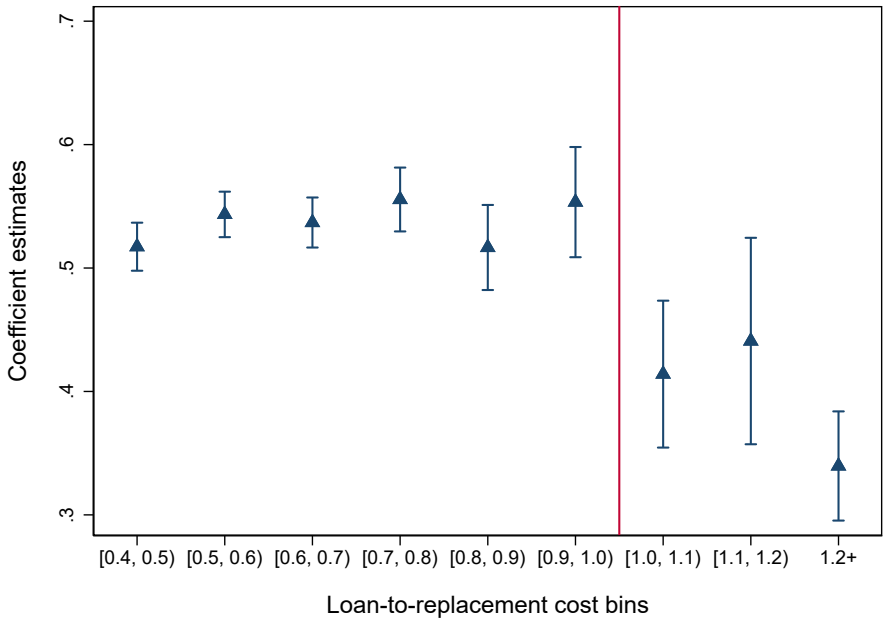

(b) Replacement cost

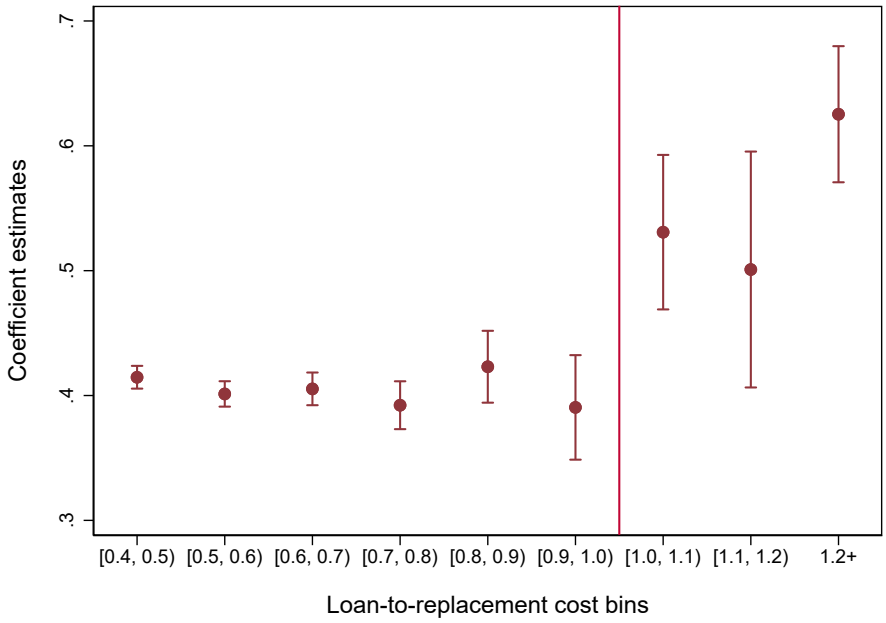

Figure C.10: Actual coverage vs. coverage implied by lender rule

This figure shows a binscatter plot of coverage amount against the required coverage implied by the lender rule: *min*(loan amount, replacement cost). The data are at the time of mortgage origination, and spans the period from 2013 to 2023.

