The Effect of Rising Insurance Premiums on Mortgage Delinquency and Household Relocation*

Shan Ge† Stephanie Johnson‡ Nitzan Tzur-Ilan§

Abstract

U.S. homeowners' insurance premiums have increased drastically, widely attributed to climate risk. Using household-level microdata and two novel instruments, we document three findings. First, households facing larger premium increases are more likely to relocate, moving to areas with lower climate risk and insurance costs. Homes with larger increases are more often acquired by investors, who obtain lower premiums. Second, higher premiums increase mortgage delinquencies. Third, the relocation effect is stronger among less financially constrained households, while the delinquency effect concentrates among more constrained ones. Rising insurance costs can threaten household financial resilience, while relocation offers a costly mechanism of adaptation.

JEL Classification: G21, G22, G5, G52, G53, R21, Q54, D14, R3
Keywords: climate risk, insurance, mortgage, delinquency, relocation

*We thank Viral Acharya, Natee Amornsiripanitch, Siddhartha Biswas, Noémie Bucourt, Tony Cookson, Julia Fonseca, Emily Gallagher, Stefano Giglio, Matthew Gustafson, Mallick Hossain, Ben Keys, Ralph Koijen, Theresa Kuchler, Lu Liu, Tatyana Marchuk, Holger Mueller, Philip Mulder, Cecilia Parlatore, Pari Sastry, Alexi Savov, Johannes Stroebel, Katherine Wagner, Jeffrey Wurgler, Moto Yogo, and seminar and conference participants at Brigham Young University, Wharton Finance, MIT, NYU, UPenn Energy Economics & Finance Seminar, Ohio State University, UIUC, University of Colorado Boulder, University of Notre Dame, University of Florida, Indiana University, University of Wisconsin-Madison, Bocconi, HEC Paris, INSEAD, Hebrew University, Freddie Mac, NBER Insurance Meeting, NYU Summer Climate Finance Conference, AFA, WFA, RADAR Climate Seminar, Young Scholars Finance Consortium at Texas A&M, USC, SMU, CEPR Climate Change and the Environment Symposium, and Sustainable Financial Intermediation Conference, for helpful comments and suggestions. We affirm that we have no material financial interests related to this research. The information presented herein (including any applicable table, chart, graph, or the like) is based on data provided by CoreLogic, ICE McDash, and the Credit Risks Insight Servicing McDash (CRISM). Those data are used as a source, but all calculations, findings, and assertions are those of the authors. We thank Dylan Ryfe and Dongyu Zuo for excellent research assistance. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

†Stern School of Business, New York University, E-mail: sg3634@stern.nyu.edu.
‡Jones Graduate School of Business, Rice University, E-mail: stephanie.g.johnson@rice.edu.
§Federal Reserve Bank of Dallas, E-mail: nitzan.tzurilan@dal.frb.org.

1 Introduction

Rising insurance premiums have become a major financial concern for U.S. households. They can push households toward costly adaptation or trigger negative financial outcomes. Understanding these effects matters for assessing the broader impact of climate change, which is widely regarded as a main driver of recent and projected increases.

Homeowners can mitigate this financial pressure through costly strategies. They may search more widely across insurers, or relocate to areas or properties with lower insurance costs. According to a Redfin survey, 10% of homeowners planning to move cite lowering insurance costs as a reason.[1] For marginal households, the present value of insurance savings can shift the cost-benefit calculation of moving. Some buyers may have advantages in holding the vacated home, and can therefore pay above the departing owner's valuation, further tipping the decision to move. Understanding this relocation response sheds light on a key adaptation mechanism to climate change.

For households unable to mitigate through these strategies, rising premiums threaten to erode financial resilience, a concern that has drawn attention from policymakers and regulators.[2] We focus on the impact on mortgages, because they dominate household balance sheets (Campbell and Ramadorai, 2025; Gomes et al., 2021), accounting for 73% of total household debt.[3] Since mortgage payments typically take priority in household budgets, mortgage performance is informative about household financial health. Rising insurance premiums reduce the liquidity available for mortgage payments, potentially increasing delinquency rates and spreading losses to the financial sector, given the systemic importance of mortgages.

We ask how rising homeowners' insurance costs affect household financial outcomes

[1] See `https://www.redfin.com/news/moving-reasons-safety-survey-2024/`.

[2] See, e.g., U.S. Department of the Treasury, Federal Insurance Office's Analyses of U.S. Homeowners Insurance Market, `https://home.treasury.gov/news/press-releases/jy2791`; Financial Stability Oversight Council, 2025 Annual Report, `https://home.treasury.gov/system/files/261/FSOC2025AnnualReport.pdf`.

[3] See `https://www.newyorkfed.org/microeconomics/hhdc`. The number is as of 12/2/2025.

and how households adapt by relocating. These questions remain largely unexplored. We use a newly available dataset that links detailed insurance policy information to household-level mortgage and credit report outcomes in the U.S. between 2013 and 2023.

We document two main findings. First, larger premium increases make households more likely to become delinquent on their mortgages. Second, households facing larger premium increases are more likely to relocate to areas with lower climate risk and lower premiums. The homes they vacate are also more likely to be acquired by investors, who obtain lower premium rates.

To identify the causal effect of premium increases, we construct two novel instruments that exploit idiosyncratic differences in the timing and size of insurers' rate changes. The first instrument is, for each household-policy-year, the average premium rate increase among policies renewing in the same month with the same insurer in other ZIP codes within the same state. The second instrument is constructed identically, except that it restricts to distant ZIP codes, those farther from the focal ZIP code than the median pairwise ZIP code distance in the state.

The relevance of our instruments stems from two mechanisms. First, premium increases follow location-level trends driven by insurers' risk assessments of that location (Keys and Mulder, 2024), and the risks being assessed are correlated within a state: ZIP codes in the same state are exposed to common weather and catastrophe events, so an insurer that revises its assessment reprices across its book within the state. Second, premium increases are subject to state regulator approvals of insurers' rate change requests (Oh et al., 2022). Therefore, a property's premium change should correlate strongly with the same insurer's premium changes elsewhere in the state in the same month.

Our instrument is also plausibly exogenous. Since our regressions include ZIP-by-renewal month fixed effects, we effectively compare borrowers within the same ZIP code who renew their insurance policies in the same month. The within-ZIP-by-renewal month variation in the instruments stems from differences across insurers based on their risk

assessment and the timing of their rate changes. Since households usually receive renewal notices with premium information in advance, they could switch insurers before renewal. To address this endogenous response, we use the insurance policy expiration month and the insurer determined in the previous policy year, so the instrument is predetermined relative to households' subsequent choices. Our instrument therefore likely satisfies the exclusion restriction: it is unlikely to affect individual borrowers' mortgage performance or relocation outcomes through channels other than insurance costs.

To absorb the direct effect of local risk, we control for climate risk bin-by-renewal month fixed effects and climate risk, in addition to ZIP-by-time fixed effects. The instrument, constructed from the same insurer's rate changes in other or distant ZIP codes, should therefore move household outcomes through the premium rather than through climate risk itself.

At the same time, our empirical strategy can still speak to the effect of climate risk on households through the insurance market. Our instrument may partially reflect insurers' repricing of climate-related risk, with its variation arising from differences in when they reprice and how they assess the risk in each state. Indeed, at the ZIP-month level, the instrument varies more in high-risk areas in recent years. Insofar as households cannot distinguish premium increases driven by our instrument from those driven by climate risk, the effect of increases attributable solely to climate risk may be similar.

We begin with delinquency. Rising premiums raise the probability of mortgage delinquency within 12 months of the premium change at renewal: when annual premiums increase by \$1,000, the probability of 60-day delinquency rises by 0.38 percentage points, a 22% increase relative to the mean of 1.7%.[4] We estimate that the average premium increase at 2023 renewals caused roughly 42,000 mortgages to be 60 days past due.[5]

[4]Throughout the paper we convert premium rates into dollars at a coverage limit of \$500,000. We use this figure because it is the round number closest to the average coverage limit at 2023 renewals, \$497,000, and because it makes the conversion transparent: a \$1,000 increase in the annual premium corresponds to a 0.2 percentage point increase in the premium rate.

[5]We multiply the Column (5) estimate of 0.019 in Table 2 by the average premium rate increase at 2023 renewals, 0.0438 percentage points of the coverage limit, which implies a $0.019 \times 0.0438 \approx 0.083$ percentage

In addition, the delinquency effects are stronger for more constrained borrowers: those with lower credit scores, and those with higher loan-to-value or debt-to-income ratios at origination. The distress extends beyond missed payments: larger premium increases also raise foreclosure rates and credit card utilization. That the effects appear across several margins suggests the delinquency results reflect genuine financial stress rather than a mechanical relationship between insurance and mortgage payments.

In the second part of the paper, we turn to how households mitigate the effects of rising premiums. One margin is to search for a cheaper insurer: in Ge et al. (2026c), we show that households facing larger premium increases are more likely to switch insurers, and switching lowers premiums. The other margin, and our focus here, is relocation.

