The Insurance Protection Gap *

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

Original Draft: May 2024
Current Draft: November 2025

Abstract

In a world of rising climate risk, how well insured are U.S. households? We provide the first quantification of the "insurance protection gap"—the extent of households' exposure to uninsured property losses, and the first estimates of households' demand elasticities. We utilize newly available, individual-level data on homeowners insurance contracts spanning the entire United States over more than a decade. We show the insurance protection gap is large, particularly among financially vulnerable lower income households. Importantly, households and mortgage lenders jointly determine insurance coverage. We find that lender exposures are largely covered, and that household equity is under-insured. To study insurance demand, we exploit quasi-exogenous regulatory shocks to insurance pricing, and find that demand elasticities vary significantly across households. High income households are price inelastic, and adjust mortgage debt to maintain coverage. In contrast, low income households are price elastic, particularly over the discretionary component of coverage that exceeds lender requirements. Low income households also adjust mortgage debt far less than high income households. We argue that these disparate patterns are likely driven by mortgage credit supply. Credit supply expands to allow high income households to finance insurance using credit, but it does not expand for low income households, reflecting a lender tradeoff between leverage and under-insurance. Our estimates imply that between 17% to 31% of households would hit binding loan-to-value (LTV) constraints and be forced to reduce coverage substantially, if prices were to continue growing at historical rates moving forward. These results imply that insurance markets may shrink even as the financial risks from natural disasters continue to increase.

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 (discussant), Lauren Cohen, Tony Cookson (discussant), Cameron Ellis (discussant), Shan Ge (discussant), Robin Greenwood, Mallick Hossain, Ben Keys, Ralph Koijen, Ulrike Malmendier, Philip Mulder, Borghan Narajabad, Monika Piazzesi, Hyeyoon Jung (discussant), Motohiro Yogo, and participants at the 2024 Yale Junior Finance Seminar, the RFS-Oklahoma Climate Conference, the UNC-Duke Corporate Finance Conference, and the UNC Wealth Inequality Conference 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.

†Columbia Business School; prs2107@columbia.edu.
‡Federal Reserve Board; Tess.Scharlemann@frb.gov.
§Harvard Business School; isen@hbs.edu.
¶Federal Reserve Board; ana-maria.k.tenekedjieva@frb.gov.

Over the last fifty years, property damage from natural disasters has dramatically risen. In the United States, an estimated 35 million homes, or one-third of the housing stock, has high exposure to natural disaster risk.[1] Property insurance is the primary way that households financially protect their homes in the face of escalating climate risks. Insurance payments are used by households to repair and rebuild homes damaged by natural disasters, helping to preserve household wealth. However, there are mounting concerns that the most vulnerable households may not have sufficient insurance coverage–what policymakers refer to as a "protection gap." Under-insurance could make households more financially exposed to catastrophic shocks, and create risks for lenders and taxpayers if disasters precipitate household default. Despite intense public concern, there is limited systematic evidence on the extent to which households are under-insured, and what factors influence their coverage choices.

An important feature of the U.S. homeowners insurance market is that households choose contract terms jointly with mortgage lenders. Nearly two-thirds of homeowners finance their home purchases with a mortgage. Lenders face substantial losses when natural disasters damage mortgaged properties and precipitate borrower default. To mitigate these risks, lenders typically require borrowers to purchase and maintain homeowners insurance throughout the life of the loan, often specifying minimum coverage amounts.[2] A key implication of this institutional arrangement is that lenders constitute a significant component of insurance demand, significantly impacting how households can adjust insurance coverage in response to rising insurance prices. Furthermore, the credit supply choices of lenders become particularly important if households need to finance insurance coverage using debt– an outcome that emerges in the textbook models of insurance demand that assume frictionless access to credit markets (e.g., Arrow (1963), Mossin (1968)). From the perspective of a mortgage lender, it may be optimal to extend credit for households to finance insurance coverage, because more coverage helps reduce default in the aftermath of disasters. However, increasing leverage also worsens default risk because of both adverse selection and moral hazard concerns, creating a tradeoff for lenders between the benefits of more insurance coverage and the costs of

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

[2] Because lenders also monitor the maintenance of coverage, they often require that premiums be paid through escrow accounts and that borrowers pre-fund up to one year of premiums at origination. Lenders are also often listed as beneficiaries on the insurance contract, so they can receive insurance proceeds after natural disasters even if the household defaults on the mortgage. For more details see here.

higher leverage. Therefore, understanding the causes and consequences of the protection gap requires taking a careful account of what households choose optimally, what is required by lenders, how households rely on debt to finance insurance, and what risks are borne by households versus lenders.

This paper provides the first comprehensive analysis of how households and lenders jointly determine homeowners insurance coverage at mortgage origination. 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 2013–2023. Exploiting quasi-exogenous price variation induced by regulatory pricing frictions, we provide the first national estimates of insurance demand elasticities. We have three main findings. First, there is an "insurance protection gap"— households' exposure to uninsured property losses is large, particularly among lower-income households. Lender exposures are largely insulated due to required coverage mandates, however, household equity remains substantially exposed. Second, high-income households exhibit inelastic coverage demand, financing higher premiums by increasing mortgage borrowing. In contrast, low-income households are price elastic, particularly for the discretionary portion of coverage that exceeds lender requirements. Third, binding loan-to-value constraints can explain these disparities. High income borrowers are able to expand debt to finance insurance, protecting their equity values. In contrast, low income households cannot do so because lenders restrict credit supply, which forces them to reduce coverage amounts instead to alleviate the burden of rising premiums. Collectively, these results show that lender requirements and credit supply significantly influence coverage choices and that the insurance protection gap varies across the income distribution.

Our paper proceeds in three parts. In the first part of the paper, we present new facts on the household insurance protection gap. We find that the average U.S. household is under-insured at mortgage origination, with only 70% of the rebuilding costs covered by the insurance contract.[3] The insurance gap is especially pronounced among the most financially vulnerable households, those with lower incomes and credit scores, and persists even for low income households facing high climate risk. Crucially, we find that coverage amounts display pronounced bunching at the loan balance, with a vast majority of households obtaining coverage above this threshold. This bunching is consistent with the notion that lenders require households to obtain insurance coverage that at least covers loan

[3]Coverage ratios are quite persistent through the life of the loan, meaning under-insurance usually begins at origination.

exposures. Such requirements ensure that lender exposures are protected, while allowing limited flexibility for coverage below the loan amount.

In the second part of the paper, we estimate insurance demand elasticities as a first step toward understanding households' willingness or ability to pay for coverage (Einav et al., 2010). To obtain plausibly exogenous price variation, we implement a Hausman-type instrument (Nevo (2001), Hausman (1996)), exploiting the fact that multi-state insurers cross-subsidize risks across states (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 underwriting losses. In contrast, in "low-friction" states, insurers can adjust prices more responsively to changing risks and demand conditions. Following Oh et al. (2021), we exploit this cross-subsidization across states by using average monthly prices in high-friction states by credit score group as an instrument for prices in low-friction states.[4] In the first stage, we regress low-friction state prices on the instrument, controlling for zip-year-credit score fixed effects and other borrower and property-level covariates. In the second stage, we regress coverage on the instrumented prices using the same specification. This design compares observably similar borrowers within the same zip code of a low friction state, who originate mortgages in different months of the same year and thus face different prices driven by regulatory timing in high-friction states. The identifying assumption is that the timing and magnitude of these regulatory decisions in high friction states are exogenous to demand conditions in low-friction states, and only effect low friction state prices through insurer supply because of the cross-subsidy.

Using this IV approach, we find that high income households are inelastic in their demand for coverage. A 1% price change leads to a statistically insignificant and economically negligible reduction in coverage. In contrast, we find that for low income households, a 1% increase in insurance prices leads to a 0.18% reduction in coverage. Thus, low income households are more elastic than high income households, implying that they have a lower willingness-to-pay (or ability-to-pay) for insurance. However, while more elastic, even low income households are not very elastic in absolute terms. This makes sense since lenders likely put limits on how much households can reduce coverage.

[4]We split borrowers into four credit score groups. The instrument varies at the credit score group - month level.

As a result, households can only reduce coverage to a smaller degree in response to price changes.

To better understand these dynamics, we look at the *discretionary* component of coverage, which we define as coverage in excess of the loan amount. This definition comes from the stylized fact we uncover that coverage seems to bunch at the loan amount. It seems that most households obtain coverage in excess of the loan amount, but that lenders are unwilling to let households take on coverage limits that are lower than the loan amount. When looking at this discretionary portion of coverage, we find that low income borrowers are highly elastic: a 1% increase in insurance prices leads to a 2.4% reduction in discretionary coverage. These results suggest that as insurance prices grow, it is the low income household's equity that will become more exposed, since lenders ensure that their exposures remain covered.

In the third part of the paper, we go beyond lender requirements on coverage, to consider how lenders enable households to adjust mortgage debt in order to finance insurance coverage. To shed light on debt dynamics, we utilize the same Hausman instrument as earlier, using prices in high friction states as an instrument for prices in low friction states. Rather than using coverage as the outcome variable, we consider measures of household mortgage debt, namely the loan-to-value ratios, and we estimate how debt adjusts in response to quasi-exogenous changes in insurance prices. We argue that the instrument is also valid in this setting, since the timing of regulatory decisions in high friction states are also unlikely to be driven by credit conditions in low friction states.

We find that plausibly exogenous increases in insurance prices lead to significant increases in loan-to-value ratios. Interestingly, the debt adjustment is nearly two times larger for *high* income households than lower income households. These results suggest that high income households are better able to finance rising premiums with debt. Taking the results of parts 2 and 3 together, in response to insurance price shocks, low income households reduce coverage and have limited debt adjustments, whereas high income households maintain coverage and increase debt balances. The differences we find between high and low income households show that insurance coverage demand elasticities must account for the debt capacity of the borrower, which is determined by lender credit supply. Our estimates imply that in a counterfactual scenario where the higher income households could not access debt markets, their coverage demand elasticities would be over 2.4 times higher.

We interpret the opposite elasticities in the insurance and mortgage contracts as reflecting

lender credit supply choices. That is, lenders extend more credit for high income households to buy insurance, but choose not to extend more credit to low income households. An alternate explanation of our elasticities is that low income households do not value insurance, and therefore they have no additional credit demand to finance insurance. This could occur if, for example, lower income households have different beliefs about risk, informational constraints, risk preferences, sophistication, disaster aid expectations etc, all of which could lead them to value insurance less than high income borrowers. To differentiate between these two competing explanations, we consider the following within-household test that exploits the panel nature of our data. We now look beyond mortgage origination, and consider how the same household adjusts insurance over time. We limit the sample to low income borrowers, and identify when a borrower transitions from low income to high income within the first three years of the mortgage. We then test how insurance demand adjusts in response to exogenous price changes for households when their income transitions from low to high income, incorporating a borrower fixed effect. We argue that the borrower fixed effect helps capture other drivers of demand that are intrinsic to the borrower and unlikely to change in such a short period of time, such as beliefs and preferences. Using this test, we find that the same borrower becomes price inelastic when they transition from low to high income, suggesting that low income households do value insurance. We also consider a separate test evaluating lender denials of mortgage applications based on the LTV and income of the borrower, and how these denials relate to insurance prices. For low income borrowers, we find that lenders are more likely to deny high LTV loans when insurance prices are high. However, we find no such effect for high income borrowers. These results provide a second piece of evidence favoring a credit supply interpretation.

The fact that households adjust mortgage debt in response to insurance prices may at first be surprising. However, the institutional features of the housing search process can explain why this occurs. Usually households begin searching for insurance only after putting a deposit down on the home, typically as part of the process of closing the mortgage. Higher than expected prices could directly exacerbate household liquidity constraints, since households have to escrow 6-12 months of premiums upfront with the mortgage lender.[5] For many households, escrow requirements on premiums could end up coming out of what they saved for the down payment, leading to a higher

[5]Canceling the housing transaction at this point could be very costly, and households may have limited alternatives if they are liquidity constrained and cannot adjust savings or consumptions to meet this requirement. Additionally borrowers often do not have the option to opt out of escrow requirements, especially low income borrowers.

loan-to-value mortgage. Furthermore, if households anticipate that such higher-than-expected prices will persist over their tenure of the house, they may also choose to pay a smaller downpayment to have liquidity buffers for smoothing these burdens over a longer period. Mortgage finance is cheaper than nearly all other forms of financing available to households (e.g., credit cards, personal loans, etc.), which could also make this prudent for households.

On the lender side, it may be advantageous to provide households with more credit to finance insurance, since higher coverage may make default less likely after disasters. However, lenders face an important trade-off because higher leverage can increase default risk in other states of the world. Thus, lenders may not enable all households to use leverage, and have incentives to ration credit. Our results suggest that lenders enable high income households to use leverage, but not low income households. This makes sense since low income households are more likely to already have very high LTV contracts, making any additional leverage too risky for the lender. For example, moving from 70 to 72% leverage is very different in terms of the added default risk than moving from 90 to 92%. Our results show that the credit supply decisions of lenders then spillover to households' insurance choices, influencing how much of their equity is protected by insurance.

We conclude by discussing the broader implications of our findings. It is likely that as climate-related risks rise, insurance premiums will keep rising in tandem with risks. Our estimates imply that, if prices continue growing at historical rates moving forward, between 17% to 31% of households would hit binding LTV limits and be forced to reduce coverage substantially, meaning insurance markets may shrink even as climate risks rise. Furthermore, these changes will reduce household financial resiliency in two ways: by increasing household leverage, and by increasing the protection gap further. Finally, there are significant distributional implications since high income households increase debt, while low income households reduce coverage. Taken together, the results suggest that both household leverage and the protection gap is likely to grow in size, posing significant challenges for the most financially vulnerable households, the insurance market, and the mortgage market more broadly. A natural policy implication is to provide more transparent insurance price information during the housing search process. Accurate pricing information is difficult to obtain until the household has already put down a deposit and spent significant resources to find a home; this makes it difficult and costly for households to adjust on other margins besides debt and coverage.

Related literature. This paper contributes to three strands of the literature. First, we add to the literature on the determinants of households' demand for insurance. Our work starts from the theoretical models of Arrow (1963), Mossin (1968), Froot et al. (1993), Rampini and Viswanathan (2013), who consider optimal insurance choices and, in the latter two papers, the impact of financial constraints on insurance demand. To bring these theories to the data, we follow the empirical approach of Einav et al. (2010) and Einav and Finkelstein (2023) of estimating demand elasticities and, to our knowledge, are the first to estimate these parameters for the U.S. homeowners insurance market.[6] In addition, we are the first to study the lender component of insurance demand, and how mortgage demand and supply adjusts in response to shocks to insurance prices. The linkages we document between debt and insurance connects to a rich literature in development economics that considers how to optimally bundle insurance and credit contracts in light of financial and informational constraints (Braverman and Stiglitz (1982), Banerjee et al. (2014), Casaburi and Willis (2018), Karlan et al. (2014), Annan (2021)). We show that many of these insights are generalizable to the average mortgage borrower in the United States, and also contribute to this literature by considering endogenous responses in credit supply. In particular, we show that there is heterogeneity in when lenders allow households to take on more debt to finance insurance.

Second, our paper contributes to the small but growing new literature studying homeowners insurance markets. Thus far, this literature has focused on the supply-side of insurance, primarily on the determinants of pricing and insurer rationing. For example, studies have documented the role played by capital market frictions (Froot and O'Connell, 1999, Jaffee and Russell, 1997), state-level price regulations (Oh et al., 2021), informational asymmetries (Boomhower et al., 2023), trust and claims' validity (Glaeser and Shleifer, 2003, 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). In contemporaneous work Cookson et al. (2023) also study insurance demand in the context of victims of the Marshall Fire in Boulder Colorado. They show that the high-income households in their sample are also under-insured, and consider the role of household shopping behavior in driving under-insurance. Relative to this literature, our

[6]There are some estimates of demand elasticities on the extensive margin for flood insurance in Wagner (2022), Collier et al. (2024), and Liao and Mulder (2021). Flood insurance covers flood damage and is only required for mortgages in flood zones, creating limited take-up nationwide (Amornsiripanitch et al., 2024). In contrast, homeowners insurance is required for all mortgages, with nearly 90% of all homeowners purchasing homeowners insurance.

novel contribution is to focus on the role of mortgage lenders and the role played by credit supply. In addition, to our knowledge, we are the first to estimate household demand elasticities, and to study homeowners insurance demand in a nationwide context with over a decade of data.

Third, our findings speak to the literature on climate risk and household financial resiliency. Thus far the literature has focused on mortgage default (Issler et al., 2020, Kousky et al., 2020, Sastry, 2025, Ge et al., 2024, Biswas et al., 2023), post-disaster resiliency (Gallagher and Hartley, 2017, Cookson et al., 2023), and house price declines (Baldauf et al., 2020, Murfin and Spiegel, 2020, Bernstein et al., 2019).[7] A key finding of our paper is that households increase their borrowing to finance insurance purchases. In that sense, we also contribute to the very large literature on the drivers of mortgage debt (e.g., Adelino et al. (2016), Bhutta and Keys (2022), DeFusco and Paciorek (2017), Defusco et al. (2019)). This literature has shown that higher debt is strongly associated with worse default outcomes (e.g., Greenwald (2018), Gupta and Hansman (2021)). This suggests that both under-insurance and rising leverage are related channels through which insurance premiums worsen household financial resiliency. Under-insurance makes the mortgage riskier during natural disasters. Leverage makes the mortgage riskier more broadly, not just in times of natural disasters. To the extent that premiums continue to rise because of rising climate risk, our results predict this will worsen household financial resiliency through both channels.