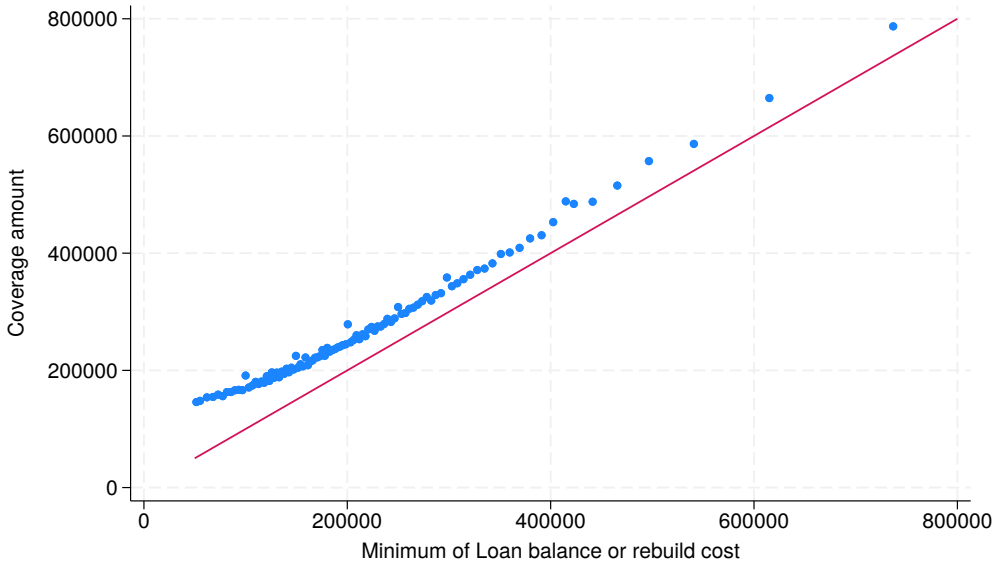

Figure C.11: Insurers' Underwriting Profitability Across State Types

This figure shows estimates of insurers' combined ratios by high and low friction states (panel a) and over time for high friction states (panel b). Combined ratio is a widely used metric to measure insurers' underwriting profitability. We compute the combined ratio as the sum of ZIP code level loss ratio (losses-to-premiums ratio) taken from Treasury's FIO data (see Appendix A), the average annual aggregate expense-to-premium ratio, and the annual loss-adjustment expense-to-premium ratio from the NAIC regulatory filings.

(a) Combined Ratios by State Type

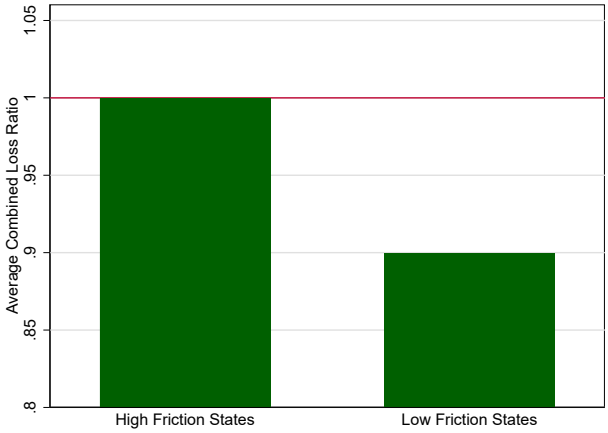

(b) Combined Ratios over time for High Friction States

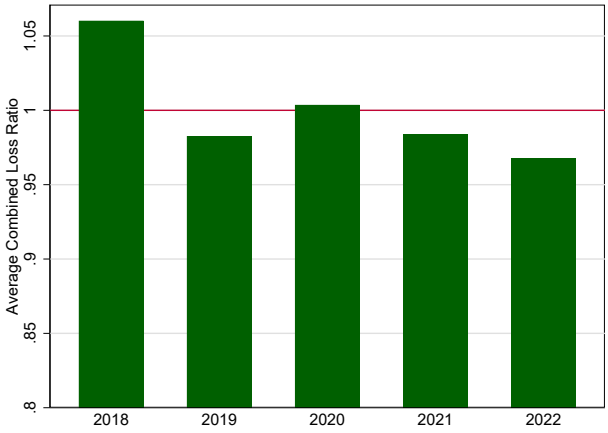

Figure C.12: Price gap between rate targets and received.

This chart uses insurer rate filings across high friction states. This graph plots the relationship between the insurance price gap and the price increases that were granted. The price gap is defined as the difference between the actuarially fair price implied by insurers models (target) and what price change regulators grant (received).

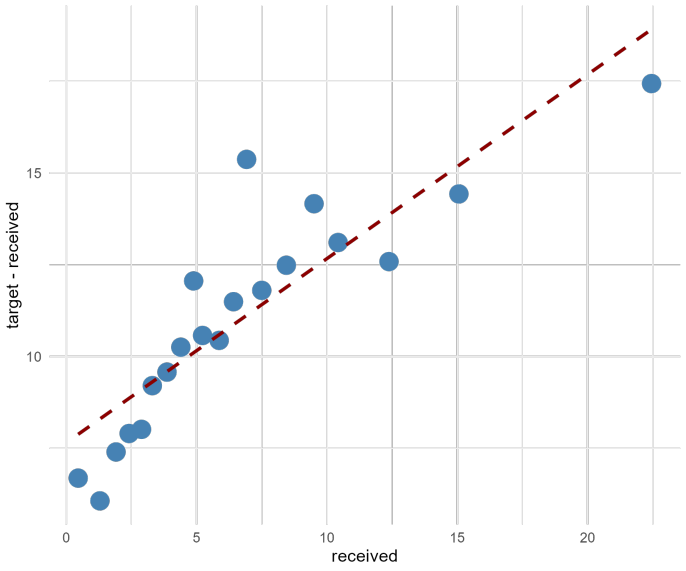

Figure C.13: Case study: Time series of two insurers' prices

We compute the time series of insurance price (premium per $1000 of coverage) between 2015 and 2020, and the concurrent rate changes, for two of the largest five insurers in California. Each dot is the median of the prices the insurer charged for the policies renewed in a given month in the McDash sample. The dashed vertical lines denote the date that each rate change took effect. Rate increases are shown in blue, and rate decreases are shown in red. Dates on the regulatory changes were collected from S&P MI's Rate Filing products.