We find that households are more likely to relocate in response to rising insurance premiums. Movers who experienced larger pre-move premium increases relocate to areas with lower climate risk and lower insurance premiums. The relocation response is stronger among households with higher origination mortgage rates relative to contemporaneous market averages, who are therefore less affected by mortgage lock-in. Mortgage rates have risen in recent years alongside large increases in insurance premiums, so lock-in may have limited the relocation response over our sample period. We expect that response to strengthen as mortgage rates fall.

A $1,000 increase in annual premiums corresponds to a 0.82 percentage point increase in relocation probability (17% of the mean). This is comparable to the mortgage lock-in effect on household relocation documented by Fonseca and Liu (2024) and Fonseca et al. (2026). The response is stronger for borrowers with higher credit scores and lower LTV or DTI ratios, who can better afford the liquidity cost of moving. For the same premium increase, non-movers pay a further $280 per year, while movers pay $1,184 per year less than those non-movers. Assuming a discount rate of 6% (Giglio et al., 2021), the present

point increase in the 60-day delinquency rate, and apply that increase to the 51 million first-lien mortgages outstanding in the U.S. (Federal Housing Finance Agency, National Mortgage Database).

value of the movers' premium savings amounts to $13,580 over 20 years.[6]

Why relocate at all, when switching is less costly? Part of the answer is that switching offsets only part of an increase: a household can leave an unfavorable carrier, but insurers serving its location price broadly the same underlying risk, so the part of an increase common to the area can be escaped only by moving. For the same $1,000 increase, switching lowers the following year's premium by about $400 (Ge et al., 2026c), substantially less than the $1,184 that movers save. Even that $400 reduction is temporary. Within a ZIP code, insurers raise premiums at different times, and those that raise later eventually exceed those that raised earlier. The insurer a household switches to is inexpensive partly because it has not yet raised its premiums, so the saving erodes over time. For some households there may be no cheaper quote: a carrier that previously charged well below its competitors may remain the cheapest in terms of premium levels even after a large increase.

Because relocation is costly, even a $13,580 saving is unlikely to trigger a move on its own. Consistent with this, the switching response is far more sensitive: the proportional increase in the probability of switching insurers is roughly 18 times the proportional relocation response. The saving instead adds to a household's accumulated reasons to relocate (a job change, a growing family) and tips those already near the margin. We formalize this move-versus-switch choice in a model, presented in Internet Appendix IA.B.

As movers vacate their previous homes, those that experienced larger premium increases are more likely to be purchased by investors. Investors pay 14% lower total premiums, a result of lower premium rates, lower coverage, and higher deductibles. Their lower rates may also come from greater market power in negotiating with insurers, or from lower costs of searching for insurance. Their lower coverage and higher deductibles could reflect lower effective risk aversion. Investors with greater financial flexibility may

[6]Every household in our sample has an active mortgage, and among these households the probability of moving, as defined in Section 8, is 0.05 (Table 1). At a constant annual rate, that implies an expected tenure of about 20 years.

be better positioned to self-insure against losses. They may also hold operational advantages, such as established contractor relationships or economies of scale in repairs, lowering the coverage they need.

Our mortgage delinquency results carry implications for household financial resilience and broader financial stability. While consistent with the literature on cash-flow shocks triggering mortgage defaults,[7] our findings reveal a distinct channel through which climate change can affect households and the financial system.

The literature on mortgage payments and delinquency does not, however, settle the question we ask, because a premium increase is a shock households can act on: they can search for a cheaper insurer, raise their deductible, cut coverage, or move. Cutting coverage lowers the bill only by transferring disaster risk to the household. The delinquency estimate therefore captures the cost remaining after mitigation, which could in principle have been close to zero.

As disasters become more severe and frequent, they push insurance rates higher, placing financial strain on those who do not directly experience them. This mechanism poses an emerging risk to the financial sector and the federal government, as we find delinquency effects in government-backed and agency mortgages, and those held in bank portfolios. Once premiums reach their projected 2055 levels, 136,000 more mortgages would become delinquent each year.[8] These results speak to the policy debate over insurance affordability as policymakers consider interventions such as means-tested subsidies.[9]

Our relocation findings have implications for long-run demographic sorting and the distribution of climate risk exposure. Financially better-off households are more likely

[7]See, for example, Gupta (2019); Ganong and Noel (2023); Indarte (2023); Kleiner et al. (2021), Ganong and Noel (2020), and Keys (2018).

[8]First Street Foundation projects premiums will rise 29.4% by 2055 on average in the U.S. (`https://www.realtor.com/news/trends/biggest-home-insurance-premium-hikes/`). Applied to the 0.48% average premium rate in our sample (Table 1), the Column (5) estimate of 0.019 in Table 2, and the 51 million first-lien mortgages outstanding in the U.S.: $0.48 \times 0.294 \times 0.019 \approx 0.27$ percentage points, or roughly 136,000 additional delinquencies per year. Mortgage count from the Federal Housing Finance Agency's National Mortgage Database, `https://www.fhfa.gov/data/dashboard/nmdb-aggregate-statistics`.

[9]For related policy proposals, see, `https://www.cbo.gov/publication/59918`.

to relocate to lower-risk areas as premiums increase, while constrained households remain in place despite higher costs. If premiums continue to rise in high-risk areas over the coming decades, this pattern may drive substantial demographic shifts, with less constrained households migrating away and more constrained households becoming increasingly concentrated in those areas. This sorting could exacerbate existing inequalities over time, as financially constrained households face both higher disaster exposure and a reduced capacity to absorb climate-related shocks.

The increased presence of investors in properties experiencing large premium increases also has long-term implications for housing markets in risky areas. Investors' lower insurance costs give them a competitive advantage in high-risk markets. As climate risks intensify and premiums continue to rise, this advantage may accelerate a shift away from individual homeownership toward investor ownership in vulnerable areas. Given evidence that homeownership facilitates wealth accumulation and generates positive community externalities (Sodini et al., 2023; Artigue et al., 2025; DiPasquale and Glaeser, 1999), such a transition could affect household wealth inequality and neighborhood amenities in climate-exposed regions over the coming decades.

We are among the first to identify the effect of property insurance premiums on households' financial distress and relocation decisions. Our work contributes to several related bodies of literature. First, we extend research on insurance and mortgage markets. Sastry et al. (2026) find that financially less constrained households borrow more through purchase mortgages in response to increases in homeowners' insurance premiums. Sastry (2026) shows that lenders require higher down payments from borrowers who under-insure, potentially due to post-disaster default concerns. Sastry et al. (2023) find higher mortgage defaults in areas with greater insurer insolvency following disasters. We demonstrate that premium increases alone, independent of actual disasters, can also elevate default risk. This mechanism may explain the finding in Ge et al. (2024) that buyers of homes facing exogenous flood insurance premium increases are less likely to

obtain mortgages, as lenders may anticipate higher default risks from elevated insurance costs. By highlighting the important effect of insurance premiums, we also contribute to a growing literature on insurance pricing and market dynamics.[10] Beyond insurance pricing, Kalda et al. (2025) and Jotikasthira et al. (2025) uncover that insurers deny and delay claim payments as another response to losses.

In addition, our study contributes to the growing body of literature examining the intersection of climate change, natural disasters, and mortgage markets. Prior papers explore the impact of disasters on mortgage delinquencies (Gallagher and Hartley, 2017; Kousky et al., 2020; Billings et al., 2022; Issler et al., 2024; Biswas et al., 2023) and the pricing of mortgage-related securities (Gete et al., 2024; Dice et al., 2024). We highlight a different channel (insurance costs) through which disaster and climate risks can affect mortgage outcomes. In so doing, we also complement recent work on climate/environmental risk in housing markets (Bernstein et al., 2019; Baldauf et al., 2020; Murfin and Spiegel, 2020; Keys and Mulder, 2020; Giglio et al., 2021; Lopez and Tzur-Ilan, 2023; Ge et al., 2024; An et al., 2024; Van Binsbergen et al., 2025).

Our relocation results also contribute to a literature on climate risk, disasters, and where households live. Boutros et al. (2026) model households' joint choice of location, tenure, and property insurance coverage. More broadly, quantitative spatial models treat moving as a central margin of climate adaptation (Bilal and Rossi-Hansberg, 2025), and related work embeds relocation in models of global warming and coastal flooding (Desmet and Rossi-Hansberg, 2015; Desmet et al., 2021; Cruz and Rossi-Hansberg, 2024).

[10] Cookson et al. (2024) uncover coverage neglect in the homeowners insurance market. Liao and Mulder (2021) study the effect of home equity on flood insurance demand. A few papers study the effect of insurance on the housing market, including Ge et al. (2024), Georgic and Klaiber (2022), Gibson and Mullins (2020), Bakkensen and Barrage (2021), Ge et al. (2026a), and Ge et al. (2026b). Papers on insurance pricing include Keys and Mulder (2024), Boomhower et al. (2024), Oh et al. (2022), Sastry et al. (2023), Blonz et al. (2026), Koijen and Yogo (2015), Ge (2022), Coen et al. (2026), Collier et al. (2026), and Solomon (2026). Papers on the National Flood Insurance Program pricing include Wagner (2022a), Weill (2022), Mulder and Kousky (2023), and Mulder (2021). Marcoux and Wagner (2023) and Wagner (2022b) study insurance market design. Jung et al. (2023) measure insurers' exposure to climate risk. Abramson et al. (2025) show that guaranty fund assessments lead solvent insurers to restrict supply in states experiencing insurer insolvencies. See Sastry and Sen (2025) for a review.