1. Institutional Background, Data, and Facts on Underinsurance

1.1. Institutional Background

For most American households, their house represents a significant portion of their wealth. The vast majority of households rely on mortgages to buy homes, and mortgage lenders require that a number of criteria be met for this financing to be provided, including a requirement that borrowers purchase homeowners insurance. As a result, the homeowners insurance product is nearly ubiquitous. Insurers sell homeowners multi-peril insurance coverage to nearly 85% of all U.S. homeowners (Jeziorski et al., 2021), representing over $15 trillion in coverage taken out annually (Oh et al., 2021).

Homeowners Insurance: The standard homeowners contract is annual and covers damages

[7] They also bear indirect exposures through labor markets (Kruttli et al., 2019) and discounts in municipal bond prices (Goldsmith-Pinkham et al., 2020).

from most natural disasters.[8] Around 90% of homeowners insurance claims are for property damage caused by weather events.[9] The insurance contract has two key characteristics: coverage and premiums.[10] Those purchasing insurance are entitled to claim payouts up to a pre-specified coverage amount if an insured loss event materializes. The premium is the amount that the insured party must pay to the insurer. It is usually paid every month, and is a function of the risk of the house, and the coverage choices of the insured. Premiums may also reflect other regulatory and operational costs borne by the insurer, which are passed onto households.

Timing of purchase: Homeowners insurance is typically purchased after the household has made its housing decision, but before the mortgage closes. By this stage, the household has already placed a deposit on the home, making it difficult—and often costly—to adjust its housing consumption. However, the household can still choose its insurance coverage level and adjust its mortgage terms, including the loan amount. For this reason, we interpret our estimated elasticities as capturing households' adjustments to coverage and debt conditional on a given housing choice, rather than responses that operate through changes in housing consumption.

Mortgage servicers: Mortgage servicers are usually responsible for collecting borrowers' monthly mortgage payments on behalf of the lender and monitoring that other requirements, such as insurance and property tax payments, are met. To simplify payments and enforce requirements, most tax and insurance payments are collected and paid using escrow accounts. 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. The majority of mortgages feature escrow deposit accounts (Anderson and Dokko, 2008), and are required by FHA, Fannie Mae, and Freddie Mac in most circumstances.[11] This feature of mortgage lending enables the data provider ICE McDash to obtain micro data on both the mortgage and the homeowners insurance

[8] An important exception is flood risk, which is carved out from standard homeowners insurance contracts and mostly provided through the government-run National Flood Insurance Program.

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

[10] Another feature of insurance contracts are deductibles, which is the amount the borrower must pay before insurers start paying for insured events. In the data, there is little variation in deductibles and deductibles are orders of magnitude smaller than coverage amounts, with most contracts ranging between $500 and $2000 dollars. In case of catastrophic losses, the amount by which households are underinsured would matter a lot more than the deductible amount they would need to pay, even at the higher end of deductible levels.

[11] 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.

contract, because mortgage servicers monitor both contracts together for each individual.

Upfront escrow requirements: When a mortgage is originated, servicers estimate the annual tax and insurance outlays associated with the mortgage, establish a monthly payment sufficient to meet those payments, and collect a reserve equal to no more than 1/6 the annual expected escrow expenses.[12] Because homeowners insurance is paid annually for the year ahead, the servicer collects the entire homeowners insurance payment as part of the initial escrow balance, making insurance burden salient at origination.

Lenders are beneficiaries to insurance proceeds: Mortgage lenders have a direct interest in homeowners insurance proceeds, because the home serves as collateral for the loan. 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 lender (or mortgage servicer).[13] The lender typically controls how those insurance proceeds are disbursed — often via an escrow mechanism, ensuring that the funds are applied appropriately, e.g., to rebuilding and repairs or if repairs are not feasible toward the loan balance.[14] These arrangements ensure that lenders are protected in the event of a natural disaster, shielding them from the possibility that households will default on the mortgage and abscond with insurance payments. 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 lender, thereby shielding the lender from potential losses.

1.2. Data

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), mortgage contract, borrower characteristics, households' debt and 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

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

[13] For example, see "*Understanding the insurance claims payment process,*" Insurance Information Institute. Report can be found here.

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

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, with more details and the merge procedure described in Appendix A.

1.2.1. 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, etc.). We merge this with McDash's mortgage performance modules, which contain data on each loan's performance history from its origination to its final payment. This includes whether the mortgage is current or in delinquency status, and records events such as prepayment, delinquency, default or foreclosure.

(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 D.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: To incorporate borrower income information, we use modeled primary-borrower income from Equifax's Personal Income Model (PIM), which uses characteristics from a borrower's credit bureau profile to predict their personal income. These data are linked to the

loan-level McDash data using Credit Risk Insights Servicing McDash (CRISM).[15]

(4) CoreLogic Deeds Module: Next, we bring in property characteristics, such as the geographic coordinates, structure age, and size of the property, using the CoreLogic Deeds database, which compiles property characteristics from tax assessment records. We merge these data to the loan-level data using a fuzzy matching algorithm we have developed based on address (ZIP code), origination date, loan amount, loan purpose, and loan type. We validate this merging procedure using the proprietary FR Y-14 data at the Federal Reserve which contains confidential borrower-level data on both the loan contract and property address.

(5) CoreLogic Climate Module: We also integrate key measures of property-level climate risk from the CoreLogic 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, CoreLogic's climate module is also widely used by lenders and insurance companies.[16] We also use CoreLogic parcel-level estimates of replacement costs, which is calculated for a snapshot in 2023.

(6) R.S. Means Construction Costs: Finally, to understand the extent to which households are under-insured, we incorporate data on local construction costs from R.S. Means company, which is widely used in the urban and housing literature (e.g., Glaeser and Gyourko, 2005) that is developed based on detailed surveys of home builders across metropolitan cities nation-wide. We apply these indices to our replacement cost estimates from CoreLogic to determine the replacement cost at mortgage origination.

Figure D.1 provides a high-level summary of the data construction, and Appendix A provides further details. We next describe the construction of our key variables.

1.2.2. Measurement of Key Variables

Coverage Adequacy and Under-insurance: To assess how insured a given homeowner is, we compare the coverage taken with the replacement cost of the property, defined as the amount it

[15] CRISM is an anonymous match between ICE-McDash mortgage servicing records and borrowers' credit files. The Personal Income Model was calibrated on an anonymous subsample of credit profiles using earnings income reported by employers.

[16] For example, the *Wall Street Journal* writes: "CoreLogic, 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.

would cost to rebuild the exact same property in case of a total loss event.[17] Specifically, we define Coverage Ratio for property p at time t as:

$$\text{Coverage Ratio}_{pt} = \text{Coverage}_{pt}/\text{Replacement cost}_{pt}. \tag{1}$$

If coverage equals rebuild cost, then a homeowner can expect insurance payments to cover all potential rebuilding costs, whereas if coverage is below rebuild cost, the homeowner is responsible for any repair costs that exceed the coverage limit. Thus, a coverage level below replacement cost indicates that the household is under-insured.

Our measure of coverage comes from the McDash Insurance module. The coverage variable we observe in our data refers to basic policy limits for the structure.[18] To estimate replacement costs, we use CoreLogic estimates and deflate them using local construction cost indices from R.S. Means to obtain the replacement cost at mortgage origination.[19] Insurers recommend that households determine coverage based on the rebuild costs rather than house prices to assess resiliency against losses. While the two are highly correlated,[20] land values and local amenities (or disamenities) affect house prices but may not impact reconstruction costs after a disaster, making rebuild cost the relevant measure.[21]

Insurance Price: We define the price of insurance as the monthly insurance premium per

[17] We use the terms replacement cost, reconstruction cost, and rebuild costs interchangeably.

[18] 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.

[19] We check that these estimates highly correlated with other measures of replacement costs. For example, we can also multiply local per square foot construction costs for an average quality home, which we take from R.S. Means, and the structure size, which we obtain from CoreLogic Deeds, to obtain Replacement Cost = Construction cost per square foot × Total square feet. This alternate procedure is based on Sastry (2025) and follows insurer recommendations. For example, the insurance trade group Insurance Information Institute (III) writes: "For a quick estimate of the amount of insurance you need, multiply the total square footage of your home by local, per-square-foot building costs. (Note that the land is not factored into rebuilding estimates.)". See Insurance Information Institute, "How Much Homeowners Insurance Do You Need", available from III Both approaches yield similar estimates.

[20] At origination, for the ZIP-year level, the regression slope is a statistically significant 83%, with an R^2 of 36%. At the individual level the slope is 67%, with an R^2 of 40%.

[21] For example, the insurance trade group III writes: "In every case, you'll want the limits on your policy to be high enough to cover the cost of rebuilding your home. The price you paid for your home—or the current market price—may be more or less than the cost to rebuild." See Insurance Information Institute, "How Much Homeowners Insurance Do You Need", available from III.

$100,000 of coverage for property p at time t, which we obtain from the McDash Insurance module.

$$\text{Insurance Price}_{pt} = \text{Monthly Premium}_{pt}/\$100{,}000 \text{ coverage}_{pt}. \tag{2}$$

Climate Risk: Our property-level climate risk exposure measure uses Corelogic's AAL – expected annual climate losses of the property – as a share of the property's rebuild cost. Because the measure is subject to model risk and has a highly skewed distribution, we identify high climate risk properties as those in the top decile of this AAL measure.[22]

High and Low Friction States: Homeowners insurance premiums are subject to extensive regulations in the U.S. at the state level. Regulators actively monitor insurers' pricing decisions and prescribe a wide range of pricing rules and guidelines. Oh et al. (2021) show that these practices lead to stark differences in pricing restrictions across states. We exploit the regulatory constraints to adjusting insurance prices to obtain exogenous variation in insurance pricing. To do so, we classify states as being high friction (highly regulated) or low friction (less regulated), following the methodology of Oh et al. (2021).[23]

1.2.3. Focus on Mortgage Origination

We highlight an important feature of the data which motivates our analysis. In particular, our analysis focuses on understanding the drivers of coverage choices and mortgage debt at *origination*.[24] We focus on origination because most of the variation in coverage arises at the time of origination, while changes in coverage ratios over the remaining life of the mortgage are limited. Figure D.5 shows a binned scatter plot of coverage ratios by loan age, after controlling for a loan fixed effect. Coverage ratios stay roughly flat at 70% throughout the life of the loan. This shows that households typically choose insurance coverage at the same time they are setting their mortgage terms. On average, they obtain less coverage, and the under-insurance resulting from that initial coverage decision persists throughout the life of the loan.

[22]We validate this measure by verifying that is strongly correlated with the probability of filing a claim and claim severity using aggregated ZIP-year level data. See Appendix A.

[23]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.

[24]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.

Key summary statistics for our final dataset is in Table 1, which we discuss in more detail below. Our dataset has roughly 3.2 million observations, with 1.8 million coming from low friction states, and 1.5 million from high friction states. Appendix A also provides additional descriptive facts.

1.3. Three Novel Facts on Homeowners Insurance Coverage

In this section, we document three new facts about homeowners insurance coverage that motivate our focus on mortgage lenders and financial constraints as drivers of insurance quantity choices.

1.3.1. *Fact 1:* The average American household is under-insured.

We start by estimating insurance coverage ratios, and find that the average American household is substantially under-insured. Table 1 shows that the average coverage ratio nationally is 69%. This means that the average household has less than 70% of their rebuild cost covered by insurance. In fact, nearly 61% of households have coverage ratios below 70%. Furthermore, more than 10% of households are *severely* under-insured having only 50% of their rebuild cost covered by insurance. Figure 1 shows how underinsurance varies spatially. 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.

1.3.2. *Fact 2:* Lower income borrowers have less coverage.

We next dig deeper into the prevalence of under-insurance by considering how under-insurance varies by borrower characteristics. Figure 2 shows how coverage ratios relate to borrower credit scores and borrower incomes.[25] We find a strong positive correlation between credit scores and coverage ratios (Figure 2a). Going from the lowest ventile of credit scores (650) to the highest ventile of credit scores (800) is associated with an increase in coverage ratios of 10 pp, i.e. coverage ratios rise from 64% to 74%. The relationship is quite linear. Figure 2b shows a similar pattern when plotting coverage ratios as a function of the primary borrower's income. Here the difference in the cross-section of borrowers is even starker. Going from the lowest ventile of income ($20K) to the highest ventile of income ($120K) is associated with a 30 pp increase in coverage ratios, from 60% to 90%. 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

[25] Credit scores are measured using FICO score at origination from ICE McDash.

financial resources of their own to smooth such shocks. Paradoxically, these are the exact households most likely to benefit from insurance, and the ones that standard models (Arrow, 1963, Mossin, 1968) would suggest should have the highest demand for insurance.

Next, we analyze the relationship between severe under-insurance (coverage ratios below 50%) and borrower characteristics, while controlling for borrower's climate risk exposure. For borrowers with higher income and higher credit scores, the data aligns with traditional economic theory: households facing higher climate risk tend to obtain more coverage and are less likely to be severely underinsured (Figure 3b and Figure 3a). In contrast, lower-income and lower-credit-score households exhibit different behavior: climate risk appears to have minimal influence on coverage choices. For example, households with annual primary borrower incomes below $40k show similar rates of severe under-insurance, regardless of their climate risk exposure. Surprisingly, at the lowest income levels (around $20k annually), households in high-risk areas are more likely to be severely under-insured compared to those in low-risk areas. Similarly, borrowers with lower 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 are the most likely to be severely under-insured. This suggests that financial constraints may outweigh risk consideration in determining insurance coverage, consistent with the broad patterns implied by Figure 1 at the state level.

1.3.3. *Fact 3:* Coverage bunches at the loan amount

We next examine the coverage mandates that lenders may impose on households. Mortgage lenders face substantial losses when natural disasters damage mortgaged properties and trigger borrower default. To protect the value of their collateral, lenders often require borrowers to maintain a minimum level of insurance coverage. There is little theoretical work on how lenders should optimally set insurance requirements. Anecdotally, the government-sponsored enterprises' guidelines for mortgage origination and servicing link required coverage amounts to replacement costs and loan balances. Private lenders may similarly require full replacement-cost coverage or establish minimum coverage thresholds. While requirements likely vary across lenders and borrower types, the data reveal a striking empirical pattern: coverage amounts exhibit pronounced bunching at the loan balance, with most borrowers holding coverage at or above this level (Figure 4). This bunching is

likely to be consistent with lender-imposed thresholds. Just as mortgage underwriting commonly relies on rules of thumb—such as minimum credit score, maximum DTI, or LTV cutoffs, our findings suggest that lenders apply a similar heuristic requiring, at minimum, insurance coverage up to the loan amount.[26] Lenders may prefer even higher coverage levels, but they appear to allow lower coverage only sparingly.

A natural question is what we might expect households to do in the absence of lender-imposed requirements. What is optimal will depend heavily on household costs of default and the extent to which households act strategically. If households face high or infinite costs of default, then they may optimally choose to fully insure, following the insights of Arrow (1963) and Mossin (1968). On the other hand, if households value the default option, and view their ownership of the home through the lens of their equity exposure, then we would expect insurance choices to cover households' equity exposures and not the component reflecting the lender's portion. Specifically, when households can default on their mortgage following a severe loss, default itself provides a form of implicit insurance (Mahoney (2015)—it limits downside risk once the property's value falls below the outstanding loan balance.[27] Consequently, households have little incentive to insure beyond their own equity share, since losses exceeding equity would be absorbed through default. This intuition implies that the strategic behavior of households would lead coverage to "bunch" around the household's equity value, rather than the loan amount.

1.3.4. Interpreting the Descriptive Facts

These stylized facts underscore the need to systematically understand the drivers of under-insurance for low income households, in particular the role played by financial constraints and lender requirements. The broad patterns across the three facts show a clear correlation between household income and coverage, and a role of lenders. Taken together, these facts are consistent with a model where households value insurance, but exhibit financial constraints that prevent them getting full insurance coverage. Additionally, they show that households do not have unlimited discretion in choosing coverage levels; lenders limit how under-insured households can be, because

[26] We verify that this threshold does not coincide with other benchmarks, such as 80 percent of replacement cost or the property value.

[27] This logic parallels the insight in Mahoney (2015) that bankruptcy protection acts as implicit health insurance by truncating individuals' financial losses from large medical expenses and thereby crowding out formal insurance demand.

they require that, at a minimum, loan exposures are covered. While these facts do not reflect casual relationships, we use them to motivate our focus on financial constraints and lenders as drivers of insurance coverage throughout the rest of the paper.

2. Framework: Insurance and Credit Demand with Borrowing Constraints

In this section, we provide a single conceptual framework that brings together the insights of the textbook Arrow (1963) and Mossin (1968) models along with the newer settings of Rampini and Viswanathan (2013) and Froot et al. (1993) that consider financial constraints.[28] Our framework introduces insurance choice into a standard dynamic consumption-savings model with borrowing constraints. The model generates predictions for how households jointly choose insurance coverage and debt, and comparative statics on how these choices depend on insurance prices and vary with borrowing constraints. The goal of the model is to provide an organizing framework to motivate and interpret our empirical tests throughout the rest of the paper.

2.1. Insurance Coverage Demand: Unconstrained Benchmark

We consider a dynamic partial equilibrium model of a risk-averse individual with utility function U who uses homeowners insurance to protect her house against disaster risk. Each period there is a chance that the house may be fully destroyed by a disaster event, denoted by a binary random variable y with 1 indicating the event of total loss. The probability of loss is known and is given by the parameter $\phi = E[y]$. M is the value of the housing consumption per period, and is taken as exogenous.[29] The individual chooses to insure a fraction $k \in [0, 1]$ of the house at the beginning of the period at the insurance price of p. If a loss occurs, the individual receives the payment kM.

To set ideas, we begin with a benchmark case where households do not face borrowing constraints and insurance prices are assumed to be actuarially fair.