(a) Insurer 1

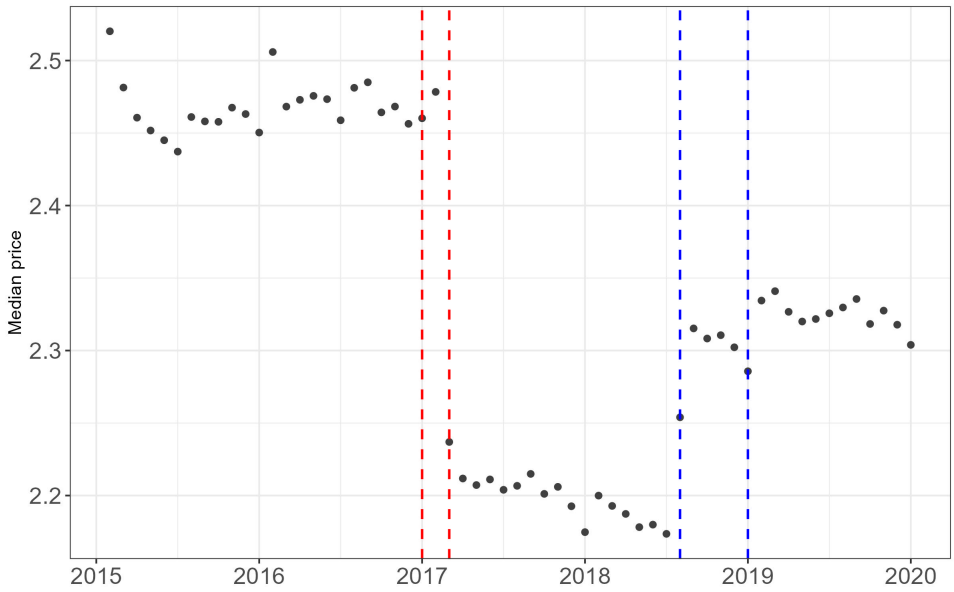

(b) Insurer 2

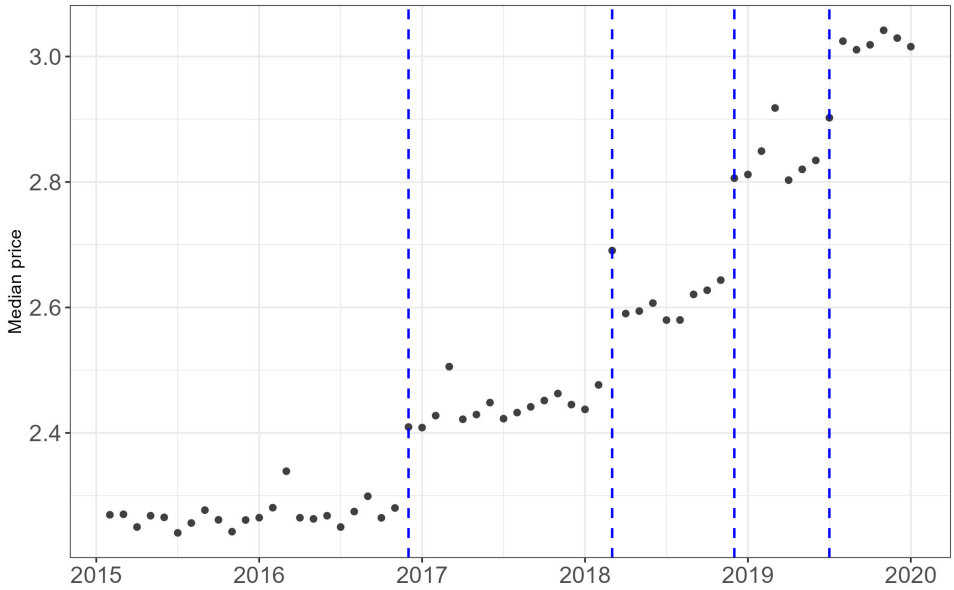

Figure C.14: Spatial Distribution of Coverage, Risk and Income Across the U.S.

The maps show state-level variation in insurance coverage ratios (panel a), climate risk (panel b), and average borrower income (panel c), respectively. Coverage ratio is defined as insurance coverage amount divided by the estimated replacement costs. Insurance coverage data are from ICE-McDash, high climate risk are homes in the top decile by CoreLogic's AAL measure, and average borrower income is from Equifax. The figure is based on the main analysis sample, outlined in Section 1.