A number of papers study the effect of disasters on displacement and rebuilding decisions (Deryugina et al., 2018; Boustan et al., 2020; Sheldon and Zhan, 2022; Aron-Dine, 2025). In concurrent work, Gruber and Solomon (2026) find that a recent reform to the National Flood Insurance Program, which directly conveyed property-level risk information and implemented a new pricing scheme, has an effect on net migration. We add to this literature in two ways. First, we identify relocation driven by the price of insuring against climate risk rather than by the realization of a disaster. Second, rather than net migration flows, we observe both sides of the reallocation: which households leave, where they go, and who buys the homes they vacate.

Work on how physical climate risk affects financial markets has emphasized direct damage to assets and the disruption of business operations as sources of loan defaults and bank losses (Acharya et al., 2023). We highlight a new channel: rising insurance premiums transmit climate risk to households and lenders even where no disaster occurs.

2 Institutional Background

Homeowners' insurance policies typically provide comprehensive coverage for the structure, contents, and temporary living expenses if the home becomes uninhabitable due to damage. These multi-peril policies typically have one-year terms with automatic renewal, subject to potential rate or term changes communicated through annual statements. Standard policies often exclude certain perils, such as floods and earthquakes. Specialized policies are available to cover these excluded risks.[11]

Around 86% of homeowners held such policies in 2023 and 2024; an industry survey puts the figure at 88%.[12] This high take-up rate is likely driven by mortgage lenders re-

[11] See https://content.naic.org/sites/default/files/publication-hoi-pp-consumer-homeowners.pdf.

[12] We compute the 86% from American Community Survey one-year county tables as the number of households reporting annual insurance premiums exceeding $100 divided by the number of households surveyed, averaged across counties; the one-year tables cover only high-population counties. For the 88% figure, see https://www.iii.org/sites/default/files/docs/pdf/2023_q2_ho_perception_of_weather_risks.pdf.

quiring insurance coverage as a precondition for loans. If insurance coverage lapses, the servicer may force-place a policy on the property at the borrower's expense. Lender and secondary market guidelines require servicers to maintain adequate insurance coverage or bear the resulting losses through indemnification or repurchase. In the same survey, 32% of homeowners reported being affected by weather events in the last five years, highlighting the importance of such coverage.

Homeowners' insurance policies typically include several key components that define the scope and extent of coverage. Homeowners can often adjust coverage limits and deductibles to change their insurance premiums. The coverage limit represents the maximum amount the insurer will pay in the event of a total loss, typically set to the estimated cost to rebuild the home. Lenders and secondary market institutions often impose specific policy requirements.[13] Deductibles are the amount the policyholder must pay out-of-pocket before the insurer begins to pay.

The cost of homeowners' insurance is primarily determined by location-specific risks, with states prone to natural disasters such as hurricanes and wildfires typically having higher premiums. For instance, Florida and Louisiana have some of the highest average annual premiums, exceeding $3,000, while Nevada has much lower average costs, $1,158 per year in 2023 (Table A.1). Nationwide, the average annual premium in our sample is $1,974 in 2023, with a standard deviation of $1,314, reflecting the wide dispersion in insurance costs across households. Other factors affecting pricing include the home's age, construction materials, proximity to fire stations, and the homeowner's claims history and credit score. Insurance companies use complex actuarial models to assess risk and set premiums, leading to variations in pricing among different insurers (Boomhower et al., 2024). Homeowners' insurance rates in the United States are subject to regulatory oversight and approval processes, which vary by state (Oh et al., 2022). Insurance

[13]For example, Fannie Mae requires property insurance coverage equal to the lesser of 100% of the insurable value of the improvements or the unpaid principal balance, so long as the latter is no less than 80% of that value, on a replacement cost basis. See `https://selling-guide.fanniemae.com/sel/b7-3-02/property-insurance-requirements-one-four-unit-properties`.

companies typically must submit rate change requests to state insurance departments or commissions for review and approval before implementing new premiums.

3 Data

3.1 Residential Mortgage Servicing (ICE McDash Analytics)

The ICE McDash Residential Mortgage Servicing Database comprises the servicing portfolios of the largest mortgage servicers in the US, and covers approximately two-thirds of installment-type loans in the residential mortgage servicing market. Loan-level attributes include borrower characteristics (credit scores, owner occupancy, and loan purpose); collateral characteristics (LTV, property type, ZIP code); and loan characteristics (loan balance and loan status). We restrict our sample to loans secured by single-family homes with non-missing ZIP code, origination date, LTV, DTI, and credit score. We also drop loans whose origination loan-to-value ratio exceeds 105 or debt-to-income ratio exceeds 65, which we treat as non-standard loans or possible reporting errors. We obtain these data for 2013–2023. Because we residualize our premium change instruments against their values from the prior two years, renewal years 2013–2015 do not enter the instrumental-variable regressions; 2014 and 2015 remain in the OLS regressions.

We rely on these data for the dependent variable of whether a mortgage is delinquent for at least 60 days. We drop loans whose servicing was transferred to another servicer within the year after a policy renewal, and loans terminated in that year for an unknown reason, as we cannot observe their outcome. Prepayment also includes refinancing.

To include a loan in our sample, we need to be able to compute the change in premium between policy renewals. Insurance information is attached to the December loan performance file, so each policy is observed once a year. The snapshot reports both the policy's effective date and its expiration date, so a single observation identifies the renewal month. The file is assembled about one month after the performance month, so the

insurance it carries is current as of the following January. A January renewal is therefore usually recorded in the December file immediately before it, while a renewal in any other month is recorded in the December file of its own calendar year. A loan enters our sample only if it is still active in that December performance month, so a loan whose policy renews in March and which is prepaid in June never enters the sample in that year. The survival required after renewal is ten months for a February renewal and one month less for each later renewal month, reaching zero in December. The assembly date is not a fixed calendar day and varies across servicers, so the preceding December file captures most January renewals but not all. A January renewal that the preceding file misses is first recorded in the December file of its own year, so the loan must survive eleven months, the longest requirement in the calendar. January is therefore the least affected renewal month on average rather than an unaffected one.

Figure IA.1 plots the share of mortgages that are prepaid by the calendar month of the insurance policy expiration date, among the loans that enter the sample, so a loan that never enters contributes to neither the numerator nor the denominator. The share is approximately 15% for policies renewing in January, falls to around 3% for February renewals, and then rises monotonically to 13% for December renewals. This is the pattern the survival requirement implies, and actual prepayment rates are unlikely to vary across insurance renewal months in this way. We therefore read the figure as measuring the share of the post-renewal year over which prepayment is observable, not a difference in prepayment behavior. To identify relocation we require mortgage prepayment information, so we restrict the relocation analyses to borrowers whose policies renew between September and January, the window over which prepayment is most accurately reported and the required survival after renewal is shortest, three months or less except for the minority of January policies noted above. Delinquency reporting does not exhibit this pattern, so the delinquency analyses use renewals in all months (Figure IA.2).

3.2 ICE McDash Property Insurance Module

We obtain data on insurance at the loan level from ICE McDash for 2013 through 2023. A loan identifier matches the insurance records to the servicing data, and the servicers sharing insurance information account for around two-thirds of the servicing sample.

Borrowers must submit proof of homeowners' insurance when the mortgage is closed and every year thereafter, and servicers record these data in the servicing system because a servicer may face penalties, such as repurchase of a defaulted loan, when coverage does not meet lender or secondary market requirements.

The ICE McDash Property Insurance data include the following variables: coverage amount, deductible amount, the date when the policy expires, and an identifier for the insurer. We restrict our sample to loans with annual insurance policies[14] that also appeared in the ICE McDash sample in the previous calendar year. We further require that the loan is current in the month before the insurance policy expires. Throughout the paper, a policy's *premium rate* is the annualized premium on the primary policy (twelve times the reported monthly premium amount) divided by the coverage limit, expressed in percentage points, and a change in the premium rate is a within-loan change between two consecutive renewals. Our sample contains around 230 insurers, identified by brand rather than by the state-level subsidiary that issues the policy. In a given year the average insurer accounts for 0.46% of the premium dollars in our sample nationally and for 1.48% within a state. The corresponding Herfindahl indexes of premium dollars average 593 nationally and 905 within a state and year. Table 1 reports their distributions.

3.3 ICE McDash Credit Risk Insights Servicing Module (CRISM)

CRISM matches the ICE McDash mortgage servicing records to anonymized credit bureau data, adding credit market outcomes such as credit card utilization and an indicator

[14] We classify a policy as annual when its expiration date falls 351–379 days after its effective date (i.e., one year $\pm$ 14 days).

of whether the borrower has moved residence, which we record when the mortgage terminates and the borrower's credit-bureau address changes, a new purchase mortgage appears in ICE McDash, or CoreLogic's transaction records show a sale of the property. Because every route requires the mortgage to terminate, a household that relocates temporarily, for example between a primary and a secondary residence, is not counted as a mover.