The individual chooses consumption, insurance, and savings to maximize her discounted lifetime utility subject to a budget constraint:

[28] Our model builds on Liu and Myers (2012), replicating their results on constraints but also adding several additional propositions related to optimal debt choice. For related literature on insurance demand with liquidity constraints, see Ericson and Sydnor (2018), Rampini and Viswanathan (2013), and Gollier (2003).

[29] We abstract from the housing choice to focus on the partial equilibrium effects of financial constraints and insurance demand. In this model, in each period, the individual stands to gain M in wealth if there is no shock, or if she is fully insured.

$$\max_{\{c_t,k_t,S_t\}} E_0 \sum_{t=0}^{\infty} \beta^t U\left(c_t\right), \tag{3}$$

subject to:

$$S_t = w_t - c_t - pk_t, \tag{4}$$

$$w_{t+1} = (1+r)S_t + (1-y_{t+1})M + y_{t+1}k_t M \tag{5}$$

where $U(.)$ is an increasing and concave utility function. The variables w_t, c_t denote wealth and consumption at period t. The parameter β is the discount factor. The variable S_t is savings (if positive) or borrowing (if negative). r is the risk-free per-period interest rate.

The budget constraint in Equation 4 says that savings or borrowing in period S_t is the individual's wealth minus consumption minus the insurance payment. The law of motion for wealth in Equation 5 says that, the next period wealth w_{t+1} is any earnings on savings (or payments on borrowings), plus the residual value of the house depending on the realization of the disaster given by y_{t+1}.

The individual's problem can be rewritten as the following Bellman equation:

$$V(w_t) = \max_{c_t,k_t,S_t} U(c_t) + \beta E_t[V(w_{t+1})], \tag{6}$$

subject to the constraints (2)–(4), and the transversality condition given by

$$\lim_{t\to\infty} \beta^t w_t = 0. \tag{7}$$

The first-order conditions with respect to c_t and k_t show that the optimal consumption and insurance coverage are chosen according to:

$$U'(c_t) - \beta(1+r)E_t V'(w_{t+1}) = 0 \tag{8}$$

$$E_t[V'(w_{t+1})(-(1+r)p + y_{t+1}M)] = 0. \tag{9}$$

We obtain the following propositions for how the unconstrained household optimally chooses insurance coverage.

Proposition 1A. The unconstrained household optimally chooses full insurance coverage ($k = 1$) when insurance price is actuarially fair.

Proposition 1B. The unconstrained household is completely price inelastic ($dk/dp = 0$) when insurance price is actuarially fair.

See Appendix B for the detailed proof. Intuitively, full insurance is optimal at fair prices because it enables risk averse households to fully offset losses from the disaster shock, allowing them to smooth consumption over time and across states. Because households are risk averse, at actuarially fair prices, they will continue to maintain full coverage and not be sensitive to small price changes. This is a similar logic to what is obtained in static frameworks of Arrow (1963) and Mossin (1968).

2.2. Insurance Coverage Demand: Borrowing Constraints

We now consider the case where borrowers also face a borrowing constraint i.e.

$$S_t \geq s. \tag{10}$$

The parameter s denotes the minimum net wealth position allowed by credit markets. For example, when $s = 0$ borrowing is not possible, whereas $s = -\infty$ implies that there is no borrowing constraint.

With this constraint, the first-order conditions with respect to c_t and k_t are now:

$$U'(c_t) - \beta(1+r)E_t V'(w_{t+1}) - \lambda_t = 0 \tag{11}$$

$$E_t[V'(w_{t+1})(-(1+r)p + y_{t+1}M)] - \lambda_t p = 0 \tag{12}$$

where λ_t is the Lagrange multiplier on the borrowing constraint. We obtain the following propositions for how the constrained household optimally chooses insurance coverage:

Proposition 2A. When the borrowing constraint binds, the household optimally chooses lower than full insurance coverage ($k < 1$) under an actuarially fair insurance price.

Proposition 2B. When the borrowing constraint binds, the household is price elastic ($dk/dp < 0$) under an actuarially fair price.

See Appendix B for the detailed proof. Intuitively, the individual would have liked to borrow more to finance full insurance and smooth consumption across states and time but is unable to do so because the borrowing constraint starts to bind. With the borrowing constraint binding, the individual must consume more from her wealth, which leads insurance coverage to fall. This intuition

is similar to the logic in richer models of financing constraints and hedging, as in Rampini and Viswanathan (2013).

Combining Propositions 1A and 2A we obtain that at actuarially fair prices, a household with a binding borrowing constraint ("constrained household") chooses less insurance than the household that does not face borrowing constraints ("unconstrained household"). We now show that this result holds regardless of whether insurance prices are actuarially fair.

Proposition 2C. All else equal, the unconstrained household, u, chooses higher insurance coverage than the constrained household, c, ($k_t^u > k_t^c$) at all levels of insurance prices.

Intuitively, this result reflects the fact that the constrained household would have liked to borrow to finance more insurance coverage, but the binding constraint leads it to have lower coverage. See Appendix B for the detailed proof.

2.3. Demand for Debt

Given the importance of financial constraints for insurance coverage choices, we now derive predictions about how household debt optimally adjusts in response to changes in insurance prices. We obtain the following result:

Proposition 3A. All else equal, the unconstrained household's borrowing is increasing in insurance prices ($dS/dp < 0$) if the household has low coverage demand elasticity (dk/dp is small). The result holds at all levels of insurance prices, including when prices are above actuarially fair.

Proposition 3B. All else equal, the constrained household's borrowing is insensitive to insurance prices ($dS/dP = 0$). The result holds at all levels of insurance prices, including when prices are above actuarially fair.

See Appendix B for the detailed proof. Intuitively, households that are unconstrained and value insurance sufficiently (i.e. have low price elasticities) would only be willing to drop coverage up to a point. This implies that their total outlay towards insurance would rise, requiring them to finance rising premiums in part by increasing borrowing. To illustrate the point further, consider the extreme case where household is completely inelastic. In that case, the borrower would maintain the same coverage levels, and would need to finance the entire increase in prices by increasing borrowing and reducing consumption. On the other hand, if the household is perfectly elastic, they would drop

coverage in tandem with prices changes, without having to adjust debt or consumption.

2.4. Empirically Testing the Model's Predictions

The model has a number of testable implications. To summarize the propositions, the model predicts that that constrained households obtain less insurance coverage and have more elastic insurance coverage demand than unconstrained ones. Furthermore, the model suggests that constrained households' debt choices are less sensitive to changes in insurance prices than unconstrained ones.

To bring the model to the data, we note that in the model, households optimally choose insurance coverage k and savings/borrowing S, taking prices p and interest rates r as exogenously given parameters. Therefore, to test the model's predictions, we cannot simply rely upon the correlations in Section 1.3 – we require *exogenous* variation in insurance prices. Our empirical design in the subsequent sections will explain how we obtain such variation.

In addition, to explore the model-implied heterogeneity by financial constraints, we require a proxy for whether the borrower is financially constrained. We use a simple approach, classifying households as "financially constrained" if they have below-median income levels. This is motivated by Figure 5, which shows that higher income households have significantly higher debt capacity than lower income households. In fact, lower income households are nearly two-times more likely to have binding LTV or DTI constraints than higher income households, suggesting that income is strongly correlated with debt capacity.

In Section 3, we test Propositions 1A, 1B, 2A and 2B, and 2C, providing some of the first estimates in the literature for households' insurance coverage demand elasticities, and hetereogeneity in elasticities by whether the contract is (i) actuarially fair, and (ii) whether borrowing constraints are likely to bind. In Section 4, we consider Propositions 3A and 3B, estimating cross-price elasticities for household debt with respect to insurance prices.

3. Elasticity of Demand for Homeowners Insurance

In this section, we exploit plausibly exogenous variation in insurance prices to estimate households' coverage demand elasticities and test the model's predictions regarding the role of borrowing constraints in households' coverage choices.

3.1. Empirical Design

Estimating demand elasticities is challenging because of the classic simultaneity bias inherent in equilibrium price-quantity outcomes, whereby prices and quantities are jointly determined through the interaction of supply and demand forces. Thus, any observed price-quantity relationships inherently reflect both insurers' cost structures (supply elasticity) and consumers' willingness to pay (demand elasticity). To isolate a demand elasticity, we require exogenous supply shocks that shift the supply curve independently of demand conditions. With such variation, we can estimate demand elasticities and more directly test the model's predictions, which takes prices as exogenous.

We obtain arguably quasi-exogenous changes in insurance prices by instrumenting insurance prices in low price regulation (low friction) states with insurance prices in high price regulation (high friction) states. Our 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.[30] We implement a similar idea by using prices in *other* states as instruments with one key difference: we only use prices in high friction states, where insurers face significantly more regulatory hurdles to changing premiums in response to changes in expected losses or demand conditions (Oh et al., 2021). While there is still variation in premiums, the timing is driven by when regulators allow insurers to adjust premiums. Moreover, the premium change often differs substantially from what is implied by insurers' actuarial assessments. As a result, price changes in high friction states are unlikely to reflect demand conditions in low friction states, as we explain in more detail below.

We construct the instrument as follows. First, we limit the sample to a subset of high friction states, as in Oh et al. (2021). We compute the average insurance price for these states for each month and for each of the four credit score categories. 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:

[30] 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.

$$\text{First Stage: } \ln P_{i,c,z,m,t} = \omega + \eta \ln \bar{P}_{HF,c,m,t} + \theta' W_{i,c,z,m,t} + \nu_{i,c,z,m,t} \tag{13}$$

$$\text{Second Stage: } \ln \text{coverage}_{i,c,z,m,t} = \alpha + \beta \widehat{\ln P_{i,c,z,m,t}} + \gamma' W_{i,c,z,m,t} + \varepsilon_{i,c,z,m,t} \tag{14}$$

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-credit score - year fixed effects as well as controls for climate risk exposure, replacement cost, borrower income, building age, borrower age, and insurer-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 CoreLogic. With ZIP - credit score - year fixed effects, we are identifying our coefficient off of the month-to-month variation in a given year for prices in high friction states. In other words, our estimation compares two borrowers in the same ZIP, having similar credit scores that just happen to close their mortgages at two different points in time during the year. These borrowers face different prices in low friction states. Importantly, we isolate the price differences that arise from exogenous regulatory approval timings in high-friction states, rather than those that arise from shifts in local demand conditions in low-friction states.

The coefficient β can be interpreted as the elasticity of demand for coverage with respect to insurance prices. An estimated $\beta < 0$ would suggest that households are price elastic, or in the notation of the model, that $dk/dp < 0$.

3.2. Identifying Assumptions

For this instrument to be valid, we require both 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, and describe these identifying assumptions in more detail in Appendix C.

Relevance: A natural question is why insurance prices in low friction states are correlated with prices in high friction states. The key logic is that regulatory pricing constraints in high-friction states cause insurers to operate at persistent financial losses, which erodes their capital buffers. This leads them to raise prices in low-friction states. We describe this key idea in two steps. First, Figure D.6a shows that insurers' combined ratios are on average 100% in high friction states. That is,

the amount insurers pay out in claims and operational expenses is roughly equal to what they earn in insurance premiums. Moreover, losses are *persistent* in high friction states, with high combined ratios ranging between 95%-110% throughout the sample (Figure D.6b).[31]

Second, insurers typically operate in multiple states and offset exposures in high friction states by adjusting premiums in low friction states (Oh et al., 2021).[32] This cross-subsidy can be directly seen through insurers' combined ratios, with insurers in low friction states being on average highly profitable (combined ratios are close to 90%). This suggests that premiums are above actuarially fair rates in low friction states, and likely below or close to actuarially fair rates in high friction states. Table 2 shows the first stage of our IV regression, where insurances prices in low friction states are regressed on prices in high friction states, as in Equation 14. The R^2 is high, consistent with the high first-stage F-statistics we obtain (reported in the tables of the second stage), indicating that the relevance condition holds strongly (Stock and Yogo, 2005). The first-stage coefficient η is also highly statistically significant indicating that prices in high friction states strongly relate to prices in low friction states.

Monotonicity: The sign of the first stage coefficient does not affect the validity of the IV or our interpretation of the second stage results, as long as the first stage relationship is consistently monotonic and in the same direction. It is worth noting that either positive and negative signs would be consistent with a regulation-based explanation for the cross-state pricing relationship. The first stage could be negative, if price increases in high friction states *fully* alleviate insurers' financial constraints, allowing them to *lower* prices in low friction states. Another possibility, which is what we observe, is that the first stage is positive. That is, price increases in high friction states are associated with price increases in low friction states. Similarly, the original Hausman intuition is that the first stage coefficient should be positive, since supply shocks (such as increases in capital costs) would lead to an increase in prices in high friction states and in low friction states. It would be a challenge for the IV if we had these opposing forces leading to non-monotonicities.

However, empirically we find a consistently positive and strong first stage. It may be surprising

[31] Combined ratio is a widely used metric to measure insurers' underwriting profitability.

[32] Note that Oh et al. (2021) find that insurers cross-subsidize homeowners insurance across states to smooth financial frictions induced by losses. The idea that insurance prices may deviate from fundamentals has also been shown more broadly in other insurance markets across a range of studies (Koijen and Yogo, 2015, Froot and O'Connell, 1999, Ge, 2022, Gao et al., 2024).

that an increase in prices in high friction states *increases* the need to raise prices in low friction states. We argue this may arise if despite price increases, insurers continue to face persistent losses, suggesting that rate increases in high friction states are not sufficient to bring prices to actuarially fair levels or mitigate insurers' constraints fully.[33] Consistent with this, our evidence on the combined ratios in high friction states show that insurers remain unprofitable, despite price increases. Furthermore, we find that price increases in high friction states are actually associated with *increasing* price gaps relative to the prices implied by insurers' actuarial projections– higher price increases actually mean that insurers have even higher rate gaps, explaining the sign of the first stage. Thus in our setting, the cross-subsidy channel and the original Hausman intuition about supply shocks both push in the same direction. This, coupled with growing risks due to development in high-risk areas and climate losses, indicates that insurance contracts in high friction states likely continue to remain below actuarially fair, consistent with Figure D.6b. In other words, receiving higher prices in high friction states does not mitigate insurers' constraints and are associated with ever growing unprofitability, requiring insurers to increase prices further in low friction states.

Another question is whether our instrument might be weak for households facing single-state insurers, or only be valid for households facing multi-state insurers that cross-subsidize. However, this is unlikely for two reasons. First, even single-state insurers in low friction states may be affected by premiums in high friction states because they compete with multi-state insurers, and may be raising prices at the same time. Thus, it is not the case that only multi-state insurers are impacted by the cross-subsidy. Second, single state insurers have a small market share (10% in the US overall). Collectively, these factors help explain our strong first stage.

Exclusion Restriction: We also require that average prices in high friction states are unrelated to demand in low friction states, and that prices in high friction states only influence coverage in

[33]There is significant anecdotal evidence in support of this interpretation. 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.

low friction states through their impact on prices in low friction states. We believe this is likely to hold because the variation in our instrument is driven by the *timing* of when state regulators in high friction states choose to allow price increases. We argue this timing is unlikely to be related to demand conditions or risk factors in low friction states. Our empirical design compares two observably similar borrowers originating their mortgage in the same zip code in different months of the year. Because of this difference in timing, one household will obtain a different price than the other, with the difference being determined by when regulators allowed prices in high friction states to change rather than changes in demand or risks in low friction states.

There are two potential challenges. First, it would be a problem for our design if regulators in high-friction states consider demand conditions in low-friction states when making rate decisions. However, we argue that this is unlikely, as prices in high friction states appear unresponsive even to risk measures within those same states. To demonstrate this, we examine two measures of risk: (i) a dummy variable identifying properties in the top decile of expected losses; and (ii) lagged ZIP-code-level loss ratios, defined as total losses divided by total premiums, which is a standard measure of risk in the literature. Table D.2 shows that insurance prices in low friction states are significantly higher (by nearly 8%) for high risk properties that are in the top decile of expected annual losses. However, in high friction states, prices are completely insensitive to risk, with average annual losses having nearly a zero pass-through to prices. The results are similar using the second measure of risk, loss ratios. Table D.3 shows that premiums in low friction states are responsive to lagged loss ratios, but this is not the case in high friction states. In other words, unlike low friction states, price changes in high friction states occur for non-fundamental reasons. Given that prices do not even reflect actuarial risks in high friction states, it is highly unlikely that prices respond to demand shocks, especially from other states.

Furthermore, Oh et al. (2021) show that while prices in low friction states respond to conditions in high friction states, the opposite is not true. That is, prices in high friction states do not respond to conditions in low friction states. Moreover, they find that the spillovers from low to high friction states occur through a financial constraints channel, explaining why an insurer's pricing problem departs from simple state-by-state profit maximization. Specifically, losses in high friction states changes insurers' marginal costs leading to insurers wanting to raise prices everywhere. However,

insurers are only able to do so in low friction states, where regulatory frictions are lower and raising prices is relatively easier. In other words, the evidence shows that it is the supply curve that shifts not the demand curve.

Second, one might be concerned that exclusion may be violated due to a learning channel, if households in low friction states learn about their own risks by monitoring regulatory decisions or prices in high friction states directly. In this case, prices in high friction states would directly impact coverage in low friction states, not just through its effect on prices in low friction states, violating the exclusion restriction. We argue that this channel is unlikely. (i) While households may learn from climate events and losses that occurs elsewhere, these disaster events do not line up with the timing of price changes in high friction states. Consistent with this, Oh et al. (2021) show that prices in high friction states do not respond to losses in their own states, contemporaneously or even in the following year. (ii) It would also need to be the case that households observe the exact timing of regulatory decisions on *pricing* in other states, data on which is relatively difficult to obtain.[34] Furthermore, a learning hypothesis predicts that demand for coverage would rise as high friction prices go up, which if anything pushes against finding that demand curves slope downwards.