(a) Coverage Ratio

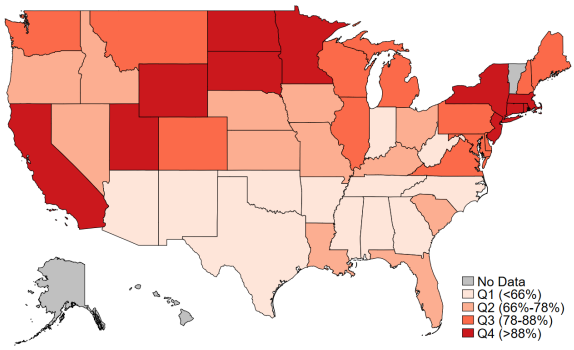

(b) Climate risk

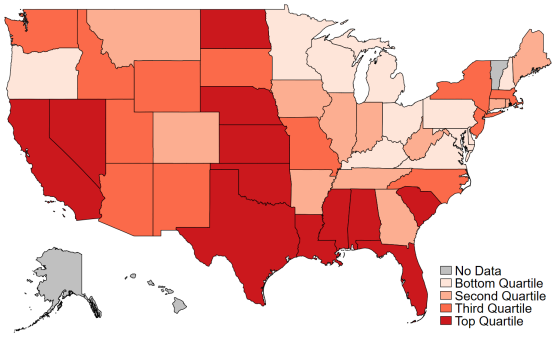

(c) Borrower Income

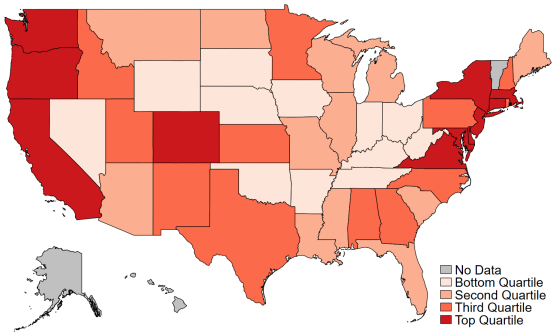

Figure C.15: Coverage Ratios by Borrower Credit Score and Income

This figure shows a binned scatterplot of coverage ratios by borrower credit score (panel a) and income (in thousands) (panel b). The relationship is plotted separately for high and low climate risk properties, where "high climate risk" indicates those in the top decile of climate risk by CoreLogic's AAL measure. Coverage ratio is defined as insurance coverage amount divided by the estimated replacement costs. The figure is based on the main analysis sample, outlined in Section 1.

(a) Borrower Credit Score

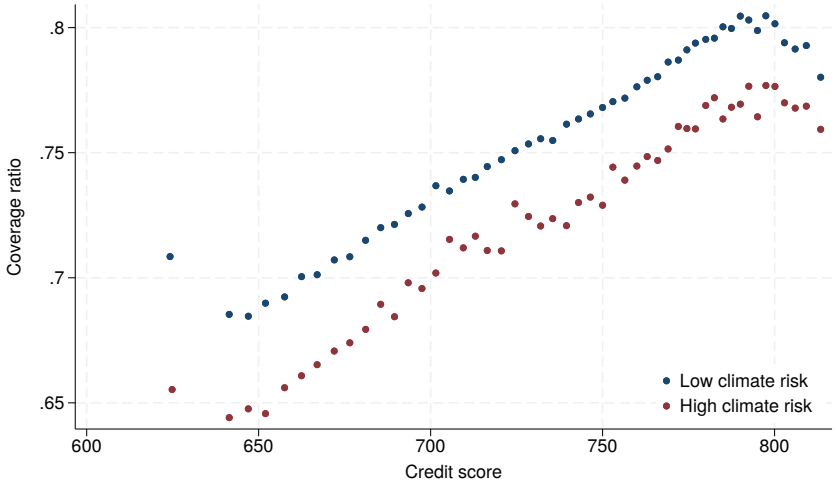

(b) Borrower Income

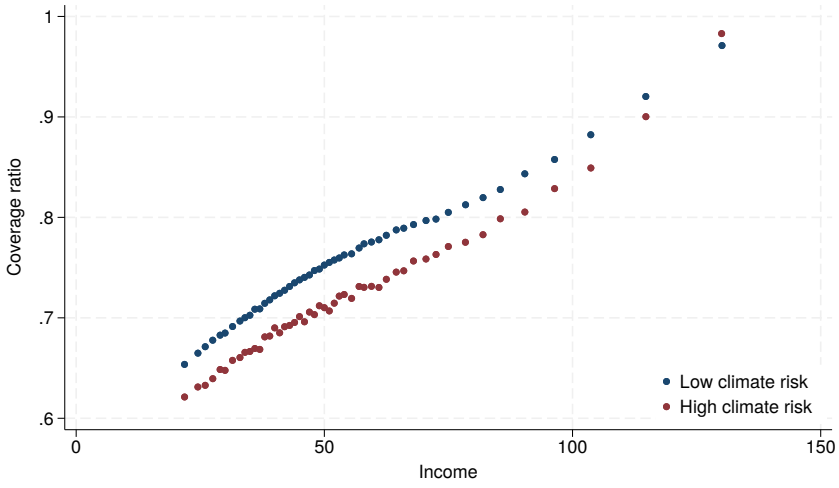

C.1. Appendix Tables

Table C.1: Comparison of McDash Insurance Module to Overall McDash Sample

This table compares summary statistics from a 10% random sample of the ICE McDash Mortgage modules, and from the subset with the Insurance module. Insurance variables are measured at the first observation in the ICE McDash loan performance data, as in our main analysis sample. The sample in this table does not include a property-level merge, so property-level data like climate risk and replacement cost are omitted. All variables are as defined in the table notes in Table 1.