The data also provide Risk Scores, which differ from credit scores and are generated by a proprietary model that predicts the likelihood of a borrower becoming delinquent on credit obligations. Within the CRISM match, we keep borrowers with a single open first mortgage in the month before the policy expires.

To measure movers' outcomes, we identify a borrower's post-renewal replacement loan from its origination timing and location, excluding candidate loans secured by the same property or Census block, and loans originated outside the window running from the policy start month to four months after the prior loan terminates. Outcomes measured on the new property's own insurance policy require that replacement loan to appear in ICE McDash, and are estimated on a smaller sample. We treat credit-card utilization above 110% as missing.

3.4 CoreLogic Solutions Real Estate Data

We use CoreLogic's tax assessor and transaction data (Deeds and Tax Data) and CoreLogic Climate Risk Data, which applies a proprietary climate model to generate structure-level measures of current climate risk on nine perils, among them Average Annual Loss (AAL). AAL is the expected annual loss to structure and contents generated by simulating many possible iterations of a given year and then calculating the average loss across all iterations. CoreLogic reports AALs as a share of the replacement cost of the structure (Wylie et al., 2025). CoreLogic does not report past measures of climate risk, so each structure has a single measure of current risk, which we apply throughout the sample period.

3.5 Matching Loan-Level Data with CoreLogic

We merge our loan-level insurance data with CoreLogic's property-level data. This provides us with additional information about the property and its transaction history. We first identify a potential set of candidate properties by requiring an exact match on the 5-digit ZIP code, as well as a mortgage with the same month of origination and a loan amount rounded to the nearest thousand dollars. We exclude piggyback and subordinate-lien records from the candidate set.

We then rank candidate properties by how many of the following criteria they satisfy: the loan amounts match exactly; origination dates match exactly; loans have the same purpose (purchase or refinance); the loan type matches (FHA, VA or conventional). Ties are broken by a distance measure based on loan amount and origination date, and each loan takes its highest-ranked property. 85% of loans have a property match, 56% have an exact match on all criteria, and 80% of matches miss at most one additional criterion.

We use the matched data to measure the climate risk exposure (proxied by AAL) at the household level. When the property-level value is missing, or when the CoreLogic match is ambiguous in that more than one candidate property maps to the loan, we use the ZIP-by-state median, both for AAL and for the other CoreLogic characteristics we use as controls, namely the year the structure was built (house age), square footage, and property value. We sort households into three bins based on the AAL measure of their location prior to each insurance renewal. We then use climate risk bin-by-time fixed effects in our regression analyses. The climate-risk (AAL) measure underlying our standardized climate-risk control is winsorized at the 98th percentile due to large right tails. Where an analysis is conducted at the ZIP level rather than the household level, our climate-risk measure is the median climate risk across the properties in the ZIP code.

Table 1 reports summary statistics for the main household, mortgage, insurance, and climate-risk variables used in our analysis.

4 Premiums Increased More in High-Climate-Risk Areas

In Figure 1, we examine how premiums have evolved in recent years in relation to climate risk. We estimate the following regression at the ZIP-by-year level.

$$\Delta(Premium\ Rate)_{z,2013\ to\ t} = \gamma_z + \delta_{c,t} + \sum_{t=2015}^{2023} \beta_t ClimateRisk_z + \alpha_t Controls_z + \epsilon_{z,t},$$

where z denotes ZIP code, c denotes the county, and t the year. The dependent variable is the cumulative premium-rate change. $ClimateRisk_z$ is the median ZIP-level climate risk. We control for ZIP and county-by-year fixed effects, as well as the ZIP's rural share and log population density from the 2010 census, each interacted with year, and we cluster standard errors by state $\times$ year. The estimation sample runs from 2014 to 2023 with the 2014 climate-risk interaction omitted, so β_t is measured relative to 2014 and gives the additional cumulative premium-rate increase between 2014 and t associated with higher climate risk. Figure 1 plots β_t and its 90% confidence interval.

The figure suggests that premiums have grown faster in areas with higher climate risk, consistent with Keys and Mulder (2024). The gap is flat through 2019 and opens from 2020 onward: by 2023, ZIP-level climate risk at the 90th rather than the 10th percentile was associated with a cumulative premium-rate increase 0.016 percentage points larger, equal to 3.4% of the mean premium rate and 12% of the mean cumulative increase.

5 Main Specification

Our household-level analyses share the following OLS specification, in which the household outcome is measured over the 12 months following the insurance renewal.

$$Y_{i,t\ to\ t+12} = \beta \times Premium\ Rate\ Increase_{i,t} + Controls + FE_{ZIP,t} + OtherFEs + \epsilon_{i,t}, \quad (1)$$

where i indexes households, t indexes the insurance renewal month, and ZIP indexes the ZIP code. Our two main dependent variables, $Y_{i,t \text{ to } t+12}$, are an indicator of whether the household is delinquent on the mortgage for at least 60 days (Section 7) and an indicator of whether the household relocates (Section 8). We apply the same specification to further outcomes: credit card usage, measured over the same 12-month window, and foreclosure, measured over 24 months to allow time for the proceeding to begin. The foreclosure indicator marks the start of a foreclosure proceeding; it does not require the proceeding to reach a foreclosure sale.

The main independent variable, *Premium Rate Increase*, is the change in the premium rate at policy renewal relative to the previous policy year, and is set to zero when the premium rate falls. Decreases often reflect households choosing to decrease coverage and increase deductibles.

We control for borrowers' credit scores in the month prior to the insurance renewal, a control whose importance Duarte and Fonseca (2026) and Blonz et al. (2026) emphasize. In addition, we control for households' climate risk exposure at their residence and for the premium rate in the policy year before the renewal. By controlling for ZIP-by-renewal month fixed effects, we effectively compare households within the same ZIP code who renewed their insurance policies in the same month. We also control for a battery of other fixed effects: climate risk tercile bin-by-insurance renewal month, mortgage origination year, loan-to-value ratio bin, credit score at origination bin, and debt-to-income ratio bin.[15] Standard errors are clustered by state $\times$ year throughout.

6 Instrumental Variables for Premium Increase

The OLS relationship between higher insurance premiums and household outcomes may reflect omitted variables. Households in financial distress, for instance, are more likely

[15]The origination bins are constructed as follows: credit-score bins are 20 points wide and capped to the $[580, 800]$ range, DTI bins are 5 points wide with a floor at 10, and LTV bins are 5 points wide with values below 75 pooled into a single group.

to have homes in disrepair and lower credit scores, which leads insurers to charge them higher premiums (Blonz et al., 2026). Financial distress also contributes independently to mortgage delinquency. We mitigate this and other endogeneity concerns in our OLS regressions with the controls and fixed effects of Section 5, chief among them ZIP-by-renewal month. We further instrument a household's premium increase with the increases the same insurer implements elsewhere in the state.

To construct our instruments, for each borrower we calculate the average premium rate increase across policies that renew in the same month by the same insurer in other ZIP codes within the same state.[16] The first, the Other-ZIP instrument, uses all such ZIP codes; the second, the Far-ZIP instrument, uses only those whose distance from the focal ZIP code exceeds the median pairwise ZIP code distance in the state. Both instruments use each policy's expiration month and insurer from the prior year rather than its renewal month and current insurer. This reduces concerns about endogenous renewal timing and insurer switching. The two instruments trade first-stage power against the plausibility of the exclusion restriction.

Our instrument construction implies a first-stage coefficient below one. The leave-out average is computed from renewals with an unchanged insurer, deductible, and coverage type, so it measures the change the prior insurer implemented. By contrast, a household's own premium rate change, the endogenous regressor, is measured from the insurer and policy terms actually in force before and after the renewal. Households respond to a premium increase by switching insurer, raising the deductible, or moving from Replacement Cost Value to Actual Cash Value coverage, and a household that responds in any of these ways realizes less than the change its prior insurer implemented. Table IA.3 shows that both instruments predict switching at the current renewal. Policyholders elsewhere in the

[16]The leave-out averages are computed over owner-occupied and investor policies (excluding second homes and policies with unknown occupancy) that renew with the same insurer and with unchanged deductible and coverage type, so that each contributing observation reflects a within-policy premium change. Policies whose carrier is recorded under certain non-standard insurer codes (such as lender-placed or master policies and unidentified carriers) are treated as having a missing insurer and are excluded both from the instruments and from the insurer-switching indicators.

state respond in the same way, and the restriction on the leave-out sample excludes those that do, so the average measures the insurer's repricing on the policies that retained their terms. The first stage remains strong. Second-stage estimates are not affected by households' own responses or by the restriction on the leave-out sample.

Next, we residualize each instrument with respect to its 12- and 24-month lagged values and year-month fixed effects. Table IA.1 presents the results: both lagged coefficients are positive and statistically significant, indicating substantial serial correlation. Residualizing therefore isolates the component of an insurer's rate change that its own recent history does not predict, so that our estimates are identified from new rate revisions rather than from the predictable continuation of earlier ones.