3.3. Estimated Demand Elasticities

We next report our estimates of household demand elasticities. In line with the conceptual framework in Section 2, we examine the extent to which the elasticity estimates vary by different proxies of household borrowing constraints. Proposition 2C in Section 2 predicts that constrained households have both lower coverage levels and more elastic coverage demand than unconstrained households. Consistent with this prediction, Figure 2b shows that lower income households have less coverage. We next test the prediction about heterogeneous demand elasticities. To do so, we use households' income and property values as proxies of borrowing constraints and estimate the elasticities of demand for these different subgroups.

Table 3 shows the main results from estimating Equation 14. We estimate that the average low income borrower is elastic with a statistically significant price elasticity estimate of -0.18. This means that a 1% increase in insurance prices leads to a 0.18% reduction in coverage. In contrast,

[34]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 regulatory decisions in high-friction states through media or social channels.

high income borrowers are inelastic: a 1% increase in prices leads to a negligible change, making low income borrowers substantially more elastic than high income ones. We find very similar results in Columns 3 and 4, when we split households by property values. Low property value households are substantially more elastic than high property value ones.

Discretionary coverage elasticities. While low income (low property value) households are more elastic than high income (high property value) households, they are not elastic in an absolute sense: most would describe an elasticity of -.18 as fairly inelastic. This inelasticity is also intuitive, since households may not have full discretion over the total coverage amount. In particular, Figure 4 shows that there is an informal rule-of-thumb where lenders require that households insure at least the loan amount. Thus, households likely only have discretion over coverage amounts that exceed the lender minimum. We therefore consider how elastic low income households are in the discretionary component of coverage, which we define as coverage in excess of the loan amount. To do so, we limit the sample to low income borrowers with coverage greater than or equal to the loan amount, as well impose the condition that replacement costs exceed the loan amount. We then use our same instrumental variable specification to estimate the elasticity of demand over the discretionary component of coverage. Column 2 of Table 4 shows the results: low income borrowers are highly elastic over the discretionary component of coverage. A 1% increase in prices leads to a 2.4% reduction in discretionary coverage. These results suggest that as prices grow, household equity would be significantly more exposed as households continue to reduce their discretionary coverage.

Results with Debt-to-Income ratio (DTI) thresholds: So far, we have used income and wealth as proxies to measure household financial constraints. One potential concern with these measures is that they may also capture other household characteristics unrelated to financing frictions. To provide additional evidence in support of the financial-constraints channel, we turn to an alternative proxy based on borrowers' debt-to-income ratios. Specifically, we estimate elasticities separately for two subsamples: borrowers with DTI ratios below 30% and those above 30%. A DTI threshold of 30% represents a key underwriting cutoff in the mortgage market, above which lenders become substantially more reluctant to extend credit. This sharp institutional feature makes DTI a useful measure of the tightness of borrowing constraints. Table D.7, column 1 shows that households facing relatively slack borrowing constraints (DTI below 30%) exhibit inelastic coverage elasticities.

In contrast, households with more binding borrowing constraints (DTI above 30% exhibit relatively higher elasticities in coverage overall (column 2) and particularly high elasticities in the discretionary component of coverage (column 3). These results reinforce our interpretation that mortgage credit access plays a central role in shaping households' insurance choices.

Robustnesss: We show that these estimates are unlikely to be explained by the actuarial fairness of the insurance contract. We obtain variation in actuarial fairness by looking at high friction states. Insurance contracts in high friction states are on average more likely to be below or close to being actuarially fair, in particular much more so than contracts in low friction states.[35] We show in Table 1 that households in high friction states are not fully insured (with coverage ratios at roughly 70%), which is inconsistent with Proposition 1A. Furthermore, Table D.5 reports the demand elasticities for high friction states using an OLS estimation.[36] Given that in these states price changes are unrelated to shifts in expected losses and demand conditions, an OLS specification can help recover the demand elasticities because insurers cannot simultaneously adjust in response to demand. The results in Table D.5 show that even with prices below or at actuarial fair levels, households are elastic. Because demand elasticities are similar in high and low friction states, it is unlikely that the broader patterns of the estimated demand elasticities are driven by the generosity of the insurance contract.

Overall, the fact that lower income borrowers are more price elastic strongly suggests that households with fewer financial resources respond to rising premiums by dropping coverage. It also suggests that, all-else-equal, households with more financial resources are better able to accommodate price increases, thus they have to drop coverage to a lesser degree. In the next section, we examine the forces that might explain these differences in elasticities by evaluating households' debt choices in response to plausibly exogenous shocks to insurance prices.

[35] This comes from the data on combined ratios, which are a measure of insurer profitability. Loss ratios in high friction states are persistently close to 1, while in low friction states they are under 1 (Figure D.6b).

[36] Because the price of insurance in high-friction states adjusts quasi-exogenously, we argue that variation in the price is not driven by fundamentals of risk or demand. That is, variation in price reflects variation in how generous the contract is, even though it remains below the actuarily-fair price.

4. Impact of Insurance Prices on Household Debt

The conceptual framework shown in Section 2 shows that debt plays an important role in supporting households' purchases of insurance coverage. We now test the model's predictions empirically.

4.1. Econometric Concern

Proposition 3A implies that households with spare debt capacity can finance large insurance price shocks by increasing their borrowing levels at origination. In addition, Proposition 3B implies that more (less) constrained households' debt choices are less (more) sensitive to insurance prices. At the extreme, when the constraint binds, debt choices are fully insensitive to insurance prices. To test these predictions, we need to estimate a cross-price elasticity of demand of mortgage debt with respect to insurance prices. The ideal experiment would randomly assign different insurance prices across households and examine how their demand for credit changes.

In this setting, the OLS regression of a household's mortgage debt on its insurance prices will likely be biased for a number of reasons. One key omitted variable of concern is the borrower's unobservable risk type– borrowers with high idiosyncratic default risk, or high climate risk, may both have higher insurance premiums and demand higher debt. In this case, the OLS estimate would be higher than the true treatment effect of higher premiums, because it would also capture the effect of riskier borrowers' demand for more mortgage debt. On the other hand, borrowers with higher unobservable risk types may not actually be able to borrow more as lenders may screen more tightly or ration credit (Sastry, 2025). If so, higher premiums would be associated with smaller loan size. In other words, the OLS estimate would be lower than the true treatment effect of higher premiums, because it would also capture lender screening and credit rationing. The intuitions are similar if households and lenders learn about disaster risk probabilities from insurance prices – households would have incentives to increase debt, and lenders would have incentives to ration debt. This discussion highlights that the selection effect of unobservable risk does not present a clear prediction on what the sign of the bias should be because the demand effect and the supply effect go in opposite directions. What is clear, is that a simple correlation between insurance premiums and the level of debt chosen would also capture such selection effects and therefore would be difficult to interpret.

A related concern is that insurance companies can endogenously react to both changing unobserv-

able risk types and changing loan demand. For example, insurers may know that lenders have lax credit standards in some areas, which can enable households to buy more coverage because they are less cash-constrained; they therefore have incentives to adjust prices optimally given this knowledge. Such endogenous supply adjustments would also bias the estimate, because the correlation with premiums and debt levels would also reflect insurance supply elasticities.

4.2. Instrumental Variables Specification

To address these endogeneity issues, we employ the same instrumental variables strategy as in Section 3. In this strategy, we exploit variation in insurance prices that are driven by changes in high friction states, where regulatory frictions prevent insurers from adjusting prices frictionlessly in response to fundamental shocks (e.g., demand or risk). Importantly, in this setting, the fact that insurers in high friction states cannot adjust prices in response to demand shocks also implies that they do not respond to credit conditions in low friction states. We can therefore use insurance prices in high friction states to instrument for insurance prices in low friction states. Thus, we consider the following 2SLS specification:

$$\begin{aligned} \text{First Stage: } \ln P_{i,c,z,m,t} &= \omega + \eta \ln \bar{P}_{HF,c,t} + \theta' W_{i,c,z,m,t} + \nu_{i,c,z,m,t} \\ \text{Second Stage: } \ln LTV_{i,c,z,m,t} &= \alpha + \beta \ln \widehat{P_{i,c,z,m,t}} + \gamma' W_{i,c,z,m,t} + \varepsilon_{i,c,z,m,t} \end{aligned} \tag{15}$$

The first stage, instrument, and control variables are the same as defined in Section 3. The coefficient η represents the pass-through of premium shocks in high friction states to low-friction states, and the estimates are reported in Table 2. In this specification, the second stage uses measures of household debt as outcome variables. The coefficient estimate of interest is β, which is the cross-price elasticity of debt with respect to insurance prices and captures how measures of debt change as insurance prices change. We note that this parameter will capture both the effects of credit demand and credit supply. Households may wish to adjust loan-to-value ratios in response to prices, but in equilibrium they can only do so if lenders agree.

Measures of debt. We measure household debt using their loan-to-value ratios (LTVs). LTV ratios are the ratio of the mortgage loan amount to the house price, and are commonly used by lenders to measure household leverage.[37] The benefit of using this measure is that there is no mechanical

[37]The denominator is either the transaction sale price or the appraised value.

relationship between rising insurance premiums and LTV ratios, which may be a challenge with other common measures (e.g., debt-to-income ratios).

4.3. Identifying Assumptions

The key idea of the IV approach is that we seek to obtain exogenous variation in the supply of insurance so that we can estimate the causal effects of insurance prices on household credit demand. From that standpoint, we require similar identifying assumptions as in Section 3. For relevance, we require that prices in high friction states are relevant for prices in low friction states. Similarly, we require that the first stage relationship is monotonic and in the same direction. However, our exclusion restriction now requires the following assumptions: (1) that prices in high friction states are exogenous to insurance demand in low friction states, and (2) that prices in high friction states are exogenous to lending markets in low friction states. Under these two exogeneity assumptions, we can interpret the IV estimates as the causal effect of insurance prices on household debt.

However, we are unable to separately identify the component of β as reflecting credit demand (household-driven) or credit supply (lender-driven) elasticity. Households may want to increase debt in response to exogenous changes in prices; however, they will only be able to do so in equilibrium if lenders are willing to allow households to take on higher leverage. Thus, our estimate will reflect the net effect of insurance prices on overall credit, driven by the cumulative net effect of changes in both credit demand and supply, i.e. the overall net causal effect of prices on LTV ratios. Section 4.5 provides more evidence on which force is more likely to be driving our results.

4.4. Results

We next report the estimates of households' cross-price debt elasticities. To consider how debt adjustment differs by the type of borrower, we examine heterogeneous effects by borrower income and property value. Proposition 3A of our model predicts that unconstrained households' debt choices are sensitive to insurance prices. In contrast, Proposition 3B predicts that constrained households' debt choices are insensitive to insurance prices. To motivate why these measures represent household constraints, we provide evidence in Figure 5, which shows the fraction of households that are at a binding debt to income or loan to value ratio levels using common thresholds used in the mortgage literature. This shows that low income households are substantially more likely to be at binding debt limits than high income borrowers.

Table 5 shows the results of the IV regressions. Columns (1) and (2) show the results when splitting the sample by income; Columns (3) and (4) show the results when splitting the sample by house prices. Consistent with the model's predictions, we find that loan-to-value ratios significantly increase with exogenous changes in insurance prices, implying that price shocks in insurance markets affect household debt choices. We note that this effect is particularly pronounced for high income borrowers. In column (1), we find that a 1% increase in insurance prices leads to a 0.7% increase in LTV ratios for low income households. In contrast, for high income borrowers, in Column (2), we find that a 1% increase in prices leads to a 1% increase in loan-to-value ratios. In other words, both high and low income borrowers increase debt in response to insurance prices, but interestingly the adjustment by high income households is much larger. We find very similar results in Columns (3) and (4), when we split households by property values. In Table D.9, we show that are estimates are robust to controlling for the interest rate on the mortgage.

Recall that higher income households also exhibited lower coverage elasticities than lower income ones (Table 3). Collectively, these results are consistent with the model's predictions about how borrowing constraints influence the joint choice of insurance coverage and debt: lower income households (that are constrained) drop coverage in response to rising premiums, while higher income households (that are unconstrained) raise debt and do not drop coverage.

4.5. Credit demand and supply

There are two possible interpretations of the opposite elasticities we observe in insurance and debt markets across high and low income borrowers. The first interpretation is that both low income and high income households seek out more credit in order to finance insurance purchases, and therefore we observe a positive debt elasticity. However, lenders are only willing to extend more credit for high income households. This could be the case because high income borrowers are more likely to be at lower LTV ratio to begin with, so additional credit extensions may not meaningfully change the risk of the mortgage enough to concern lenders (Figure 5). On the other hand, lenders may be more wary of extending credit to low income households, because such households are already more likely to obtain higher LTV ratio mortgages. For these borrowers, additional credit extension could greatly increase the risk of the loan, since it is well known that default risk is non-linear in LTV ratios and increase significantly above certain thresholds (such as 80%, 90%, etc.). Lender credit rationing

of low income households could thus explain the disparate adjustments of LTV ratios, making the seemingly disparate results due to differences in *credit supply* adjustments.

On the other hand, another interpretation could be that the disparate patterns emerge because high income households value insurance whereas low income households do not. This could occur because, for example, low income households have different beliefs about climate change or their own risks, preferences, financial literacy, which in turn lower their overall willingness-to-pay for insurance. Under this explanation, the lower elasticities for LTV ratios along with the high elasticities in insurance reflect the fact that low income households simply value insurance to a lesser degree, and thus their credit demand does not change since they do not seek additional credit to finance insurance purchase. Under this interpretation, the disparate results in LTV ratios thus reflect differences in *credit demand* by households.

To help shed light on these two mechanisms, we consider a test to understand the degree to which low income households value insurance, distinguishing between whether the elasticities reflect a lower *willingness-to-pay* or alternatively a lower *ability-to-pay* that is a function of lender credit rationing. In this test, we exploit the panel nature of our dataset, looking past origination to see insurance choices over a longer period. Because insurance is an annual contract, insurance coverage terms can change even after origination. To understand the role of ability-to-pay, we look at households which start out as low income, but then transition to high-income within the first 3-5 years of the mortgage. We then examine how their demand elasticities change when they transition income groups, and include a loan fixed effect so that we compare the same borrower over time. The idea of this test is that other fundamental characteristics of the borrower, such as their beliefs, risk aversion, and literacy, are unlikely to change in such a short period of time, and can thus be captured by the loan fixed effect. We also include a loan age fixed effect, so that we compare households at the same point of the mortgage life cycle.

We implement this test by adjusting our specification to regress coverage on both the instrumented price, as well as the interaction of the instrumented price with a dummy variable which indicates that the household transitioned from low income to high income. We include a loan fixed effect, loan age fixed effect, year fixed effect, as well as the time-varying current FICO score of the borrower as controls. Table 6 shows the results: low income households who do not transition to high income

remain elastic over their coverage, with elasticities that roughly match the overall elasticities we find for low income borrowers. However, households that transition to high income status become *inelastic* in their coverage demand. These results show that low income households indeed value insurance, and are consistent with the idea that binding LTV constraints or limited access to credit restrict their ability-to-pay for insurance.

Interpretation of Magnitudes: To interpret these magnitudes, we consider the following back-of-the-envelope calculation. We ask if prices were to continue growing at historical rates moving forward, what fraction of currently unconstrained households would start to hit binding loan-to-value limits. Our estimates imply that between 17% to 31% of previously unconstrained would hit LTV constraints. If so, they would also be forced to reduce coverage substantially in the same way as constrained households.

Denial Rates: In addition, we also consider other proxies of credit supply directly, including lender denial rates. We use publicly available administrative data from the Home Mortgage Disclosure Act. After 2018, this dataset began to include information on the loan-to-value ratios of mortgage applications, along with other rich information about borrower characteristics. We can therefore evaluate whether lenders are more likely to deny high-LTV mortgages of lower income borrowers than higher income borrowers when insurance prices are high. This interpretation would support the idea that both groups request high-LTV loans in response to high insurance prices, but the heterogeneity in LTV adjustments we find in Table 5 reflect lender rationing of high-LTV loans for lower income borrowers. Because of data limitations, we cannot implement our two-stage-least-squares procedure directly: specifically, we do not observe the month of application, the borrower credit score, or the insurance price associated with the loan application. Instead, we consider the reduced form of the IV, regressing a dummy for whether the mortgage is denied on the average prices in high friction states in that year, and including rich controls and fixed effects. Table D.8 shows the results. We see that lenders are more likely to deny high-LTV loans for low income borrowers when insurance prices in high friction states are high; however, there are no adjustments in credit supply for high income borrowers. This provides additional evidence supporting a credit supply interpretation of the results.

5. Conclusion

This paper provides some of the first loan-level national estimates on household under-insurance by connecting state-of-the-art mortgage, climate, and insurance databases at a property level. Although nearly all mortgage borrowers hold homeowners insurance policies, we we document that, on average, households do not have *enough* coverage. Coverage ratios are the lowest for borrowers with low incomes and low credit scores. Importantly, we find that lender exposures are largely protected, leaving households' equity particularly vulnerable to uninsured losses. Using a natural experiment exploiting regulatory shocks to insurance prices, we show that rising premiums cause low income households to drop coverage, while high income households maintain coverage and increase their debt, taking out mortgages with higher loan-to-value ratios. We provide suggestive evidence that the disparate results across high and low income households connect to lender credit supply choices, and their willingness to extend credit in order to finance household insurance coverage. A natural policy implication is that more transparent insurance price information during the housing search process can help households be more informed in advance. Accurate pricing information is difficult to obtain until the household has already put down a deposit and spent significant resources to find a home; this makes it difficult and costly for households to adjust on other margins besides debt and coverage. Taken together, our findings suggest that rising insurance premiums are simultaneously increasing both household leverage and under-insurance, forcing lenders to navigate a trade-off between credit and insurance risk within the mortgage contract.