	McDash Insurance Module		Overall McDash Sample	
	(1) Mean	(2) SD	(3) Mean	(4) SD
Coverage Amount	284085.6	153112.7		
Annual Insurance Premium	1109.3	606.3		
Insurance Price	37.1	21.0		
Annual Income	55.0	24.4	50.9	23.2
Loan Amount	243336.3	152177.6	257857.1	178836.7
DTI Ratio	35.8	9.46	35.1	10.0
LTV Ratio	86.3	13.4	84.9	14.4
Credit Score	735.1	51.3	736.9	52.2
Share with high DTI	0.69	0.46	0.65	0.48
Share with high LTV	0.64	0.48	0.59	0.49
N	538093		822914	

Table C.2: Drivers of replacement costs

This table shows shows the relationship between the ratio of replacement costs to house prices at origination with observable property characteristics. Fixed effects and controls included as indicated. Borrower controls refer to DTI, LTV, and credit score. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Ratio of Replacement Cost to Market Value			
	(1)	(2)	(3)	(4)
Building Age ∈ (20Y, 40Y]	0.071***	0.105***	0.104***	0.100***
	(0.003)	(0.002)	(0.002)	(0.002)
Building Age ∈ (40Y, 60Y]	0.141***	0.203***	0.201***	0.194***
	(0.004)	(0.003)	(0.003)	(0.003)
Building Age ∈ (60Y, 80Y]	0.161***	0.253***	0.251***	0.245***
	(0.005)	(0.004)	(0.004)	(0.004)
Building Age > 80Y	0.289***	0.362***	0.360***	0.354***
	(0.007)	(0.007)	(0.007)	(0.006)
ln Building Size (Sq. ft.)		0.210***	0.217***	0.240***
		(0.004)	(0.004)	(0.004)
ZIP-Year FE	Y	Y	Y	Y
Insurer-Year FE	N	N	Y	Y
Borrower Controls	N	N	N	Y
Observations	2,931,118	2,917,677	2,917,567	2,917,567
Adjusted R^2	0.656	0.666	0.669	0.674

Table C.3: Household characteristics around L/RC = 1

This table provides summary statistics for borrower characteristics above and below the $L/RC = 1$ threshold. We limit to borrowers within 5% of the threshold.

	L/RC ∈ [.95,1]		L/RC ∈ (1, 1.05]	
	Mean	SD	Mean	SD
DTI Ratio	36.2	9.22	36.2	9.20
LTV Ratio	87.0	10.6	86.9	10.6
Credit score	743.8	47.6	744.9	47.1
Share of High Climate Risk	0.100	0.30	0.11	0.31
Borrower age	40.1	11.4	40.0	11.2
Building age	43.8	28.9	45.4	29.1
Building (Sq ft)	1920.9	824.9	1907.5	815.0
Price per Sq. ft.	216.8	61.4	228.4	66.3
House Price	411,300.6	212,027.1	428,692.0	217,197.6
Insurance Price	33.77	19.57	33.50	19.57
Observations	60,690		50,335	

Table C.4: Confounds are smooth through the kink

This table reports estimates from the specification in Equation 4 for the indicated outcome variables, using only zip-year fixed effects as controls. The sample is restricted to observations within a 5% bandwidth around $L/RC = 1$. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	$W_{i,z,t}$								
	(1) DTI ratio	(2) FICO	(3) LTV ratio	(4) Borrower age	(5) Price psqft	(6) Building size	(7) Building age	(8) AAL	(9) Insurance price
Loan amount	-0.00	0.00	0.03***	0.15	0.35***	1.07***	1.01	-0.00	-0.02
	(0.008)	(0.040)	(0.009)	(0.508)	(0.040)	(0.311)	(0.830)	(0.000)	(0.011)
$LTRC > 1\times$ Loan amount	0.00	0.06	0.02	-0.02	-0.00	-0.00	-1.42	-0.00	0.02
	(0.012)	(0.057)	(0.013)	(0.841)	(0.067)	(0.475)	(1.260)	(0.000)	(0.016)
Replacement cost	-0.01	0.04	-0.04***	-0.13	-0.34***	4.19***	-0.93	0.00	-0.01
	(0.008)	(0.039)	(0.008)	(0.493)	(0.039)	(0.302)	(0.806)	(0.000)	(0.010)
$LTRC > 1\times$ Replacement cost	-0.00	-0.06	-0.02	0.02	0.01	-0.00	1.40	0.00	-0.02
	(0.012)	(0.058)	(0.013)	(0.846)	(0.067)	(0.477)	(1.270)	(0.000)	(0.016)
ZIP-Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	80378	80378	80378	80378	77566	77566	80378	80378	80378
Adjusted R^2	0.102	0.139	0.394	-0.006	0.643	0.874	0.819	0.588	0.678

Table C.5: Are households anchoring to house price?