Our instruments should satisfy the relevance restriction for two reasons. First, premium increases follow location-level trends driven by the risks insurers assess (Keys and Mulder, 2024), and those risks are correlated within a state: ZIP codes in the same state are exposed to common weather and catastrophe events, so an insurer that revises its assessment reprices across its book within the state. Second, premium increases follow state regulatory approvals of each insurer's rate change requests (Oh et al., 2022). Because premium changes affect homeowners only at renewal, a property's premium change should correlate with the adjustments the same insurer implements elsewhere in the state in the same renewal month. The Other-ZIP instrument uses all other ZIP codes in the state, so it is computed from the largest set of contributing policies and, as we show later, delivers the stronger first stage.

Our instruments likely satisfy the exclusion restriction for two reasons. First, our main regressions include ZIP-by-renewal month fixed effects, so we compare borrowers in the same ZIP code who renew in the same month. Within a ZIP code and renewal month, the average premium changes we use vary because insurers differ in when they reprice and when regulators approve their filings, not because of factors related to individual borrowers' mortgage delinquency or the other outcomes we examine. Second, our controls

and fixed effects largely remove potential confounders, including the time-varying effect of climate risk.

Some of the properties entering the Other-ZIP average, however, are located in the same metropolitan area as the focal household. Our ZIP-by-renewal month fixed effects absorb any local shock common to the cell, so proximity threatens the exclusion restriction only through the part that differs across insurers. For example, a carrier may write disproportionately in the neighborhoods a recent storm hit, or file a rate change on loss experience drawn partly from the focal household's own market, while also covering a distinctive set of borrowers there.

The Far-ZIP instrument further mitigates this concern, because the ZIP codes it averages over are more likely to be in different housing markets and outside the same disaster footprints. Distant ZIP codes in the same state can still share conditions with the focal household's ZIP code, though to a lesser extent. The average is computed from fewer policies, so the first stage is weaker, but its Kleibergen-Paap Wald F statistic remains far above conventional weak-instrument thresholds. Where we report a single instrument, we use the Far-ZIP one, because it is more likely to satisfy the exclusion restriction.

Figure IA.3 reports covariate-balance regressions of each instrument on predetermined household, house, and loan characteristics. The coefficients are small, and most are statistically insignificant despite our large sample size. The largest is on the lagged premium rate: a one-standard-deviation increase in that rate is associated with a change of about 0.02 standard deviations in the instrument. Every other coefficient is well below 0.01. Climate risk is among these characteristics, and neither instrument is correlated with it: the coefficient is 0.002 standard deviations for the Other-ZIP instrument and -0.0003 for the Far-ZIP instrument, neither statistically significant. Our instruments are therefore essentially orthogonal to these characteristics conditional on the fixed effects.

Even though our instrumental-variable results should not be driven by climate risk, our instruments can reflect insurers' repricing of climate-related risk. Insurers incorporate

updated assessments of a location's disaster risk at different times, so within a ZIP code and renewal month two carriers may quote different premiums for the same property because one has repriced the risk and the other has not. The ZIP-by-renewal month and climate risk bin-by-renewal month fixed effects should largely absorb the time-varying effect of climate risk and any recent change in it, so we argue the instruments move delinquency or relocation only through premiums. The risk-related content of that variation is informative about the effect of climate-driven repricing. Because households cannot easily distinguish the drivers of the increases they face, premium changes driven purely by climate risk are likely to have effects similar to those we document.

Two pieces of descriptive evidence support this account of the variation, in which insurers reprice the same climate risk at different times. Figure 2 plots, by renewal year and separately for each instrument, how within-ZIP-by-renewal-month dispersion across insurers relates to a ZIP code's climate risk. The relationship is positive in every year: insurers diverge more in how much they raise premiums in higher-risk areas, which is the variation our instruments exploit. The relationship strengthens over the sample period, nearly quadrupling between 2016 and 2023. Figure IA.4 shows that, within a ZIP code, insurers raise premiums at different times, and that insurers repricing later reach higher cumulative increases than those repricing earlier.[17] Our instruments can therefore reflect differences in the *timing* of repricing.

Both patterns indicate that a large part of these timing differences reflects the staggered pass-through of a shock common to the location rather than idiosyncratic movements in a single carrier's rates. A household whose insurer reprices early therefore faces an increase that the insurers serving its location will also implement. It observes its own premium rather than the quotes its neighbors receive, so it cannot separate the location-specific and insurer-specific components of the increase as our fixed effects do, and a

[17] Within each ZIP code, we rank insurers by how early they implement their eventual premium increase and split them into an earlier and a later half. The comparison assumes the two halves start from the same premium level. The earlier half raises rates faster at first, so for much of the sample a household insured with one of them can find a cheaper carrier among the later half.

larger increase raises its assessment of what insuring the location will cost. Switching to a carrier that has not yet repriced postpones such an increase rather than removing it. This is why the variation we use can affect relocation, in addition to insurer switching.

As we show later, the Far-ZIP first-stage coefficient is the smaller of the two. Internet Appendix IA.A explains why.

7 Effect of Premiums on Mortgage Delinquencies

Table 2 reports the effect of premium increases on mortgage delinquency. Column (1) estimates the specification of Section 5 by OLS, with the dependent variable being an indicator for being at least 60 days delinquent over the 12 months following the insurance renewal. Columns (2) and (3) report the first and second stages of the instrumental-variable specification with the Other-ZIP instrument; Columns (4) and (5) repeat both with the Far-ZIP instrument. The OLS coefficient on *Premium Rate Increase* is positive, suggesting that households facing larger premium increases are more likely to become delinquent.

Columns (2) and (4) report the first stage. Both instruments predict premium increases, with positive and statistically significant coefficients of 0.885 and 0.674: households face larger premium increases when their insurer raises premiums more in other or distant ZIP codes within the same state. The Kleibergen-Paap Wald F-statistics are large, indicating that the instruments are strong. Figure 3 shows bin-scatters of the premium rate increase against each instrument, after partialling out the controls and fixed effects: Other-ZIP in Panel A and Far-ZIP in Panel B.

Columns (3) and (5) present the second stage. The coefficients on the instrumented premium increase are positive and statistically significant, suggesting that premium increases cause higher mortgage delinquency rates. Figure 4 plots the same relationship as bin-scatters of delinquency against the fitted premium rate increase, from the Other-ZIP first stage in Panel A and from the Far-ZIP first stage in Panel B. In Column (5), a $1,000 increase in premium, equivalently a 0.2 increase in the premium rate, corresponds to a

0.38 percentage point increase in the 60-day delinquency rate, 22% of the sample mean.

The OLS coefficient in Column (1) of Table 2 is larger than both instrumental-variable estimates, consistent with endogeneity bias: households experiencing financial difficulties are likely to face larger premium increases, owing for instance to deferred maintenance, and to have elevated delinquency rates.

Columns (1) and (2) of Table 3 repeat this specification at two other delinquency thresholds: the same $1,000 increase raises the 30-day rate by 0.70 percentage points and the 90-day rate by 0.30 percentage points.

Five studies of mortgage payment shocks provide a benchmark for this magnitude. In the deeply underwater subsample of Fuster and Willen (2017), where the one-year incidence of 60-day delinquency is 17%, $1,000 of annual mortgage payment relief lowers that incidence by about 1.5 percentage points. In the conforming sample of Di Maggio et al. (2017), where the baseline incidence is 2.3%, the same $1,000 lowers it by about 0.16 percentage points. Our estimate is 0.38 percentage points on a baseline of 1.7%.

The 30- and 90-day estimates allow three further comparisons at matched outcome definitions. A $1,000 increase raises the 30-day delinquency rate by 0.70 percentage points in our sample, against 0.80 percentage points in Wong (2024), who studies escrowed property tax increases. It raises the 90-day rate by 0.30 percentage points, against 0.16 percentage points in Abel and Fuster (2021), who study borrowers refinancing into lower-rate mortgages through HARP, and 1.4 percentage points in Ganong and Noel (2020), who study payment reductions for distressed borrowers.[18]

Our estimate is neither the largest nor the smallest of this group in levels; once scaled by the baseline rate, it is comparable to Abel and Fuster (2021) and roughly 30% to 2.4 times larger than the other four.[19]

[18]Internet Appendix IA.C reports the figure each study provides, its source, and the arithmetic that places all six estimates at a common dose of $1,000 and, except for Wong (2024), at a one-year horizon.

[19]Similar to our paper, Mello (2024) also shows that even small unexpected fines can increase financial distress, especially among financially fragile households, and find that receiving fines for traffic violations averaging $195 is associated with a $34 increase in unpaid bills in collections. Gallagher et al. (2019) find that receiving health insurance subsidies is associated with a roughly 25% decline in the delinquency rate. Their

7.1 Delinquency Effect Across Investors

A natural question is whether the effect of insurance premiums on mortgage delinquencies is present for FHA and GSE loans. This would imply a credit risk for the federal government and GSEs, particularly if insurance premiums keep rising. We also examine private-label MBS and banks' portfolios, where higher delinquency is a credit risk to investors and to banks.

In Table A.2, we repeat the instrumental-variable specification across types of mortgages. We use Ginnie Mae loans in Column (1), Fannie Mae loans in (2), Freddie Mac loans in (3), private mortgages that are securitized in (4), and those held in banks' portfolios in (5). In Columns (6) and (7), we split conventional mortgages by whether they carry private mortgage insurance. The results suggest that the effect holds across subsamples, although it is not statistically significant for privately securitized mortgages, potentially due to a smaller sample size. Because the effect is present in government-backed and agency mortgages and in mortgages held in banks' portfolios, rising premiums pose credit risk to the federal government, the GSEs, and banks.