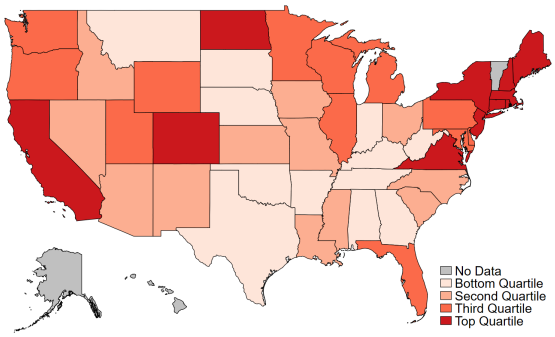
Tables and Figures

Figures

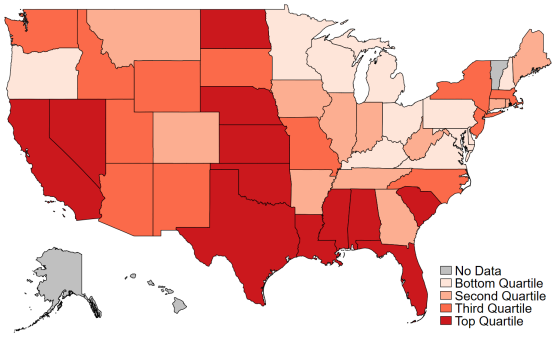
Figure 1: Spatial Distribution of Coverage, Risk and Income Across the U.S.

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, obtained by multiplying structure size and local construction costs from the R.S. Means Company. 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 1.2.

(a) Coverage Ratio



(b) Climate risk



(c) Borrower Income

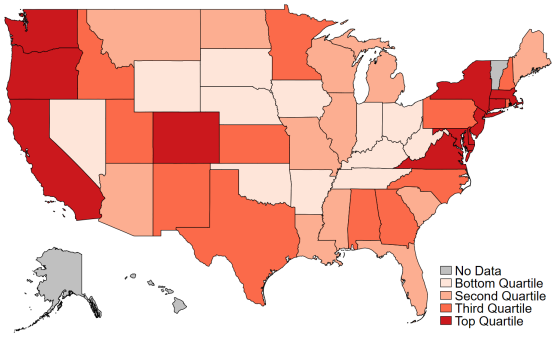
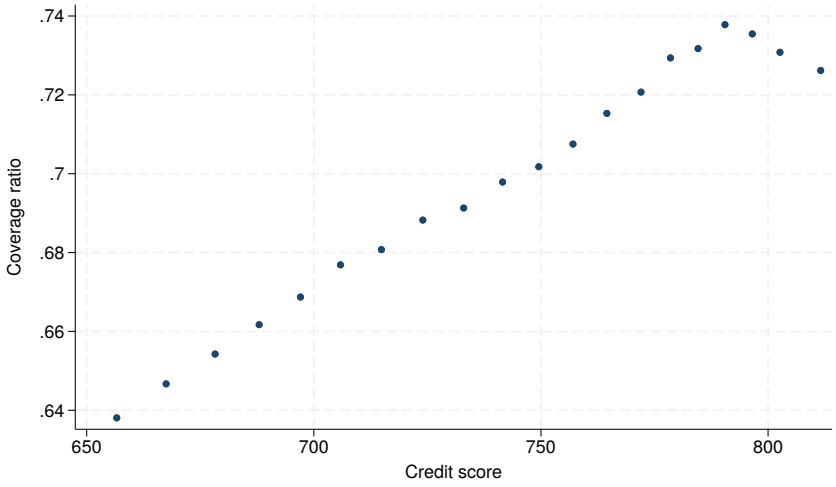


Figure 2: Coverage Ratios by Borrower Credit Score and Income

This figure shows a binned scatterplot of coverage ratios by borrower credit score (panel a) and income (in thousands) (panel b). 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 1.2.

(a) Borrower Credit Score



(b) Borrower Income

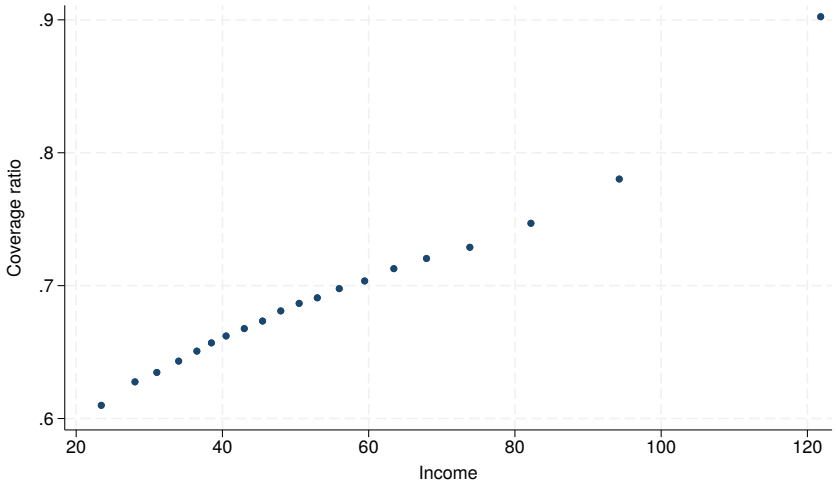
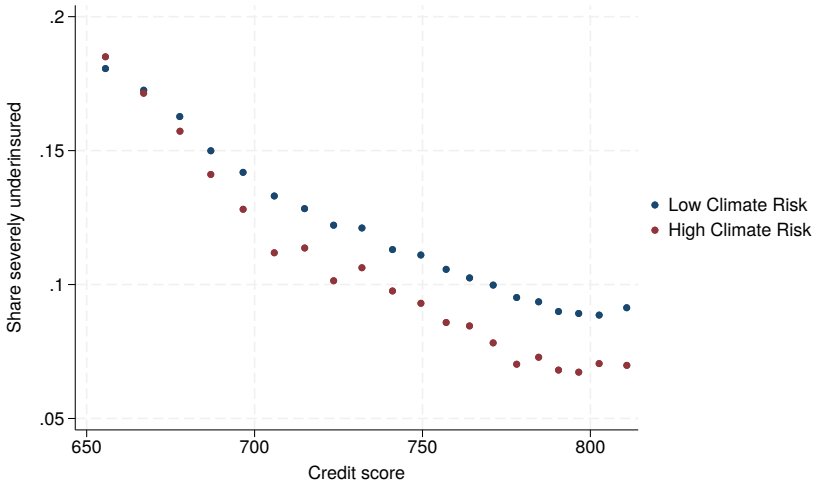


Figure 3: Severe Under-insurance by Borrower Income, Credit Score, and Climate Risk

This figure shows a binned scatter plot of the share of severely under-insured borrowers (coverage ratio of less than 50%) by borrower's credit score (panel a) and by 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, obtained by multiplying structure size and local construction costs from the R.S. Means Company. The figure is based on the main analysis sample, outlined in Section 1.2.

(a) Borrower Credit Score



(b) Borrower Income

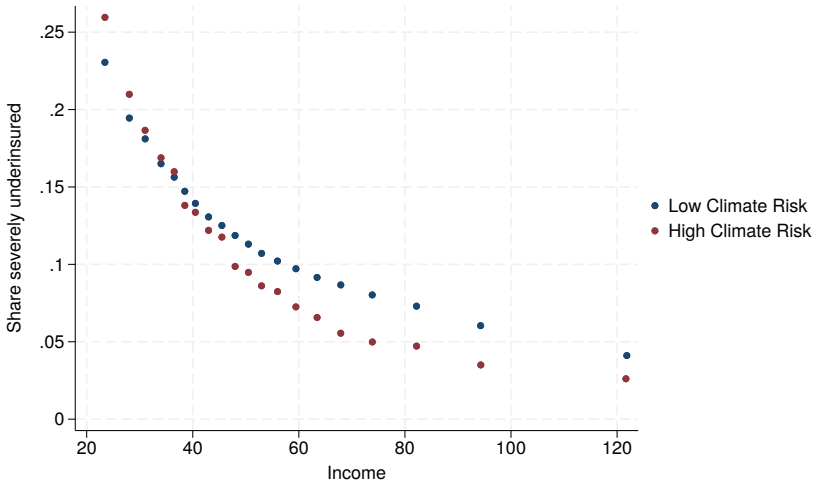


Figure 4: Discretionary and mandated coverage levels

This figure shows a histogram of the ratio of coverage amounts to loan amounts at origination. The figure is based on the main analysis sample, outlined in Section 1.2 which we limit to low income borrowers with replacement costs that exceed property values. Low income borrowers are households with income below the sample median.

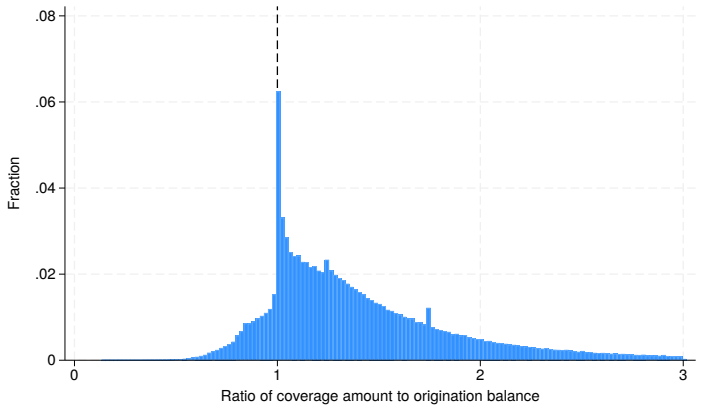
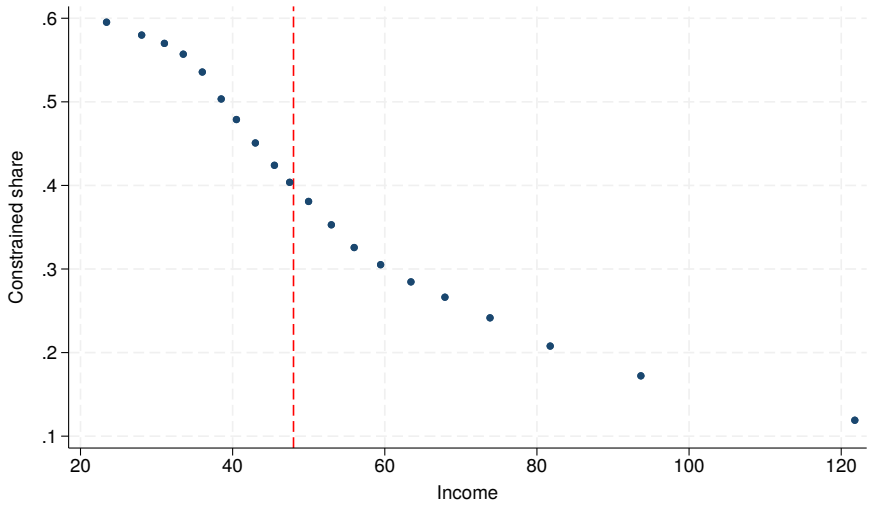


Figure 5: Mortgage Market Constraints and Borrower Income

This figure shows a binned scatter plot of borrower income (in thousands) and a measure of mortgage market borrowing constraints. We define borrowers to be constrained if they are either at the binding LTV constraint (LTV exceeds 95%) or at the binding DTI constraint (DTI exceeds 45%). The figure shows the constrained share, i.e. the fraction of borrowers in a given income ventile that are constrained. The red vertical line indicates the median income in the sample.



Tables

Table 1: Summary Statistics

This table presents the summary statistics for the main analysis sample, outlined in Section 1.2. 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, obtained by multiplying structure size and local construction costs from the R.S. Means Company. Share underinsured refers to fraction of borrowers with insurance coverage ratios below 70% and share severely underinsured refers to fraction of borrowers with insurance coverage ratios below 50%. Insurance price is the monthly insurance premium per $100,000 of coverage. Share with high DTI refers to the fraction of borrowers with debt-to-income ratios above 31%. Share with high LTV refers to the fraction of borrowers with LTV ratios above 80%. Interest rate is the rate on the mortgage loan. The market interest rate is the monthly average of the primary mortgage market survey (PMMS) 30-year mortgage rate. 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
Share with high climate risk	0.096	0.29	0.089	0.28	0.11	0.31
Coverage Amount	275152.4	140191.1	272149.1	140297.1	278880.7	139969.7
Coverage Ratio	0.69	0.20	0.69	0.20	0.68	0.19
Share underinsured	0.61	0.49	0.61	0.49	0.61	0.49
Share severely underinsured	0.12	0.33	0.12	0.32	0.13	0.33
Annual Insurance Premium	1123.6	605.3	1138.3	632.2	1105.4	569.8
Insurance Price	38.1	20.7	39.1	21.9	36.9	19.1
Annual Income	54.0	23.5	54.2	23.6	53.7	23.3
Building Size (Sq. Feet)	1968.7	817.2	1984.2	823.3	1949.5	809.3
Loan Amount	233331.4	136840.8	234202.9	137573.6	232249.5	135918.0
DTI Ratio	35.9	9.52	35.7	9.59	36.2	9.42
LTV Ratio	86.9	13.0	86.7	13.2	87.2	12.6
Credit Score	733.2	51.7	734.0	51.9	732.2	51.4
Share with high DTI	0.69	0.46	0.68	0.47	0.70	0.46
Share with high LTV	0.66	0.47	0.66	0.47	0.67	0.47
Interest Rate	4.07	0.63	4.06	0.62	4.09	0.63
Market Interest Rate	3.95	0.49				
N	3271262		1811792		1459470	

Table 2: 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 13. 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 1.2. All specifications include ZIP-credit score-year fixed effects. Specifications include 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 income-year fixed effects, and borrower age fixed effects. Standard errors are clustered at the county level. The sample is as outlined in Section 1.2, but limited to the sample of loans in low friction states. We show the first stage for each sub-sample considered in the paper: above and below median income (columns 1-2), and above and below the median house price (columns 3-4). Note: *p<0.1; **p<0.05; ***p<0.01.

	$\ln P_{izcmt}$			
	Low Income	High Income	Low House Price	High House Price
	(1)	(2)	(3)	(4)
$\ln P_{HF,cmt}$	0.0508***	0.0415***	0.0489***	0.0433***
	(0.00247)	(0.00346)	(0.00277)	(0.00297)
Observations	821096	815162	821886	811698
Zip-FICO-Year FE	Y	Y	Y	Y
Income-Year FE	Y	Y	Y	Y
Insurer-Year FE	Y	Y	Y	Y
Property Controls	Y	Y	Y	Y
Borrower Controls	Y	Y	Y	Y
Adj. R^2	0.702	0.697	0.701	0.710

Table 3: Second-Stage Results: Coverage Demand Elasticities

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 14. We present the results for two sub-samples: those with above versus below median income (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 1.2. 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 2. Standard errors are clustered at the county level. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln coverage$_{izcmt}$			
	Low Income	High Income	Low House Price	High House Price
	(1)	(2)	(3)	(4)
$\widehat{\ln P_{iczmt}}$	-0.181***	-0.0142	-0.132***	0.0136
	(0.0346)	(0.0558)	(0.0334)	(0.0423)
ZIP-FICO-Year FE	Y	Y	Y	Y
Income-Year FE	Y	Y	Y	Y
Insurer-Year FE	Y	Y	Y	Y
Property Controls	Y	Y	Y	Y
Borrower Controls	Y	Y	Y	Y
Number of Observations	821096	815162	821886	811698
Adjusted R-squared	0.524	0.526	0.403	0.434
First-Stage F	422.6	143.4	310.6	212.2

Table 4: Second-Stage Results: Discretionary Coverage Demand Elasticities

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 14, except we consider in Column 2 the discretionary component of coverage. We limit the sample to low income households, those with below-median incomes, and the sample of low friction states. In Column 1, the dependent variable is log of coverage. In Column 2, the dependent variable is log of the discretionary 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 1.2. 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 2. Standard errors are clustered at the county level. Note: *p<0.1; **p<0.05; ***p<0.01.

	(1) ln coverage$_{izcmt}$	(2) ln discretionary coverage$_{izct}$
$\widehat{\ln P_{iczmt}}$	-0.181***	-2.428***
	(0.0346)	(0.247)
ZIP-FICO-Year FE	Y	Y
Income-Year FE	Y	Y
Insurer-Year FE	Y	Y
Property Controls	Y	Y
Borrower Controls	Y	Y
Number of Observations	821096	639522
First-Stage F	422.6	334.6

Table 5: Second-Stage Results: Mortgage Debt Elasticities

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 15. We present the results for two sub-samples: those with above versus below median income (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 loan-to-value ratio. 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 1.2. 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 2. Standard errors are clustered at the county level. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln LTV $_{izcmt}$			
	Low Income	High Income	Low House Price	High House Price
	(1)	(2)	(3)	(4)
$\widehat{\ln P_{iczmt}}$	0.718***	1.176***	0.699***	1.167***
	(0.0452)	(0.105)	(0.0485)	(0.0915)
ZIP-FICO-Year FE	Y	Y	Y	Y
Income-Year FE	Y	Y	Y	Y
Insurer-Year FE	Y	Y	Y	Y
Property Controls	Y	Y	Y	Y
Borrower Controls	Y	Y	Y	Y
Number of Observations	821096	815162	821886	811698
Adjusted R-squared	-1.978	-3.582	-1.905	-3.515
First-Stage F	422.6	143.4	310.6	212.2

Table 6: 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 log 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 1.2. 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 2. Standard errors are clustered at the county 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

References

Adelino, M., A. Schoar, and F. Severino (2016, July). Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class. *Review of Financial Studies 29*(7), 1635–1670.

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.