This table provides two placebo tests to understand whether the sensitivity of coverage to loan amount and replacement cost may be driven by house prices. The first considers the pass-through of house prices to coverage around the underwriting threshold of loan balance equal to replacement cost ($L/RC = 1$) from the defacto rule. The second considers the pass-through of house price to coverage around a placebo threshold where market value (MV) equals replacement cost ($MV/RC = 1$). Controls and fixed effects are included as indicated and defined in Table 2. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Coverage$_{i,z,t}$	
	Kink: $LTRC = 1$	Kink: $MV/RC = 1$
	(1)	(2)
House price	0.309***	0.429***
	(0.0117)	(0.0534)
$LTRC > 1\times$ House price	-0.0166	
	(0.0115)	
Replacement cost	0.594***	0.413***
	(0.0143)	(0.0517)
$LTRC > 1\times$ Replacement cost	0.0338*	
	(0.0136)	
$HPRC > 1\times$ House price		0.0177
		(0.0777)
$HPRC > 1\times$ Replacement cost		-0.0162
		(0.0776)
Borrower-Year FE	Y	Y
Structure Controls	Y	Y
ZIP-Year FE	Y	Y
Insurer-Year FE	Y	Y
Observations	98,396	146,410
Adjusted R^2	0.841	0.856

Table C.6: Changes in Pass-through to Coverage at Mortgage Rule Kink: Robustness to Bunching Households

This table shows the results from estimating a regression kink deign as shown in Equation 4 where we consider how the sensitivity of coverage amount to loan amount and replacement cost changes at the loan to replacement cost (LTRC) threshold of 1. We consider the most stringent specification from Table 2 with the most controls and fixed effects but vary the bandwidth limit to include loans within 15% (column 1), 10% (column 2), and 5% (column 3) of the $LTRC = 1$ threshold. All variables and fixed effects are as defined in Table 2, except our sample excludes borrowers that bunch at the loan balance, where bunching is defined as in Table 7. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	$\text{Coverage}_{i,z,t}$		
	(1) ± 15%	(2) ± 10%	(3) ± 5%
Loan amount	0.580***	0.605***	0.696***
	(0.014)	(0.027)	(0.089)
$LTRC > 1\times$ Replacement Cost	-0.120***	-0.134***	-0.285**
	(0.022)	(0.044)	(0.127)
Replacement Cost	0.383***	0.362***	0.278***
	(0.012)	(0.026)	(0.087)
$LTRC > 1\times$ Replacement Cost	0.122***	0.134***	0.286**
	(0.023)	(0.044)	(0.127)
Borrower-Year FE	Y	Y	Y
Structure Controls	Y	Y	Y
ZIP-Year FE	Y	Y	Y
Insurer-Year FE	Y	Y	Y
Observations	282,007	176,762	80,060
Adjusted R^2	0.823	0.823	0.825

Table C.7: Household Coverage Ratio: Lender Underwriting Rules vs. Borrower Characteristics

This table reports estimates from linear regressions of borrowers' coverage ratios (C/RC) on their loan to replacement cost ratio (L/RC), defined as the loan amount divided by the First Street property level replacement cost estimate. Column (1) includes no controls or fixed effects; subsequent columns sequentially add property and borrower characteristics as well as fixed effects as indicated in the table. Borrower controls comprise indicator variables for debt to income (DTI), origination credit score (from ICE-McDash), and loan to value (LTV) buckets, along with a high climate risk indicator based on the CoreLogic risk measure. Property controls include building age, building size, and price per square foot. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Coverage Ratio							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
L/RC	0.549***	0.605***	0.617***	0.618***	0.624***	0.622***	0.618***	0.626***
	(0.0100)	(0.0206)	(0.0252)	(0.0252)	(0.0252)	(0.0252)	(0.0252)	(0.0252)
Medium DTI					-0.0195***			-0.0180***
					(0.000703)			(0.000665)
High DTI					-0.0247***			-0.0223***
					(0.000814)			(0.000751)
Very High DTI					-0.0297***			-0.0263***
					(0.000879)			(0.000788)
Medium FICO						0.00404***		0.00318***
						(0.000448)		(0.000446)
High FICO						0.0102***		0.00848***
						(0.000520)		(0.000505)
Very High FICO						0.0275***		0.0233***
						(0.000863)		(0.000770)
High Climate Risk							-0.00281***	-0.00238***
							(0.000645)	(0.000645)
Constant	0.404***	0.416***	0.360***	0.370***	0.390***	0.357***	0.370***	0.377***
	(0.00600)	(0.00572)	(0.00666)	(0.00665)	(0.00618)	(0.00682)	(0.00665)	(0.00636)
Property Controls	N	Y	Y	Y	Y	Y	Y	Y
ZIP-Year FE	N	N	Y	Y	Y	Y	Y	Y
Insurer-Year FE	N	N	N	Y	Y	Y	Y	Y
Observations	3,215,282	3,035,099	3,002,325	3,002,225	3,002,225	3,002,225	3,002,225	3,002,225
Adjusted R^2	0.439	0.439	0.558	0.575	0.577	0.576	0.575	0.577

Table C.8: Sensitivity of Insurance Prices to Climate Risk by Regulatory Friction

This table presents the results of a property-level OLS regression to measure the sensitivity of insurance prices to climate risk depending on whether properties are in high vs. low friction states. The dependent variable is log insurance price (monthly premiums per $100,000 of coverage). High climate risk refers to properties in the top decile of climate risk by CoreLogic's AAL measure. AAL is a property-level measure of climate risk which refers to expected losses due to damages from climate disasters. We group states into two types: high and low friction, following the classification in Oh et al. (2021). High friction are states in the top tercile of regulatory friction and low friction are all the remaining states. Both specifications control for the property's replacement cost, which is obtained by multiplying structure size and local construction costs from the R.S. Means Company. Standard errors are clustered at the county level. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln P_{izt}	
	(1)	(2)
High Climate Risk = 1	0.0860***	0.0859***
	(0.0112)	(0.0112)
High Climate Risk = 1 × High Friction = 1	-0.0765***	-0.0777***
	(0.0123)	(0.0124)
ZIP-Year FE	Y	Y
Rebuild Cost Controls	N	Y
Observations	3238688	3238688
Adjusted R-squared	0.605	0.617