7.2 Foreclosure and Credit Card Usage

In Column (3) of Table 3, the instrumental-variable estimate shows that premium increases also raise the probability of foreclosure, measured over the 24 months after renewal to allow for the time the process takes. The estimate suggests that a $1,000 increase in premiums raises this probability by 0.053 percentage points, roughly 22% of its mean.

Column (4) suggests that a $1,000 increase in premiums raises credit card utilization by 0.2 percentage points of the credit limit. Column (5) replaces the dependent variable with the dollar change in credit card balances and the endogenous variable with the dollar premium increase, keeping the instrument the same. The estimate suggests that a $1,000

mechanism differs from ours, as it operates through reduced exposure to out-of-pocket medical expenditure risk rather than through direct insurance cost burdens.

increase in premiums raises credit card balances by $76 over the policy year.

These results indicate that rising insurance premiums affect household finances beyond delinquency, reaching foreclosure and credit card borrowing.

7.3 Heterogeneous Effects on Delinquency

We examine how the effect of premium increases on mortgage delinquency varies with household financial constraints. If premium increases are negative liquidity shocks, the effect should be concentrated among financially constrained households.

To test this prediction, Table 5 sorts households into three bins by credit score (Columns 1–2), loan-to-value ratio (Columns 3–4), and debt-to-income ratio (Columns 5–6). Each column replicates our main OLS specification and adds interactions between *Premium Rate Increase* and indicators for the second and third bins of the sort variable. The cutoffs are placed at fixed-effect bin boundaries, so that no fixed-effect bin spans two of the three bins. Odd-numbered columns report mortgage delinquency and even-numbered columns report relocation, the latter of which we discuss in Section 8.2.

The coefficient on *Premium Rate Increase* remains positive and statistically significant across all columns, so the baseline effects are present among households in the omitted, lowest bin of each sort variable. In Column (1), both interaction terms are negative and statistically significant, with magnitudes that largely offset the main effect. These results suggest that households in the second and third credit score bins show little delinquency response to premium increases, consistent with financially flexible households being able to absorb these shocks without becoming delinquent.

Two further proxies for financial constraints yield the same conclusion. The interactions are positive and significant for the middle and top bins of LTV (0.010 and 0.053 in Column 3) and of DTI (0.025 and 0.043 in Column 5), so the delinquency response to premium increases is 5.4 and 3.9 times as large for the most levered and most debt-burdened households as for the least.

7.4 Robustness

We address three concerns about the delinquency estimates. First, some insurers may be matched to households that are more likely to become delinquent, and may also raise premiums more. Columns (1)–(3) of Table A.3 repeat the OLS and both two-stage least-squares specifications with insurer-year fixed effects.

Second, households in disaster-affected areas may experience larger premium increases alongside higher delinquency, both consequences of the disaster. Our specification compares borrowers in the same ZIP code who renew in the same month, and both instruments are constructed from insurers' pricing changes in other or distant ZIP codes. Columns (1)–(3) of Table A.4 address the concern directly. Columns (1) and (2) estimate the effect separately for ZIP codes that experienced a natural disaster in the current or preceding two calendar years and for those that did not; Column (3) pools the two samples with an interaction between *Premium Rate Increase* and a recent-disaster indicator.

Third, the effect may reflect inattention rather than financial strain: households may fail to notice premium increases and underfund their escrow accounts.[20] Table A.5 restricts the sample to households that adjusted their deductible percentage or changed between Replacement Cost Value and Actual Cash Value coverage at the same renewal; these policy adjustments suggest awareness of the increase. We further restrict the sample to households that did not switch insurers, so that the increase we instrument is one they continued to bear. Our main delinquency results still hold in this sample. The range of outcomes affected argues against inattention as well: larger premium increases also raise foreclosure rates and credit card utilization (Section 7.2). Failing to notice a higher escrow payment would not produce responses across these margins, whereas financial strain would.

In all three tables, the estimated effect of *Premium Rate Increase* on delinquency remains

[20] Insurers notify homeowners of premium increases, but inattention remains plausible: Liu (2019) finds that mortgage borrowers are markedly less cost-sensitive to fees than to the salient interest rate, and that lenders price accordingly.

positive and statistically significant.

8 Effect of Premiums on Relocation

Section 7 suggests that premium increases push households into financial distress. Households can mitigate an increase by switching insurers, or by relocating. If premium increases lead households to leave high-risk locations, insurance pricing is a channel through which climate risk reallocates where people live, with consequences for the composition of those neighborhoods. We therefore examine whether households facing larger premium increases are more likely to move.

Table 4 estimates the main specification in Equation (1) on the matched ICE McDash–Credit Risk Insights Servicing (CRISM) sample, restricted as described in Section 3.1 to borrowers whose insurance renews between September and January, for whom the prepayment that identifies a move is most accurately observed. Column (1) presents the OLS regression. The dependent variable equals one if the borrower's mortgage terminates within 12 months of insurance renewal and the termination is accompanied by a change in the borrower's credit-record address, a replacement purchase loan, or a matched sale of the insured property; the primary independent variable is *Premium Rate Increase*. Columns (2) and (4) present the first-stage results of the instrumental variable regressions, and Columns (3) and (5) the second-stage results. Our instruments remain highly relevant in this smaller sample.

The OLS coefficient on *Premium Rate Increase* is positive and statistically significant: households facing larger premium increases are more likely to relocate. The coefficients are similar between the OLS and the instrumental variable regressions. Based on Column (5), a \$1,000 increase in premium corresponds to a 0.82 percentage point increase in relocation probability, equivalent to 17% of the mean. Figure 5 presents the corresponding bin-scatter plots of moving probability against the instrumented premium rate increase, fitted by the Other-ZIP first stage in Panel A and the Far-ZIP first stage in Panel B.

The mortgage lock-in effect on relocation documented by Fonseca and Liu (2024) provides a benchmark for this magnitude. To compare effects in dollar terms, premium rate changes must be scaled by the ratio of loan amount to coverage amount (mean of 1/3.3 in our sample). A one percentage point increase in mortgage rate deltas is thus equivalent to a 0.3 percentage point increase in premium rates. Our 2SLS estimate implies that a 0.3 percentage point premium rate increase raises moving probability by 1.2 percentage points. For the same dollar change in moving incentives, Fonseca and Liu (2024) find that a one percentage point increase in mortgage rate deltas raises the moving rate by 0.69 percentage points over their 2010–2024 sample and by 1.20 over 2022–2024.

8.1 Moving versus Switching Insurers

Why would a household bear the large cost of relocation rather than search for a cheaper insurer? Three considerations can explain it.

First, switching recovers only part of a premium increase: for a $1,000 increase, Ge et al. (2026c) estimate a reduction of about $400 in the following year's premium. A household can leave an unfavorable carrier, but insurers serving its location price broadly the same underlying risk, so a search does not offset the part of an increase that is common to the area. The relief is partly temporary: because insurers in the same ZIP code reprice at different times (Section 6), a household leaving an insurer that has already raised its rates selects another insurer that has not yet done so. The new insurer raises its rates subsequently, and the initial saving declines. For some households a search may yield no relief. A large increase does not mean a high premium: a carrier that previously charged well below its competitors can remain the cheapest even after raising its rates, leaving the household no cheaper quote to switch to.

Relocation, by contrast, is not limited in these ways: it changes the location-specific risk being priced. It also delivers what a search would have delivered, because a relocating household purchases a new policy on a different property and often selects a new

insurer: 38% of movers in our data change insurer. For movers with larger pre-move increases, premiums fall by close to the full increase (Section 8.5).

Second, households are unlikely to move for insurance alone. The present value of the savings adds to their accumulated, idiosyncratic reasons to relocate, such as a job change or a growing family, so it changes the decision only for households already near their moving threshold. Consistent with this, the switching response is far more sensitive: the same $1,000 increase raises the probability of switching insurers by 19.0 percentage points, 306% of its mean, roughly 18 times the proportional relocation response (Column 5 of Table IA.2).

Third, potential buyers may value a property above its original owner, which can create a gain from selling. Some buyers may have a comparative advantage in owning properties with large premium increases, as we discuss in Section 8.6. Original owners may also extrapolate from recent increases and expect larger future increases than buyers do. We formalize the choice between the two responses in a model, presented in Internet Appendix IA.B.

8.2 Heterogeneous Effects on Relocation

We now examine how the effect of premium increases on relocation varies with household financial constraints, using the same sorts as Section 7.3. Constrained households may benefit more from escaping high-premium areas, but relocation entails substantial upfront costs that they may be unable to cover. Less constrained households may also find relocation easier if they have greater flexibility to change jobs or work remotely.

Column (2) of Table 5 reports the credit score sort for relocation. Both interaction terms are positive and statistically significant: the relocation response is 76% larger for the second credit score bin and 133% larger for the third than for the lowest. Sorting by loan-to-value and debt-to-income points in the same direction. In Columns (4) and (6), the top-bin interactions are negative and statistically significant: households in the lowest

LTV and DTI bins, the least constrained, are 43% and 50% more sensitive to premium increases than the most constrained.