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.

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. Keys (2022). Moral hazard during the housing boom: Evidence from private mortgage insurance. *Review of Financial Studies*.

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, and J. Weill (2024). Pricing protection: Credit scores, disaster risk, and home insurance affordability. *SSRN Working Paper Number 5026600*.

Bolton, P., H. Chen, and N. Wang (2011, October). A Unified Theory of Tobin's q, Corporate Investment, Financing, and Risk Management. *Journal of Finance 66*(5), 1545–1578.

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.

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.

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 A. Paciorek (2017, February). The interest rate elasticity of mortgage demand: Evidence from bunching at the conforming loan limit. *American Economic Journal: Economic Policy 9*(1), 210–40.

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*.

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., 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.

Gao, J., S. Ge, L. D. W. Schmidt, and C. Tello-Trillo (2024). How Do Health Insurance Costs Affect Low- and High-Income Workers? Working papers.

Ge, S. (2022). How do financial constraints affect product pricing? Evidence from weather and life insurance premiums. *The Journal of Finance 77*(1), 449–503.

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*.

Glaeser, E. L. and A. Shleifer (2003). The inefficiency of law. *The Quarterly Journal of Economics 118*(2), 387–436.

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

Gollier, C. (2003). To insure or not to insure?: An insurance puzzle. *The Geneva Papers on Risk and Insurance Theory 28*(1), 5–24.

Greenwald, D. (2018). The mortgage credit channel of macroeconomic transmission. *MIT Sloan Research Paper No. 5184-16*.

Gupta, A. and C. Hansman (2021, 04). Selection, leverage, and default in the mortgage market. *The Review of Financial Studies 35*(2), 720–770.

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*.

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.

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.

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.

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.

Liu, Y. and R. J. Myers (2012). The dynamics of insurance demand under liquidity constraints and insurer default risk. *IFPRI Discussion Paper 01174*.

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

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.

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

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

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

Stock, J. and M. Yogo (2005). Testing for weak instruments in linear iv regression. *Identification and Inference for Econometric Models*, 80–108.

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 D.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. Figure D.1 provides a high-level schematic of how we combine loan-level information on mortgage, insurance, and borrower characteristics, with property-level databases. 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 1.2, 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 the evolution of premiums and coverage over the entire same period. Figure D.2b shows that annual premiums have increased substantially. The median household had a premium of about $800 in 2013, which increased to roughly $1500 in 2023 at the end of the sample (87% increase). The increase at the 75th percentile is from $1200 to about $2000 (67% increase), rising faster than the rate of inflation. These patterns are consistent with other papers showing rising premiums in the recent period (Mulder and Keys, 2024). During the same period, we also see coverage levels increase, but at a much slower rate (Figure D.2a). The median household increased coverage from $217,000 in 2013 to $284,000 in 2023 (31% increase). Similarly households at the 75th percentile increased coverage from $295,000 to $375,000 (27% increase).

The patterns in the time series mask significant variation in the cross-section. Figure D.3 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 1, 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 1c) are also states with the highest coverage ratios (above 73% from Figure 1a). 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 Coverage Ratios and Insurance Prices

We next provide examine which types of loan-level variables drive the variation in coverage ratios and insurance prices. We consider the following broad categories of variables. zip-year fixed effects, housing (rebuild costs), climate, mortgage contract, and insurance contract. In Figure D.4a, we find that zip-year fixed effects alone explain 40% of the variation in coverage choices. Adding our proxy for rebuild costs increases the R^2 by another 30%. Rebuild cost's high explanatory power and its high correlation with house prices further help validate our use of this proxy. We also find that mortgage variables have high explanatory power: together with zip-year fixed effects, they increase the R^2 to 77%. Interestingly, climate variables or insurance-contract variables do not explain as much of the variation. Finally, the full set of variables explain 85% of the variation in coverage choices. Taken together, the results suggest that mortgage variables, which include several types of borrower characteristics that proxy for credit constraints, play an important role in explaining the variation in coverage choices, a hypothesis we explore in great detail in the paper. We use this result to motivate our detailed look at the relation between borrower characteristics, mortgage contract, and insurance coverage choices.

In Figure D.4b 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 D.5 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 3

B. Proofs

B.1. First Order Conditions

We can derive the first-order conditions (FOCs) with respect to c_t and k_t by differentiating the Bellman equation in Equation 6 subject to the constraints and the transversality condition.

FOC w.r.t. c_t:

$$\frac{\partial}{\partial c_t}\left[U(c_t) + \beta E_t V(w_{t+1})\right] = 0$$

$$U'(c_t) + \beta(E_t V'(w_{t+1})\frac{\partial w_{t+1}}{\partial c_t}) + \lambda_t(\frac{\partial S_t}{\partial c_t}) = 0$$

$$U'(c_t) - \beta(1+r)E_t V'(w_{t+1}) - \lambda_t = 0$$

This yields the FOC in Equation 8.

FOC w.r.t. k_t:

$$\frac{\partial}{\partial k_t}\left[U(c_t) + \beta E_t V(w_{t+1})\right] = 0$$

$$E_t\left[V'(w_{t+1})\frac{\partial w_{t+1}}{\partial k_t}\right] + \lambda_t(\frac{\partial S_t}{\partial k_t}) = 0$$

$$E_t\left[V'(w_{t+1})(-(1+r)p + y_{t+1}M)\right] - \lambda_t p = 0$$

This yields the FOC in Equation 9. We represent this optimality condition with the function $G(c_t, k_t)$.

B.2. Propositions

Lemma. $G(c_t, k_t)$ is decreasing in k_t.

Proof of Lemma. Differentiating $G(c_t, k_t)$ with respect to k_t we get:

$$\frac{\partial G}{\partial k_t} = E_t\left[V''(w_{t+1}) \cdot \frac{\partial w_{t+1}}{\partial k_t} \cdot (-(1+r)p + y_{t+1}M)\right]$$

Note that $\frac{\partial w_{t+1}}{\partial k_t} = (-(1+r)p + y_{t+1}M)$. Substituting this in we get:

$$\frac{\partial G}{\partial k_t} = E_t\left[V''(w_{t+1}) \cdot (-(1+r)p + y_{t+1}M) \cdot (-(1+r)p + y_{t+1}M)\right]$$

This derivative is negative because (i) $V''(w_{t+1})$ is negative due to the concavity of the value function (risk aversion). (ii) $(-(1+r)p+yM)^2$ is a squared term. Thus, $G(c_t, k_t)$ is indeed decreasing in k.

Proof of Proposition 1A. We first show that $k_t = 1$ is a solution to the optimality condition in Equation 9, and then show that it is the unique solution.

When there is no borrowing constraint ($s = -\infty$), we have the Lagrange multiplier on the borrowing constraint $\lambda_t = 0$. This implies that Equation 9 becomes

$$G(c_t, k_t) = E_t\left[V'(w_{t+1})(-(1+r)p + y_{t+1}M)\right] = 0. \tag{B.16}$$

Full insurance is a solution. Let's assume that $k_t = 1$ then

$$w_{t+1} = (1+r)(w_t - c_t - p) + M. \tag{B.17}$$

Note that Equation B.17 holds with probability 1. That is, wealth is period $t+1$ is invariant to the loss event y as the individual is fully insured.

Substituting Equation B.17 into Equation B.16 we get:

$$V'(w_{t+1})E_t\left[-(1+r)p + y_{t+1}M\right] = 0,$$

where $V'(w_{t+1})$ comes out of the expectation as there is no uncertainty.

For actuarially fair insurance, we have that $p = \phi M/(1+r)$, where ϕ is the probability of loss given by $E_t(y_{t+1}) = \phi$. Thus, we have:

$$V'(w_{t+1})\left[-(1+r)(\phi M/(1+r)) + \phi M\right] = 0,$$

$$V'(w_{t+1})\left[-\phi M + \phi M\right] = 0,$$

$$V'(w_{t+1}) \cdot 0 = 0.$$

This shows that FOC is satisfied when $k_t = 1$, showing that $k_t = 1$ is a solution.

Uniqueness. We next show that it is the only solution. From the Lemma, we know that $G(c_t, k_t)$ is decreasing in k_t. Given that $G(c_t, 1) = 0$, we can conclude that $k_t = 1$ is the unique solution. To see why note that when $k_t < 1$, then $G(c_t, k)$ would still be positive as it is decreasing in k_t and equal to zero when $k_t = 1$, violating the optimality condition $G(c_t, k_t) = 0$. Assuming $k_t > 1$ (over-insurance) is not feasible, $k_t = 1$ is the unique solution.

Proof of Proposition 1B. *Demand elasticities.* As long as insurance remains actuarially fair, the optimal insurance choice will be full insurance ($k_t = 1$). This implies that the demand for insurance is perfectly inelastic (price elasticity = 0) at the actuarially fair price.

Proof of Proposition 2A. When borrowing constraints are binding, we have $\lambda_t > 0$. We evaluate $G(c_t, k_t)$ at $k_t = 1$. When $k_t = 1$, w_{t+1} is given by Equation B.17.

$$G(c_t, 1) = V'(w_{t+1}) E_t \left[-(1+r)p + y_{t+1}M\right] - \lambda_t p.$$

Since we assume, actuarially fair insurance,

$$G(c_t, 1) = V'(w_{t+1}) \left[-(1+r)(\phi M/(1+r)) + \phi M\right] - \lambda_t p$$

$$G(c_t, 1) = V'(w_{t+1}) \cdot 0 - \lambda_t p$$

$$G(c_t, 1) = -\lambda_t p < 0$$

This holds because $\lambda_t > 0$ (binding constraint) and $p > 0$.

From the Lemma, we know that $G(c_t, k_t)$ is decreasing in k_t. Since $G(c_t, 1) < 0$ and $G(c_t, k_t)$ is decreasing in k_t, the optimal k_t that satisfies $G(c_t, k_t) = 0$ must be less than 1. Therefore, with a binding borrowing constraint and actuarially fair insurance, the optimal insurance coverage is less than full insurance ($k < 1$).

Proof of Proposition 2B. *Demand elasticities.* To derive the expression for demand elasticity, we differentiate $G(c_t, k_t)$ with respect to p and with respect to k and apply the implicit function

theorem, $\frac{\partial k}{\partial p} = -\frac{\partial G}{\partial p} \ / \ \frac{\partial G}{\partial k}$.

$$\frac{\partial G}{\partial p} = E_t\left[V''(w_{t+1})(-(1+r)p + y_{t+1}M)\frac{\partial w_{t+1}}{\partial p}\right] - E_t\left[V'(w_{t+1})(1+r)\right] - \lambda_t$$

$$\frac{\partial G}{\partial k} = E_t\left[V''(w_{t+1})(-(1+r)p + y_{t+1}M)\frac{\partial w_{t+1}}{\partial k}\right]$$

From the Lemma, we know that the denominator $\frac{\partial G}{\partial k}$ is negative due to the concavity of the value function. We need the numerator $\frac{\partial G}{\partial p}$ to be negative (ensuring a negative $\frac{\partial k}{\partial p}$).

The second and the third term are negative, as

$$E_t\left[V'(w_{t+1})(1+r)\right] > 0, \lambda_t > 0.$$

The first term has three parts. The first part $V''(w_{t+1}) < 0$, and the third part $\frac{\partial w_{t+1}}{\partial p} < 0$. The part in the middle will be negative when $y_{t+1} = 0$ and will be positive when $y_{t+1} = 1$. Therefore, when $y_{t+1} = 0$, we have the entire first term negative, implying $\frac{\partial G}{\partial p} < 0$. Instead, when $y_{t+1} = 1$, we may have the first term as positive. However, note that the expectation is likely to be dominated by the higher probability event of no loss $y_{t+1} = 0$, implying that the first term is likely negative and $\frac{\partial G}{\partial p} < 0$. Therefore, we can conclude that $\frac{\partial k}{\partial p}$ is negative, implying a negative demand elasticity for insurance coverage with respect to price, as long as the probability of loss is not very high.[38]

Proof of Proposition 2C. Let k_t^u and k_t^c be the optimal coverage levels for the unconstrained and constrained cases respectively. From the first-order conditions:

$$\text{Unconstrained} \quad E_t\left[V'(w_{t+1}^u)(-(1+r)p + y_{t+1}M)\right] = 0$$

$$\text{Constrained} \quad E_t\left[V'(w_{t+1}^c)(-(1+r)p + y_{t+1}M)\right] - \lambda_t p = 0$$

Since $V'(w)$ is decreasing in w (due to concavity of V), and $E_t\left[-(1+r)p + y_{t+1}M\right] < 0$ (assuming

[38] We require $\phi < \frac{V''(w_{t+1}^0)p(1+r)\frac{\partial w_{t+1}^0}{\partial p}}{V''(w_{t+1}^1)(M-p(1+r))\frac{\partial w_{t+1}^1}{\partial p} + V''(w_{t+1}^0)p(1+r)\frac{\partial w_{t+1}^0}{\partial p}}$, where w_{t+1}^y is the optimum wealth evaluated at $y = 0$. Note that ϕ is a value between 0 and 1, because it takes the form of $\frac{f_0}{f_1+f_0}$ where both f_0 and f_1 are positive due to (i) the concavity of V and (ii) $\partial w_{t+1}/\partial p < 0$, since wealth decreases in insurance prices.

$p > E_t[y_{t+1}]\, M/(1+r))$, we must have $w^u_{t+1} > w^c_{t+1}$ to satisfy these conditions. This implies $k^u_t > k^c_t$.

Proof of Proposition 3A. *Unconstrained borrowers.* Differentiating the budget constraint with respect to p_t we get

$$\frac{\partial S_t}{\partial p_t} = -\frac{\partial c_t}{\partial p_t} - k_t - p_t \frac{\partial k_t}{\partial p_t},$$

since, $\frac{\partial w_t}{\partial p_t} = 0$ as wealth w_t is given at t.

We next evaluate the signs of each of the three terms.

(i) $\frac{\partial k_t}{\partial p_t} < 0$, as shown above. This makes the last term positive.

(ii) $k_t > 0$ so the middle term is negative.

(iii) We show below that $\frac{\partial c_t}{\partial p_t} < 0$. So, the first term is positive.

Thus, the overall sign of $\frac{\partial S_t}{\partial p_t}$ depends on the relative magnitudes of these three terms. When an individual is relatively inelastic, such that coverage does not change by a large amount as insurance price changes, the third term will be small. This also implies that the first term will be relatively small, implying that the middle term will dominate. If so, $\frac{\partial S_t}{\partial p_t} < 0$. Thus, as insurance prices increase S_t would fall, i.e. borrowing would rise. In contrast, when an individual is highly elastic, changes in insurance prices would bring large changes in insurance demand, making the positive term dominate. If so, $\frac{\partial S_t}{\partial p_t} > 0$. Thus, as insurance prices increase S_t would rise, i.e. borrowing would fall.

Sign of $\frac{\partial c_t}{\partial p_t}$: We differentiate $G(c_t, k_t)$ with respect to p and with respect to c and apply the implicit function theorem, $\frac{\partial c}{\partial p}$ = - $\frac{\partial G}{\partial p}$ / $\frac{\partial G}{\partial c}$.

As shown above, $\frac{\partial G}{\partial p} = E_t\left[V''(w_{t+1})(-(1+r)p + y_{t+1}M)\frac{\partial w_{t+1}}{\partial p}\right] - E_t\left[V'(w_{t+1})(1+r)\right] - \lambda_t < 0$ (see Proposition 2).

$$\frac{\partial G}{\partial c} = E_t\left[V''(w_{t+1})(-(1+r)p + y_{t+1}M)\frac{\partial w_{t+1}}{\partial c}\right]$$

$$\frac{\partial G}{\partial c} = -(1+r)\cdot E_t\left[V''(w_{t+1})(-(1+r)p + y_{t+1}M)\right]$$

Note that the first part $V''(w_{t+1}) < 0$, and the third part $\frac{\partial w_{t+1}}{\partial c} < 0$. The part in the middle will

be negative when $y_{t+1} = 0$ and will be positive when $y_{t+1} = 1$. Therefore, when $y_{t+1} = 0$, we have the entire term as negative, implying $\frac{\partial G}{\partial c} < 0$. Instead, when $y_{t+1} = 1$, we have the term as positive. However, note that the expectation is likely to be dominated by the higher probability event of no loss $y_{t+1} = 0$, implying that as long as the probability of loss is not very high, $\frac{\partial G}{\partial c} < 0$.[39]

Proof of Proposition 3B. *Constrained borrowers.* In the constrained case, debt is pinned down by the borrowing constraint as $\lambda_t > 0$, implying $S_t = s$. An increase in insurance prices does not make the budget constraint slack ($\partial\lambda_t/\partial p_t > 0$), implying S remains at s. Thus, $\partial S_t/\partial p_t = 0$.

[39]We require $\phi < \frac{V''(w^0_{t+1})p(1+r)}{V''(w^1_{t+1})(M-p(1+r))+V''(w^0_{t+1})p(1+r)}$. Note that ϕ is a value between 0 and 1, because it takes the form of $\frac{f_0}{f_1+f_0}$ where both f_0 and f_1 are less than 0 due to the concavity of V.

C. Additional Details on the Instrumental Variables Specification

Our ideal instrument shifts the supply side of the market (e.g. insurer financing costs) without shifting the demand side of the market (e.g. household preferences or beliefs). While we build on a commonly used IV strategy in the IO literature, our setting and interpretation are different. To explain our instrument further, we provide additional details, including a structural equation interpretation and identifying assumptions. We write the structural equations and the state-time level, rather than the household level, to simplify notation and focus on the key identification assumptions.

Let P denote the insurance price in a low friction state s, and Q denote quantity (in our set up this is insurance coverage) in a low friction state s. Our cost-shifting instrument is FC_t, referring to insurer financing costs. W is a vector of other exogenous supply shifters, and X is a vector of other exogenous demand shifters. δ_y denotes a year fixed effect. Subscripts denote a household in low friction state s originating a mortgage in month-year t.