Table C.9: Sensitivity of Insurance Prices to Loss Ratios by Regulatory Friction

This table presents the results of a ZIP-level OLS regression to measure the sensitivity of insurance prices to loss ratio depending on whether properties are in high vs. low friction states. The dependent variable is the average annual insurance premium per policy in ZIP code z in year t. The variable of interest are the ZIP's lagged loss ratio – the average loss ratio from reporting insurers in the previous year – and a dummy for whether the ZIP is in a high-friction state. We group states into two types: high and low friction, following the classification in Oh et al. (2021). High friction are states in the top tercile of regulatory friction and low friction are all the remaining states. We include a fixed effect for a categorical variable that indicates market size where indicated. The data are from Treasury FIO. Standard errors are clustered at the ZIP level. Note: *p<0.1; **p<0.05; ***p<0.01.

	Premiums per policy$_{z,t}$	
	(1)	(2)
Loss Ratio$_{z,t-1}$	0.627***	0.608***
	(0.194)	(0.196)
High Friction=1 × Loss Ratio$_{z,t-1}$	-1.514***	-1.589***
	(0.447)	(0.406)
ZIP FE	Y	Y
Year FE	Y	Y
Market Size FE	N	Y
Number of Observations	89884	89884

Table C.10: OLS Results: Effect of Insurance Prices on Coverage (High Friction States)

This table presents the results of the following OLS regression:

$$\ln \text{coverage}_{i,c,z,t} = \alpha + \beta \ln P_{i,c,z,t} + \gamma' W_{i,c,z,t} + \varepsilon_{i,c,z,t}.$$

The dependent variable is ln coverage amount. The independent variable is log insurance price (monthly premiums per $100,000 of coverage). The analysis uses the sample of high friction states. High and low friction states are as defined in Section 1. Specifications include ZIP-year fixed effects and controls for property and borrower characteristics, as indicated. The exact list of control variables is as outlined in Table 4. Standard errors are clustered at the county level. The sample is as outlined in Section 1. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln coverage$_{izct}$	
	(1)	(2)
ln P$_{izct}$	-0.258***	-0.252***
	(0.0131)	(0.0131)
ZIP-Year FE	Y	Y
Property Controls	Y	Y
Borrower Controls	N	Y
Number of Observations	1446756	1446756

Table C.11: First-Stage Results: Instrumenting Insurance Prices in Low Friction States using High Friction States

This table presents the first stage of the two-stage least squares regression shown in Equation 5. The analysis uses the sample of low friction states. The dependent variable is log insurance price (monthly premiums per $100,000 of coverage), and the instrument is the log of average prices in high friction states for each credit score group in each month. High and low friction states are as defined in Section 1. All specifications include ZIP-credit score-year fixed effects and controls for property and borrower characteristics where indicated. Property-level controls include a dummy for whether the property is classified high climate risk (average annual loss is in the top decile), building age fixed effects, and the property's replacement cost. Borrower controls include house price, debt-to-income ratios, and borrower age fixed effects. We show the regression separately for each sub-sample considered in the paper. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	$\ln P_{i,c,z,m,t}$						
	(1) L/RC < 80%	(2) L/RC >= 80%	(3) High DTI	(4) Medium DTI	(5) Low DTI	(6) High House Price	(7) Low House Price
$\ln \bar{P}_{HF,ct}$	0.145***	0.145***	0.133***	0.157***	0.169***	0.139***	0.147***
	(0.004)	(0.004)	(0.006)	(0.007)	(0.008)	(0.006)	(0.005)
Zip-Year-Fico FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Insurer-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Property Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Borrower Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,237,437	1,237,437	375,612	365,688	305,466	542,247	603,505
Adjusted R^2	0.771	0.771	0.794	0.801	0.806	0.778	0.768

Table C.12: Overall Coverage Demand Elasticities without Replacement Cost Control

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 6. Column 1 shows the baseline elasticities on the full sample. Columns 2-3 shows sub-samples by whether loan to replacement cost ratios are above or below 80%. The analysis uses the sample of low friction states. The dependent variable is ln coverage amount. The independent variable is ln insurance price (monthly premiums per $100,000 of coverage), instrumented using average prices in high friction states for each credit score group in each month. All specifications include ZIP-credit score-year fixed effects, and additional fixed effects and controls for property and borrower characteristics, as indicated. Control variables are as defined in Table 4, except in this table replacement cost controls are excluded. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln coverage$_{i,c,z,m,t}$		
	(1) Overall	(2) Low L/RC	(3) High L/RC
ln $\widehat{P_{i,c,z,m,t}}$	-0.053***	-0.071***	0.039
	(0.015)	(0.015)	(0.061)
ZIP × Year × FICO FE	Y	Y	Y
Insurer × Year FE	Y	Y	Y
Property Controls	Y	Y	Y
Borrower Controls	Y	Y	Y
First Stage F Stat	1200	1190	100
Observations	1,633,602	1,237,437	307,734
Adjusted R^2	0.610	0.642	0.588