Less constrained households are therefore better positioned to act on relocation incentives, while premium increases push constrained households into delinquency instead (Section 7.3).

8.3 Relocation Effect Depends on Mortgage Rate Lock-In

Relocation requires the household to give up its existing mortgage and obtain a new one. The cost of doing so depends on the origination rate relative to the rate currently available: a household paying an above-market rate loses little, while a household paying a below-market rate forfeits that rate. Fonseca and Liu (2024) show that this lock-in reduces household mobility. The relocation response to premium increases should therefore be stronger for households whose origination rate is high relative to current rates.

Table 6 tests this prediction, repeating the specification of Table 4 with an interaction between *Premium Rate Increase* and the difference between the origination rate and the average 30-year mortgage rate in the month of insurance renewal. The interaction is positive in all three columns. It is statistically significant in the OLS regression (Column 1) and in the Far-ZIP instrumental-variable regression (Column 3), and less precisely estimated in the Other-ZIP regression (Column 2). A one standard deviation larger rate gap, 1.64 percentage points, raises the relocation response to a premium increase by 0.011 in Column (3), 21% of the main effect. These estimates suggest that the least locked-in households respond most strongly.

Mortgage rates have risen in recent years alongside large increases in insurance premiums, so lock-in may have limited the relocation response we estimate. We expect that response to strengthen as mortgage rates fall and lock-in weakens.

8.4 Robustness

The first two concerns addressed in Section 7.4, insurer selection and recent disasters, also apply to the relocation results. Columns (4)–(6) of Table A.3 add insurer-year fixed effects to the OLS and both two-stage least-squares specifications, and Columns (4)–(6) of Table A.4 repeat the disaster split. The coefficient on *Premium Rate Increase* remains positive and statistically significant in both tables.

8.5 Premium Increases Drive Relocation to Safer Areas

If households move due to premium increases, they should achieve premium savings after moving. We first show that they do, then ask how their destinations differ. Table 7 presents second-stage results for mover outcomes. In Column (1), the dependent variable is the change in premium rates from policy year t to $t+1$, with the initial *Premium Rate Increase* on the right-hand side measured from policy year $t-1$ to t.

We use the full sample, which contains both movers and stayers, and add an interaction between the initial *Premium Rate Increase* and an indicator for moving within 12 months after renewal.[21] The positive coefficient on *Premium Rate Increase*, 0.280, indicates that for non-movers, larger premium increases from $t-1$ to t are associated with further premium increases from t to $t+1$. This pattern can emerge when households exhibit sticky insurer choice and certain insurers implement consecutive rate increases more often than others.

Movers show the opposite pattern. The coefficient on *Premium Rate Increase* $\times$ *Moved* is -1.184 and statistically significant, more than offsetting the positive main effect, so for movers a larger premium increase is followed by a larger subsequent premium reduction. A \$1,000 increase in premium therefore leads movers to reduce their subsequent premium rate by 0.24 percentage points more than non-movers. That is a saving of \$1,184 per year.

[21] The instrumental-variable specification therefore has two endogenous variables and two first stages. The interaction between our instrument and the moving indicator is the second instrument in both first stages.

Holding the saving constant and discounting at 6% (Giglio et al., 2021), the present value over 20 years is $13,580.

Movers with larger pre-move premium increases may save on premiums because they relocate to safer areas. Columns (2)–(5) of Table 7 examine their destinations and the insurance markets they move into: climate risk and average premium rates in Columns (2) and (3), and two alternative routes to lower insurance costs in Columns (4) and (5).

In Column (2), the dependent variable is the change in climate risk. The destination and origin values are measured at the property level where the loan appears in McDash with a good CoreLogic match, and at the ZIP median otherwise, with the destination ZIP taken from the borrower's credit-bureau address. Columns (2)–(5) restrict the sample to movers and, because the mover sample is much smaller, replace two fixed effects of our main specification with pre-move state-by-year and climate risk bin-by-year fixed effects, retaining our other controls. The coefficient on *Premium Rate Increase* is negative and statistically significant. The magnitude suggests that movers with a $1,000 larger pre-move increase lower their climate risk exposure by 17% of the mean.

Our measure of climate risk can be noisy: it may not accurately reflect a property's risk, or insurers' perception of that risk. Column (3) therefore uses an alternative proxy for movers' change in climate risk exposure: the change from their pre-move premium rate to the average in their new ZIP code, assuming that premium rates largely reflect insurers' assessments of local risks and average house quality. Using average premiums reduces the measure's sensitivity to lower teaser rates for new customers. The new ZIP average premium can also be correlated with factors unrelated to climate risk, which Columns (4) and (5) address.

The coefficient on *Premium Rate Increase* in Column (3) is also negative and statistically significant. The magnitude suggests that for the same $1,000 larger pre-move increase, the average premium rate in the new ZIP code is 0.24 percentage points lower than the mover's old rate, half the average premium rate of 0.48.

Movers with larger pre-move premium increases may also lower their insurance costs by moving to states where insurers charge less. Column (4) replaces the dependent variable with changes in insurers' average premium rates in other ZIP codes within the same state, with the state changing for across-state movers. For the 62% of movers who keep their insurer, this change is that insurer's rate in the new state after the move minus its rate in the old state before the move. The coefficient on *Premium Rate Increase* is statistically insignificant. This gives no evidence that movers with larger pre-move increases relocate to states where insurance is cheaper for reasons such as construction costs or regulation, which would appear in insurers' average premium rates in other ZIP codes of the destination state.

They can also reduce insurance costs by moving to ZIP codes with a more competitive insurance market. Column (5) replaces the dependent variable with the change in the insurer Herfindahl-Hirschman Index (HHI) from the old to the new ZIPs. The coefficient on *Premium Rate Increase* is small and statistically insignificant, if anything positive, providing no evidence of relocation to less concentrated insurance markets.

Taken together, movers who face larger pre-move increases relocate to ZIP codes with lower climate risk and lower average premium rates. Neither the rate-setting of their post-move insurer nor the concentration of the destination market changes systematically with pre-move premium increases.

8.6 Investors Purchase Houses with Large Premium Increases

Who buys the homes that these movers vacate? We hypothesize that buyers are more likely to be investors when the premium increase is larger, because investors may have greater market power or lower search costs in obtaining lower premiums.

We test this hypothesis in Table 8, where each observation is a property transaction. We link a vacated home's sale to the departing borrower's premium change using CoreLogic transactions that occur within about two months of that borrower's mortgage termi-

nation and in the twelve months following the policy start month. We classify the buyer as an investor when CoreLogic's owner-transfer records identify the grantee as a market investor *and* the purchase is financed outside the standard mortgage channel, meaning either an all-cash purchase or a mortgage from a private-party lender.[22]

Column (1) reports the OLS regression of this indicator on *Premium Rate Increase*. Columns (2) and (3) present the first and second stages of the 2SLS using the residualized Other-ZIP instrument, and Columns (4) and (5) repeat them with the Far-ZIP instrument. Because the transaction sample is much smaller than the loan-year samples of Sections 7 and 8, we replace two fixed effects of our main specification with a 5-digit ZIP fixed effect and a climate risk bin-by-renewal year fixed effect, retaining our other controls.

The coefficient on *Premium Rate Increase* is positive and statistically significant in Column (1) and both instrumental-variable results. The magnitude in Column (5) suggests that a $1,000 increase in premiums raises the probability that the buyer is an investor by 0.90 percentage points, a 22% increase relative to the sample mean.

8.6.1 Investors' Insurance Outcomes and Advantages

Table 9 examines whether investors hold advantages in lowering insurance costs and systematically make different insurance choices. Each observation is an insurance policy year, so a property enters once for each annual policy we observe. The main independent variables are an investor indicator and its interaction with the Census block-group-year average of premium rate increases relative to the prior year, standardized to have mean

[22]We build this indicator from CoreLogic's owner-transfer records rather than from McDash's non-owner-occupancy designation, which is observed only for buyers who can be matched back to a McDash record and would therefore restrict the analysis to a subsample of transactions. Market investors are identified by applying a name-pattern classification to the grantee name in CoreLogic's owner-transfer records: a buyer is institutional if its name identifies a business entity, and is counted as a market investor if it is institutional and is not a government agency, nonprofit, trust or estate, religious organization, or homeowner association. The remaining institutional categories are real-estate and property-investment firms, financial institutions, construction and development companies, generic corporate entities, and so on. Requiring the financing condition in addition to the investor classification restricts the indicator to a more sophisticated subset of investor purchases. It also excludes owner-occupant households that purchase through a limited liability company or a similar entity and finance the purchase with a standard mortgage, which the name-pattern classification alone would count as investors.

zero and standard deviation one. We use McDash's investor designation rather than the CoreLogic buyer classification of Table 8 because it does not require the property to have transacted during our sample period and matched between CoreLogic and McDash, and is therefore available for almost all the mortgages that can be linked to the insurance records. Because the insurance information comes from the servicers of these mortgages, the sample covers only properties financed with a standard mortgage. Many large investors do not use such financing and fall outside it. The investors we exclude may hold larger advantages than the ones we estimate. Search costs, bargaining power with carriers, and the capacity to bear losses without insurance all improve with portfolio size, and investors who do not borrow through a standard mortgage may not face the lender-imposed coverage minimums that apply to the properties we observe.