Suppose we seek to estimate the elasticity of demand in low friction states ε^D. The challenge with regressing quantities on prices is that our unobserved structural errors ξ_{st} and η_{st} are correlated because of the classic simultaneity problem, inducing a bias in the OLS estimation.

$$\textbf{Demand:} \quad \ln Q_{st}^D = \varepsilon^D \ln P_{st} + \gamma' X_{st} + \delta_y + \xi_{st} \tag{C.18}$$

$$\textbf{Supply:} \quad \ln Q_{st}^S = \varepsilon^S \ln P_{st} + \lambda FC_t + \theta' W_{st} + \eta_{st} \tag{C.19}$$

We consider insurer financing costs, FC_t as our supply shifter, influencing an insurer's supply curve but not the household's demand curve. Importantly for us, we assume that the insurer's financing cost is a function of the price it offers in high friction states, yielding $FC_t(\bar{P}_t^{HF})$. Our empirical strategy compares two observably equivalent borrowers (with the same characteristics X) that obtain a mortgage in different months of the same year. These two borrowers obtain different insurance prices due to price changes in high friction states, which in turn influence insurers' financing costs by influencing their future expected capital buffers and loss reserves. We explain the details of this intuition and the key identifying assumptions in more detail below.

Relevance. Why are insurer's prices in one state a function of their prices in other states? While this may make sense in a traditional product market context (Hausman (1996)), where there are common supply shocks to marginal costs (e.g. an input cost), it may be at first surprising in our insurance context, where one may expect insurers to optimize on a state-by-state basis and simply set prices based on local expected losses and markups. However, deviations from state-by-state optimization occur because insurers maximize profits in each state subject to a common across-state capital constraint. In other words, while expected losses are at a state level, financing costs are at the insurer level. This breaks the state by state separability of the insurer's optimization problem.

(1) Capital Constraints. First, the insurance literature (Froot and O'Connell, 1999, Koijen and Yogo, 2016) shows that insurers face capital constraints due to costly external financing constraints. Therefore, even in a perfectly competitive market, financing costs (or capital costs) lead insurers to add a "loading" (or surcharge) above expected losses in the price they charge. A related point is that an insurer's external financing costs are themselves a function of their equity. These financing costs typically spike after natural disasters, demonstrating insurer financial constraints, and incentivize insurers to maintain large reserves and equity buffers to smooth disaster shocks, limit their need to go to capital markets, and/or reduce the financing costs they obtain when trying to borrow or issue equity. As a result, the prices that an insurer offers *elsewhere* influence their overall capital buffer, affecting the prices that they charge in a given state.

(2) Cross-subsidies. This relationship with capital constraints becomes particularly important if insurers expect to operate at a loss in particular areas, what we call *high friction states* (Oh et al. (2021)). In the data, we see that combined ratios in high friction states are usually equal to 1–meaning insurers on average pay out in claims what they earn in premiums. In some particular years and states, these ratios can be much larger than one, meaning that insurers are operating at a loss. This result stems from the stringent pricing regulation in high friction states, that prevents insurers from charging fully actuarially fair rates. Importantly for us, operating at a loss in some states has the effect of eroding an insurer's capital, since claims must be paid out from retained earnings. Lower capital would, all else equal, increase an insurer's financing cost. Furthermore, in order to maintain their regulatory capital buffers and operate at a loss in some states, they need to charge more in other states. Regulatory capital constraints do bind for insurers - which suggests that

on the margin they would accept marginally higher financing costs and lower capital buffers. Taken together, operating at a loss in high friction states can lead to an increase in an insurer's marginal cost of operating *everywhere* by impacting their financing costs. This impact on marginal costs leads to a change in optimal pricing in low friction states. For example, the price that an insurer charges in the high friction state in California will influence their overall financing cost by impacting their equity capital, which in turn impacts the price that they can charge in Virginia, a low friction state.

Consistent with this idea, Oh et al. (2021) provide evidence that insurers tend to *cross-subsidize* risks across states by showing that loss events in high friction states cause asymmetric pricing adjustments: insurers adjust prices only in low friction states but not in high friction states. In contrast, when losses occur in low friction states, prices adjust locally and not in high friction states. Furthermore, consistent with the financing frictions channel, they show that these asymmetric spillovers are pronounced for insurers that are financially more constrained.

Building on this intuition, in our structural equations framework, we assume that average prices in high friction states $\bar{P}_t^{HF}$ shift financing costs FC_t, by influencing insurers' equity position. Functionally, this operates very similarly to the classic Hausman intuition of reflecting a common shock to marginal costs, such as a change in the cost of inputs (Hausman (1996), Nevo (2001)), except here the input is the cost of financing.

Because of these features, the instrument satisfies the relevancy condition required for an IV design: $cov(FC_t, P_{st}|X_{st}, W_{st}, \delta_y) \neq 0$.

Exclusion/Exogeneity. Crucially, a second important assumption for this approach is that prices in high friction states are excludable and exogenous: that financing costs only influence demand in low friction states through their influence on prices in low friction states. In our set up, this is reflected by the assumption that FC_t does not appear in the demand function (Equation 1), and that $cov(FC_t, \xi_{st}|X_{st}, \delta_y) = 0$: that financing costs are exogenous to the unobserved structural error in the demand equation. Because we include a year fixed effect, this requires, in particular, that average prices in high friction states do not vary *within the year* due to demand shocks in low friction states, such as a nationwide common demand shock.

We argue that this assumption is likely to hold in our setting because the variation in $\bar{P}_t^{HF}$ within the year is driven by the random timing and decisions of regulators in the high friction states.

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 adjust prices –either raise them or lower them– without obtaining regulatory approval from the relevant state insurance commissioner.

The question then becomes what drives the timing and decisions of regulators in high friction states. We argue that demand shocks *in other states* are unlikely to influence the timing of the regulatory process as well as their final rate decisions. That is, it is unlikely to be correlated with other unobserved variables in the structural error term ξ that would influence demand.

For example, a standard concern in the IO literature 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. Crucially, in our setting, even nationwide advertising campaigns are unlikely to be a concern, because of the fact that insurance companies in high friction states cannot immediately adjust their prices in response to the advertising campaign. They must still wait for the state insurance commissioners to deliberate on their rate request and issue a final decision, even if household demand changes in response to the advertisements.

Another concern in the climate context is that, as disaster-related losses mount over time, perhaps regulators could get more lenient over time or more concerned about climate change; this could also coincide with households similarly becoming more concerned with climate change. Such aggregate trends would pose a challenge, and as a result we include a year fixed effect, to ensure that our identification limits to within-year variation.

A third 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, a large fire event in California could lead to more concerns about climate risk in Virginia. Again, what is crucial for our setting is that prices in high friction states do not immediately adjust with the fire event. Oh et al. (2021) show that prices in high friction states do not change in the immediate aftermath of natural disasters in high friction states. This is because of the fact that, once again, insurance companies require regulatory approval to adjust prices, and regulatory approval in high friction states takes time to obtain. Thus,

[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.

the *within-year* variation in $\bar{P}_t^{HF}$ will also not be driven by climate events in high friction states.

Monotonicity: A third question is how the financing cost function FC is defined. For monotonicity, it must be the case that λ is the coefficient stable and that $\frac{\partial FC_t}{\partial \bar{P}_t^{HF}}$ also has a stable sign. One may be concerned about whether this holds; the standard Hausman intuitions of common supply shocks would lead to a positive correlation between prices in low and high friction states. However, an opposing channel based on a cursory cross-subsidy intuition would be that price increases in high friction states would lead to a *lower financing cost* by improving insurers' equity buffers, which would in turn lead to lower prices. This would mean that the IV may not be monotonic, sometimes having a positive first stage and sometimes having a negative first stage. Importantly, while intuitive, this latter logic is in fact missing an important consideration about the cross-subsidy: insurers' expected losses are also changing over time. If expected losses were held fixed, then of course a higher price in high friction states would improve insurers' equity position and lower financing costs by lowering the losses they face. However, if insurers' expected losses are also changing over time, which is the case in the insurance setting in recent years, then higher prices in high friction states does not automatically indicate a higher equity buffer.

What is important for the financing cost function is not $\bar{P}^{HF}$ by itself, but what is the expected effect of the high friction state's price change on *the price gap*, that is the deviation between the price they can charge and their expected losses $\bar{L}$: $PGAP_t = (\bar{P}_t^{HF} - \bar{L}_t^{HF})$. A higher price gap would lead to an increase in FC_t; a lower price gap would lead to a lower FC_t. So for monotonicity, the question becomes what is the covariance between high friction state prices and the price gap $cov(PGAP_t, \bar{P}_t^{HF})$? Stated another way, how do prices in high friction states actually influence the price gap? If this covariance is negative, then we have a monotonicity issue; if the covariance is positive, then both the traditional Hausman and the cross-subsidy channels through which the IV is relevant push in the same direction.

We show empirically that this covariance is positive: when regulators in high friction states allow larger price increases, they are associated with even larger price deviations from the amount the insurers required. When insurers receive 10%, their model implied actuarially fair rates are higher (e.g., they require a 40% rate increase). Instead, when insurers receive 2%, their model implied actuarially fair rates are lower (e.g., they require a 10% rate increase). As a result, price increase in

high friction states are actually associated with a larger price gap, indicating an even larger capital erosion and thus an even *higher* financing cost.

Thus, both the traditional Hausman intuition as well as this price gap intuition would suggest that prices in low friction and high friction states are positively correlated. Indeed, Table 2 shows that we find a strong positive correlation in our first stage that $cov(P_{st}, \bar{P}_t^{HF}) > 0$. Therefore, monotonicity is less likely to be a concern.

Timing of adjustment. A key question is about the timing of insurer's adjustments in low friction states. One possibility is that insurers adjust prices immediately in low friction states when their estimate of expected losses in a high friction state changes, without waiting for regulatory clarity. In this scenario, insurers would form an expectation of the rate change they expect regulators to grant, as well as the timing of when they expect to receive the rate change. They would then adjust prices in low friction states accordingly based on these expectations. Thus, price adjustment in low friction states would be a function of the *expected* price gap: $(\mathbb{E}[P^{HF}] - \bar{L}^{HF})$. Insurers would then adjust again when regulators actually make a decision; the second adjustment would depend on how accurate their expectations are, captured by $\bar{P}^{HF}_{actual} - (\mathbb{E}[P^{HF}] - \bar{L}^{HF})$. On the other hand, insurers could also wait for regulators in high friction states to make a decision, and then adjust prices in low friction states when the regulatory uncertainty is resolved.

It is very standard in models of financing frictions for agents to anticipate the friction ex-ante and build a precautionary buffer (Bolton et al., 2011). As as result, insurers may build an upfront precautionary pricing buffer, particularly in low friction states. One may question why insurers choose ex-post adjustments (wait for regulatory approval or losses) when they have the ability to incorporate a pricing buffer upfront. Importantly, these models also imply that insurers should adjust pricing ex-post, even if they have taken actions to mitigate the frictions ex-ante. This is because building a buffer large enough to perpetually avoid future financial constraints can be prohibitively expensive. Indeed, almost all of these models include some "frictions" that prevent the ex-ante buffers from becoming so large that the financial frictions never actually bind in the future (Bolton et al., 2011). For example, one common reason may be the costs of having an internal liquidity buffer. In our context, regulation itself may prevent insurers from building a large buffer upfront. While low friction states are less stringent, regulators still require rates to be justified by past losses and

insurers' financial constraints (which requires clarity about high friction state prices). This may limit insurers' ability to preemptively adjust initial rates by a large amount. Particularly, the uncertainty about the regulators' decision and the extent to which financial frictions may bind in the future, may prevent insurers from fully updating rates ex-ante.

D. Additional Tables and Figures

Figure D.1: Data Construction Procedure

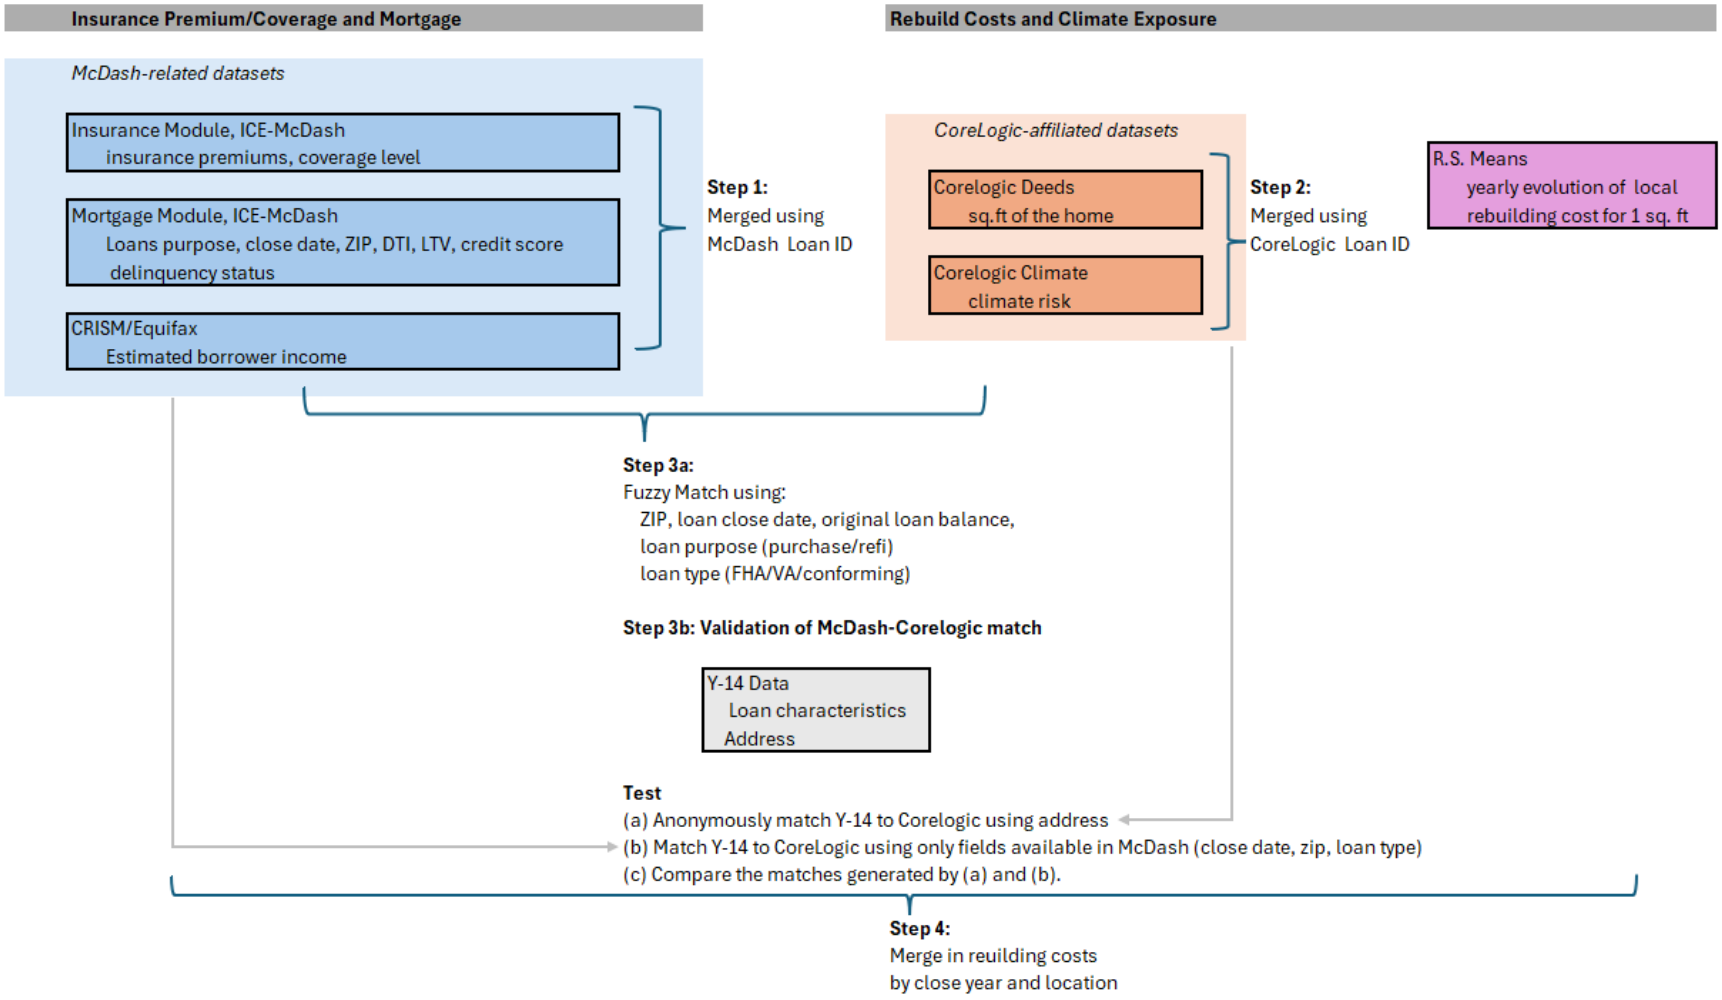
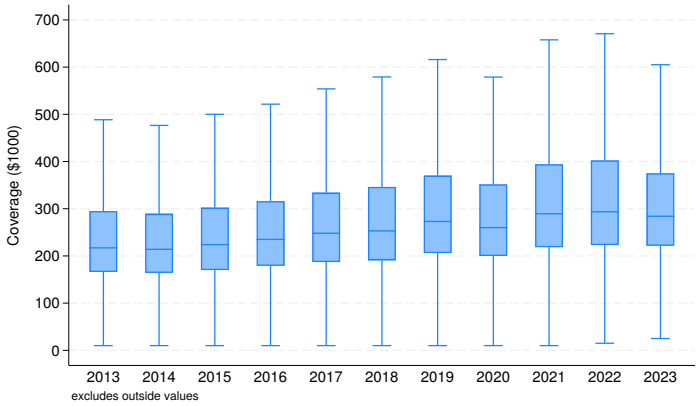


Figure D.2: Evolution of Insurance Coverage and Premiums Over Time

This figure shows a box plot time series of insurance coverage levels (nominal values reported in thousands) in panel a and annual insurance premiums in panel b. The shaded area indicates the inter-quartile range, with the median marked with a horizontal line. The whiskers indicate the upper and lower adjacent values. The data are from ICE-McDash. The figure is based on the main analysis sample, outlined in Section 1.2.