Table C.13: Overall Coverage Demand Elasticities with Calendar Month FE

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 6. Column 1 shows the baseline elasticities on the full sample. Columns 2-3 shows sub-samples by whether loan to replacement cost ratios are above or below 80%. The analysis uses the sample of low friction states. The dependent variable is ln coverage amount. The independent variable is ln insurance price (monthly premiums per $100,000 of coverage), instrumented using average prices in high friction states for each credit score group in each month. All specifications include ZIP-credit score-year fixed effects, and additional fixed effects and controls for property and borrower characteristics, as indicated. Control variables are as defined in Table 4, except this specification also includes calendar month fixed effects. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln coverage$_{i,c,z,m,t}$		
	(1) Overall	(2) Low L/RC	(3) High L/RC
ln $\widehat{P_{i,c,z,m,t}}$	-0.077***	-0.110***	0.051
	(0.013)	(0.013)	(0.062)
ZIP × Year × FICO FE	Y	Y	Y
Insurer × Year FE	Y	Y	Y
Property Controls	Y	Y	Y
Borrower Controls	Y	Y	Y
Calendar Month FE	Y	Y	Y
First Stage F Stat	1226	1239	100
Observations	1,610,605	1,237,437	307,734
Adjusted R^2	0.695	0.724	0.583

Table C.14: How bunching at loan balance varies with prices: robustness

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 6. The dependent variable is an indicator equal to one if coverage is close to the loan balance. In Panel (a), coverage is defined as close to the loan balance if it lies within the interval [0.95, 1.05]. In Panel (b), coverage is defined as close to the loan balance if it lies within the interval [0.99, 1.01]. The independent variable is log insurance price, defined as monthly premiums per $100,000 of coverage, instrumented using average prices in high friction states for each credit score group in each month. Column 1 shows the baseline elasticities using the same sample, controls, and fixed effects, as described in Table 4. Column 2 additionally controls for DTI and LTV ratios. Columns 3 and 4 show the semi-elasticities for sub-samples of loan to replacement cost ratio. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

Panel (a): Coverage within [0.95, 1.05] of loan balance

	Bunching at loan amount (Y/N)		
	(1) Overall	(2) Low L/RC	(3) High L/RC
$\ln \widehat{P_{i,c,z,m,t}}$	0.437***	0.497***	-0.171
	(0.030)	(0.030)	(0.112)
ZIP × Year × FICO FE	Y	Y	Y
Insurer × Year FE	Y	Y	Y
Property Controls	Y	Y	Y
Borrower Controls	Y	Y	Y
First Stage F Stat	1227	1237	100
Observations	1,610,605	1,237,437	307,734
Adjusted R^2	-0.099	-0.113	0.002

Panel (b): Coverage within [0.99, 1.01] of loan balance

	Bunching at loan amount (Y/N)		
	(1) Overall	(2) Low L/RC	(3) High L/RC
$\ln \widehat{P_{i,c,z,m,t}}$	0.156***	0.172***	0.039
	(0.017)	(0.017)	(0.075)
ZIP × Year × FICO FE	Y	Y	Y
Insurer × Year FE	Y	Y	Y
Property Controls	Y	Y	Y
Borrower Controls	Y	Y	Y
First Stage F Stat	1227	1237	100
Observations	1,610,605	1,237,437	307,734
Adjusted R^2	-0.036	-0.038	-0.012

Table C.15: Insurance demand elasticities: Evidence from within-person income changes

This table shows the second stage estimates of the two-stage least squares regression. The analysis uses the sample of low friction states and now considers mortgages within 3-5 years of origination, as indicated in the table. The dependent variable is ln coverage. The independent variables of interest are log insurance price (monthly premiums per $100,000 of coverage), and log insurance price interacted with a dummy for *High Income Transition*, which equals one if the household was defined as low income at the beginning of the sample but high income anytime within the first 3-5 years of the mortgage. Both regressors are instrumented using average prices in high friction states and average prices in high friction states interacted with the *High Income Transition* dummy. High and low friction states are as defined in Section 1. All specifications include ZIP-credit score-year fixed effects, and additional fixed effects and controls for property and borrower characteristics, as indicated. Control variables are as defined in Table 4. Standard errors are clustered at the county-year level. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln coverage$_{izcmt}$	
	Within 5Y of origination	Within 3Y of origination
	(1)	(2)
$\ln P_{iczmt} \times \widehat{\text{HighIncomeTransition}}_{it}$	0.0132**	0.0193**
	(0.00587)	(0.00964)
$\ln \widehat{P_{iczmt}}$	-0.0726***	-0.200***
	(0.0253)	(0.0556)
Sample	Low Income at Origination	
Loan FE	Y	Y
Loan Age FE	Y	Y
Current FICO FE	Y	Y
Year FE	Y	Y
Replacement Cost Control	Y	Y
High Income Transition FE	Y	Y
Number of Observations	3197998	2687540
Adjusted R-squared	0.0792	0.136
Overall First-Stage F	90.56	25.45