(a) Coverage Levels



(b) Insurance Premiums

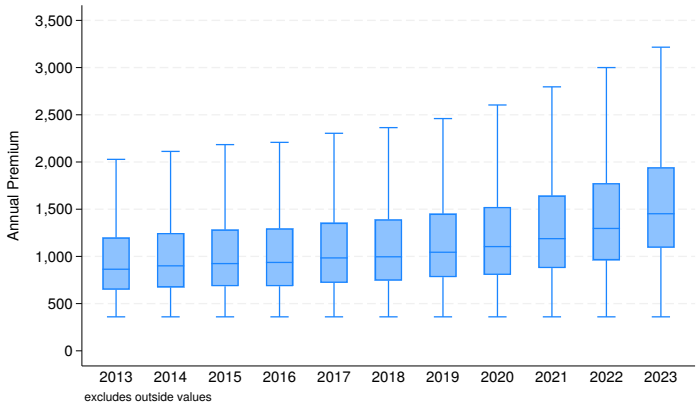


Figure D.3: Histogram of Coverage Ratio

This figure shows a histogram of coverage ratios, defined as insurance coverage amount divided by the estimated replacement cost, obtained by multiplying structure size and local construction costs from the R.S. Means Company. The insurance coverage data are from ICE-McDash. The figure is based on the main analysis sample, outlined in Section 1.2.

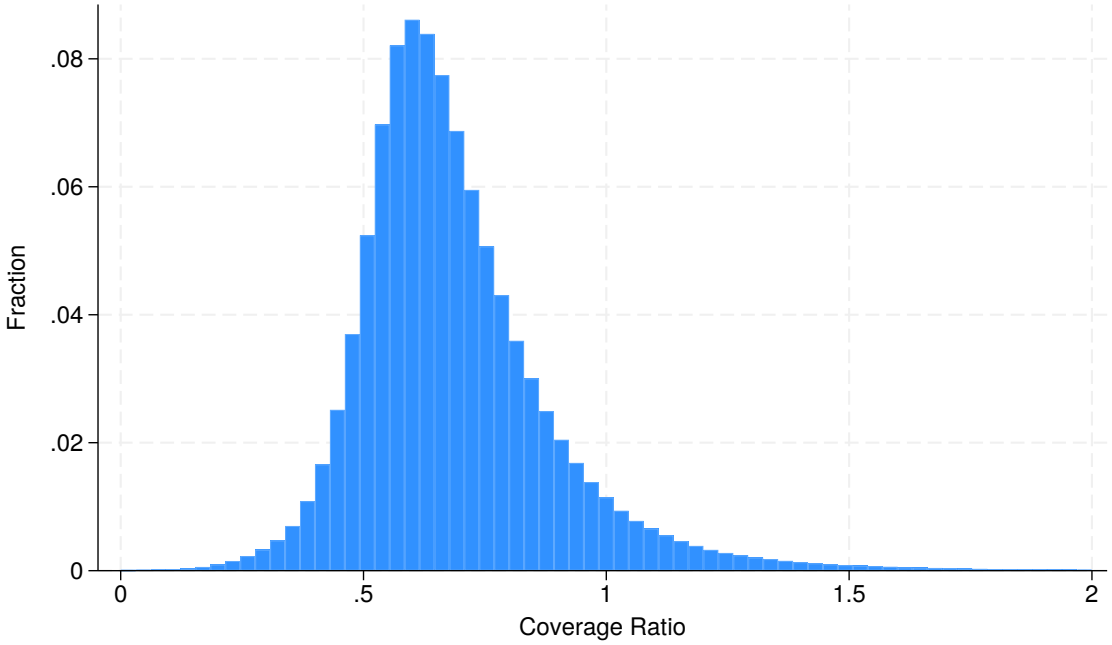
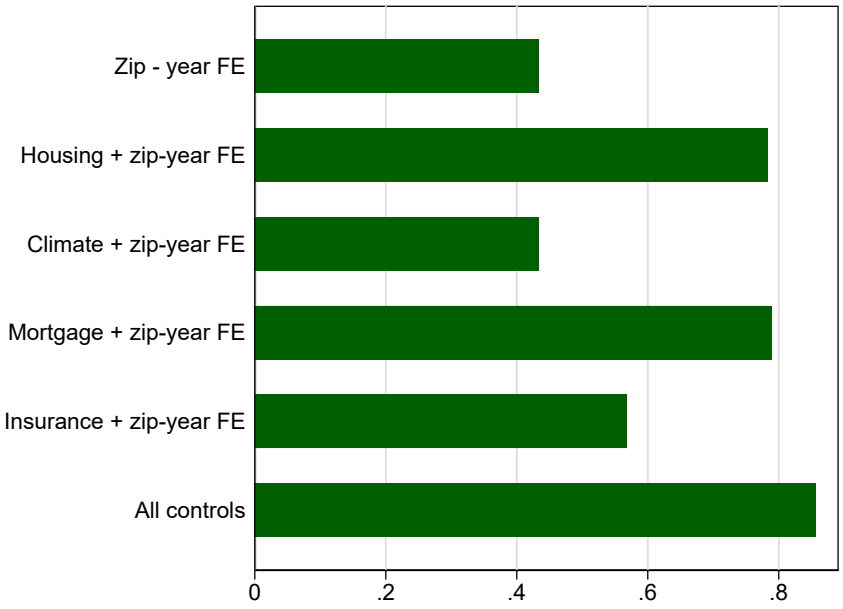


Figure D.4: Explaining the Heterogeneity in Insurance Coverage and Prices

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, 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 1.2.

(a) Insurance Coverage



(b) Insurance Price

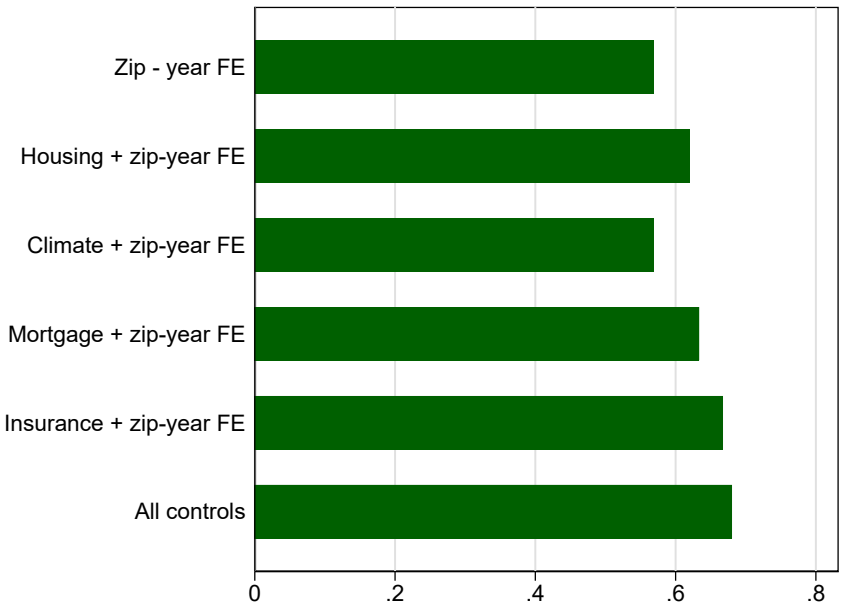


Figure D.5: Coverage Ratios by Loan Age

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 1.2, which we observe through time to assess changes in coverage ratio.

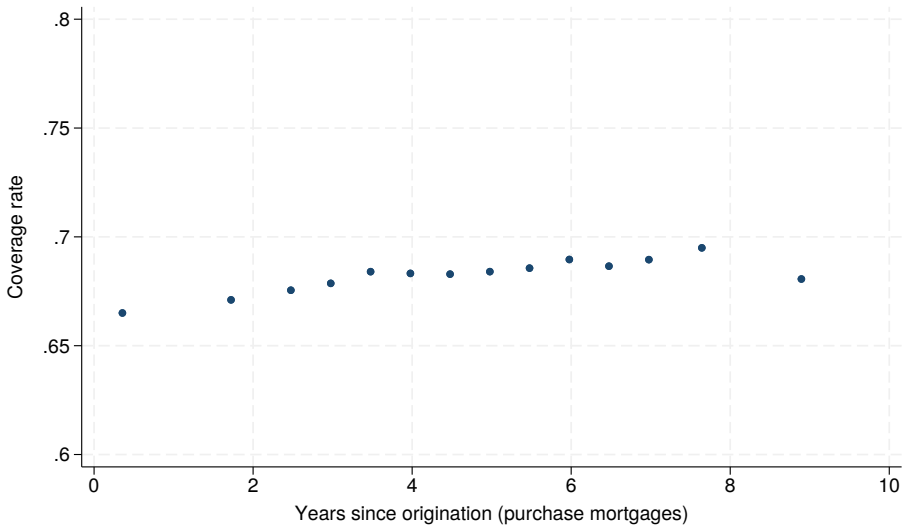
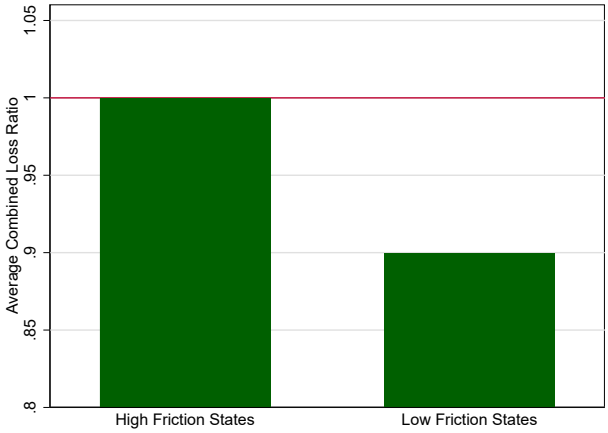


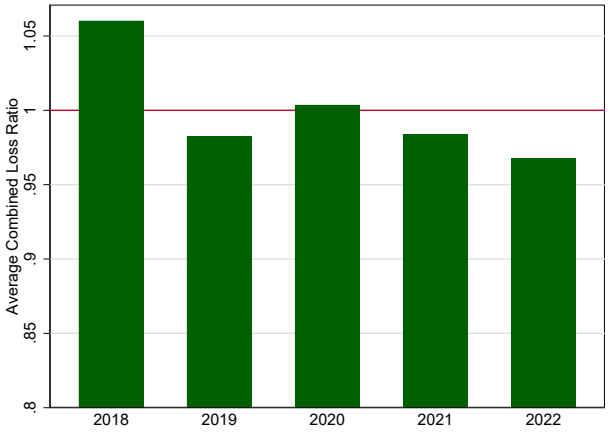
Figure D.6: Insurers' Underwriting Profitability Across State Types

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



D.1. Appendix Tables

Table D.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 D.2: 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
Number of Observations	3238688	3238688
Adjusted R-squared	0.605	0.617

Table D.3: 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 D.4: OLS Results: Effect of Insurance Prices on Coverage (Low Friction States)

This table presents the results of the following OLS regression:

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

which is the simple OLS equivalent of the 2SLS in Equation 14. The dependent variable is log coverage amount. The independent variable is log insurance price (monthly premiums per $100,000 of coverage). The analysis uses the sample of low friction states. High and low friction states are as defined in Section 1.2. 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 2. Standard errors are clustered at the county level. The sample is as outlined in Section 1.2. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln coverage$_{izct}$	
	(1)	(2)
ln P$_{izct}$	-0.235***	-0.228***
	(0.00745)	(0.00744)
ZIP-Year FE	Y	Y
Property Controls	Y	Y
Borrower Controls	N	Y
Number of Observations	1791932	1791932

Table D.5: 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 log 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 1.2. 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 2. Standard errors are clustered at the county level. The sample is as outlined in Section 1.2. 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 D.6: OLS Results: Effect of Insurance Prices on Mortgage Debt (Low Friction States)

This table presents the results of the following OLS regression:

$$\ln Debt_{i,c,z,t} = \alpha + \beta \ln P_{i,c,z,t} + \gamma' W_{i,c,z,t} + \epsilon_{i,c,z,t}$$

which is the simple OLS equivalent of the 2SLS in Equation 15. The dependent variables in columns 1 and 2 are the log of the loan's DTI ratio and in columns 3 and 4 the log of the loan's LTV ratio, both taken as of origination. The independent variable is log insurance price (monthly premiums per $100,000 of coverage). The analysis uses the sample of low friction states. High and low friction states are as defined in Section 1.2. 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 2. Standard errors are clustered at the county level. The sample is as outlined in Section 1.2. Note: *p<0.1; **p<0.05; ***p<0.01.

	ln DTI$_{izct}$		ln LTV$_{izct}$	
	(1)	(2)	(3)	(4)
ln P$_{izct}$	0.0445***	0.00828***	0.0288***	0.00803***
	(0.00318)	(0.00182)	(0.00168)	(0.00103)
ZIP-Year FE	Y	Y	Y	Y
Property Controls	Y	Y	Y	Y
Borrower Controls	N	Y	N	Y
Number of Observations	1791932	1791932	1791932	1791932

Table D.7: Second-Stage Results: Coverage Demand Elasticities using Alternate Proxy for Debt Capacity

This table shows the second stage estimates of the two-stage least squares regression shown in Equation 14. We present the results for two sub-samples: those with debt-to-income ratios that are below 30% (Column 1) and above 30% (Column 2-3), which represents a key underwriting threshold in the mortgage market. We use this as an alternative proxy of financing constraints. The analysis uses the sample of low friction states. In columns 1 and 2, the dependent variable is log coverage amount. In column 3, the dependent variable is the log of discretionary coverage, which is coverage in excess of the loan balance. 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 1.2. 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 2. Standard errors are clustered at the county level. Note: *p<0.1; **p<0.05; ***p<0.01.

	(1)	(2)	(3)
	ln coverage$_{iczmt}$	ln coverage$_{iczmt}$	ln discretionary coverage$_{iczmt}$
	Low DTI	High DTI	High DTI
$\widehat{\ln P_{iczmt}}$	-0.00721	-0.114***	-1.150***
	(0.0363)	(0.0170)	(0.114)
ZIP-FICO-Year FE	Y	Y	Y
Income-Year FE	Y	Y	Y
Insurer-Year FE	Y	Y	Y
Property Controls	Y	Y	Y
Borrower Controls	Y	Y	Y
Number of Observations	357913	1144797	790982
Adjusted R-squared	0.694	0.696	0.210
First-Stage F	197.1	271.4	356.3

Table D.8: Effect of Insurance Prices on Lender Denials of High LTV Mortgage Applications by Income Group

This table shows the effect of insurance prices on mortgage lender denial rates. This sample comes from the public Home Mortgage Disclosure Act sample, which is annually available administrative data on the near-universe of residential mortgage applications. The dependent variable of the regression is a dummy variable indicating that a mortgage application was denied. The independent variable is the annual average of our instrument, the log of average prices in high friction states. We limit the sample as indicated to high loan-to-value mortgages, defined as mortgages with LTVs that exceed 90% (Columns 1-2) or LTVs that exceeds 96% (Columns 3 and 4). Borrower controls refer to the borrower's age, property value, and debt-to-income ratios. We also include census tract fixed effects. We show the results for subsamples split the sample by income as indicated. Standard errors are shown in parentheses. Note: *p<0.1; **p<0.05; ***p<0.01.

Sample:	Application Denied (Y/N)			
	(Mortgage LTV>90%)		(Mortgage LTV>96%)	
	High Income	Low Income	High Income	Low Income
	(1)	(2)	(3)	(4)
ln $\bar{P}_t^{HF}$	0.000918	0.169***	-0.0309	0.201***
	(0.0315)	(0.0217)	(0.0520)	(0.0270)
Borrower Controls	Y	Y	Y	Y
Tract FE	Y	Y	Y	Y
Number of Observations	72621	205389	34498	146383
Adjusted R-squared	0.0479	0.0722	0.0574	0.0733

Table D.9: Robustness: 2SLS for Mortgage Debt with Interest Rate Controls

This table presents the second stage of the two-stage least squares regression shown in Equation 15, where we show that the results are robust to controlling for the interest rate of the mortgage. The analysis uses the sample of low friction states. The dependent variables in columns 1 and 2 are the log of the loan's LTV ratio, both taken as of origination. 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 1.2. 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 2. In addition, we control for the interest rate of the mortgage. Standard errors are clustered at the county level. The sample is as outlined in Section 1.2, but limited to the sample of loans in low friction states. Note: *p<0.1; **p<0.05; ***p<0.01.

	$\ln \text{LTV}_{izct}$	
	(1)	(2)
$\widehat{\ln P_{izct}}$	0.614***	0.695***
	(0.0232)	(0.0407)
ZIP-Year FE	Y	Y
Property Controls	Y	Y
Borrower Controls	N	Y
Interest Rate Control	Y	Y
Number of Observations	1791932	1791932
First-stage F	744.4	462.3