Additional controls include climate risk, pre-renewal credit score, an indicator for properties under five years old, log square footage, and log assessed property value. Fixed effects include house age (10-year) bins, Census block-group-by-year, climate risk bin-by-year, origination year, credit score bin, LTV bin, and DTI bin.

Investors do not select properties randomly, but the controls partially absorb the property characteristics relevant to insurance outcomes. Census block-group-by-year fixed effects further mitigate endogeneity concerns by comparing homes within the same block group in the same year.

Column (1), where the outcome is log premiums, yields a negative and significant coefficient on the investor indicator: at the mean local premium rate increase, investors pay 14% lower total premiums. Column (2), where the outcome is the premium rate, corroborates this pattern: investors pay lower premium rates, around 8% relative to the sample average. Column (3), where the outcome is the premium rate increase from the prior renewal, similarly shows slower premium growth for investors than for other homeowners in the same Census block group, by 19% of the mean.

The interaction terms in Columns (1)–(3) are all negative and statistically significant,

so each advantage is larger in blocks where premiums have increased more. Between two Census block groups one standard deviation apart in the average premium increase, the investor advantage widens by 3.2% of total premiums in Column (1), by 5% of the mean premium rate in Column (2), and by 17% of the mean premium rate increase in Column (3).

The dollar premium advantages may reflect investors' choices of coverage: lower coverage limits reduce total premiums. In Column (4), the dependent variable is the log coverage limit. Both the investor indicator and the interaction term carry negative, statistically significant coefficients: investors select coverage limits 7% lower than owner-occupants, and lower still when local premiums are rising faster.

Investors can also choose higher deductibles to lower premium rates and lower total premiums. In Column (5), the dependent variable is the deductible relative to the coverage limit. The investor indicator is positive and statistically significant, indicating that investors choose deductible ratios 21% above the sample mean. The interaction with local premium increases is also positive and statistically significant, so this gap widens where premiums are rising faster.

Investors may also obtain lower premium rates through superior price search and greater bargaining power, including scale discounts for insuring multiple properties. Their tendency to select lower coverage amounts and higher deductibles may reflect three factors that reduce the protection they want. First, investors may have greater financial flexibility, reducing their effective risk aversion and making high-deductible plans more attractive. Second, smaller mortgage balances may lower lender-imposed insurance minimums. Third, investors may face lower disaster repair costs by maintaining contractor relationships, employing dedicated repair staff, or handling repairs directly.

Investors can obtain lower coverage limits through three channels. The first is the endorsements that scale a policy's limit over time. Owner-occupants commonly purchase extended or guaranteed replacement cost riders (Cookson et al., 2024), which pay above

the stated dwelling limit, together with inflation-guard provisions that escalate that limit automatically at each renewal in step with local construction costs. An owner-occupant's limit therefore can exceed insurers' estimated rebuilding costs, while an investor may not purchase these riders. The second is the settlement basis: investors may choose to accept actual cash value rather than replacement cost settlement. Under the former, the insurer pays the depreciated rather than the full cost of repair, and policies are typically issued at limits anchored to the property's depreciated value rather than to the cost of rebuilding it new. The third is sorting across carriers. Insurers differ in the coverage they extend for an otherwise comparable property (Cookson et al., 2024). Investors could sort toward insurers that assign lower estimated rebuilding costs in order to reduce their premiums.

Taken together, the results in Table 9 indicate that investors hold advantages over owner-occupants in containing insurance costs, and that these advantages widen as local premium increases become larger. They therefore offer a partial explanation for why homeowners facing larger premium increases may be more likely to sell to investors: investors may value these properties more highly than owner-occupants do.

9 Conclusion

This paper examines how rising homeowners' insurance premiums affect household financial outcomes. Using household-level microdata and two novel instruments, we document three main findings. First, households facing larger premium increases are more likely to relocate to areas with lower climate risk and lower insurance costs. Homes experiencing larger increases are more likely to be acquired by investors, who obtain lower premium rates and purchase lower coverage limits. Second, larger premium increases raise rates of mortgage delinquency. Third, these effects vary systematically with financial constraints, measured by credit score, loan-to-value ratio and debt-to-income ratio at origination: the relocation response is stronger among less constrained households, who are better able to afford the upfront costs of moving, while the delinquency effect is

stronger among more constrained households.

Our findings on mortgage delinquency reveal a channel through which climate change can affect households and the broader financial system. As disasters become more severe and frequent, they push insurance rates higher, placing financial strain on households that do not directly experience them. This mechanism poses an emerging risk to both the federal government and the private financial sector. These results bear on the policy debate over insurance affordability, in which policymakers are weighing interventions such as means-tested insurance subsidies.

Our results on household relocation carry long-term implications. If premiums continue to rise in high-risk areas, the differential relocation response we document may reshape residential sorting, with less constrained households relocating away and more constrained households left increasingly concentrated in those areas. In addition, investors' competitive advantage in these areas, stemming from lower premiums, may accelerate a shift away from individual homeownership. Given evidence that homeownership facilitates wealth accumulation and generates positive community externalities, these dynamics could exacerbate inequality and erode local amenities in high-risk areas over the coming decades.

References

Abel, Joshua and Andreas Fuster, "How Do Mortgage Refinances Affect Debt, Default, and Spending? Evidence from HARP," *American Economic Journal: Macroeconomics*, 2021, *13* (2), 254–291.

Abramson, Boaz, Parinitha Sastry, Ishita Sen, and Ana-Maria Tenekedjieva, "Insurance Guaranty Funds," *Available at SSRN 5806662*, 2025.

Acharya, Viral V, Richard Berner, Robert Engle, Hyeyoon Jung, Johannes Stroebel, Xuran Zeng, and Yihao Zhao, "Climate stress testing," *Annual Review of Financial Economics*, 2023, *15* (1), 291–326.

An, Xudong, Stuart A Gabriel, and Nitzan Tzur-Ilan, "Extreme Wildfires, Distant Air Pollution, and Household Financial Health," Working Paper, Federal Reserve Bank of Philadelphia 2024.

Aron-Dine, Shifrah, "Rebuild or Relocate? Recovery after Natural Disasters," Working paper, University of California, Berkeley 2025.

Artigue, Heidi, Patrick Bayer, Fernando V Ferreira, and Stephen Ross, "Does homeownership matter? The long-term consequences of losing a house during the Great Recession," Technical Report, National Bureau of Economic Research 2025.

Bakkensen, Laura A and Lint Barrage, "Going underwater? Flood risk belief heterogeneity and coastal home price dynamics," *The Review of Financial Studies*, 2021.

Baldauf, Markus, Lorenzo Garlappi, and Constantine Yannelis, "Does Climate Change Affect Real Estate Prices? Only If You Believe in It," *The Review of Financial Studies*, 2020, *33* (3), 1256–1295.

Bernstein, Asaf, Matthew T. Gustafson, and Ryan Lewis, "Disaster on the Horizon: The Price Effect of Sea Level Rise," *Journal of Financial Economics*, 2019, *134* (2), 253–272.

Bilal, Adrien and Esteban Rossi-Hansberg, "Anticipating Climate Change Across the United States," NBER Working Paper 31323, National Bureau of Economic Research 2025.

Billings, Stephen B, Emily A Gallagher, and Lowell Ricketts, "Let the rich be flooded: The distribution of financial aid and distress after Hurricane Harvey," *Journal of Financial Economics*, 2022, *146* (2), 797–819.

Biswas, Siddhartha, Mallick Hossain, and David Zink, "California wildfires, property damage, and mortgage repayment," *Federal Reserve Bank of Philadelphia Working Paper*, 2023, (23-05).

Blonz, Joshua, Mallick Hossain, Benjamin J Keys, Philip Mulder, and Joakim A Weill, "Pricing protection: Credit scores, disaster risk, and home insurance affordability," Technical Report, National Bureau of Economic Research 2026.

Boomhower, Judson, Meredith Fowlie, Jacob Gellman, and Andrew Plantinga, "How are insurance markets adapting to climate change? risk selection and regulation in the market for homeowners insurance," Technical Report, National Bureau of Economic Research 2024.

Boustan, Leah Platt, Matthew E Kahn, Paul W Rhode, and Maria Lucia Yanguas, "The effect of natural disasters on economic activity in US counties: A century of data,"

Journal of Urban Economics, 2020, *118*, 103257.

Boutros, Michael, Nuno Clara, and Pierre Mabille, "Natural Disasters, Property Insurance, and Housing Markets," Working paper, University of Toronto, Duke University, and INSEAD 2026.

Campbell, John Y and Tarun Ramadorai, "Household Finance in Retrospect and Prospect," *Journal of Finance: Insights and Perspectives*, 2025. Forthcoming. NBER Working Paper No. 34621.

Coen, Jamie, Alessandro Gavazza, Daniel Gottlieb, and Tarun Ramadorai, "Regulating Inaction: The Case of Price Walking," 2026. CEPR Discussion Paper No. 21357.

Collier, Benjamin L., Cameron M. Ellis, Anran Li, and Adam Solomon, "Capital Costs, Reinsurance, and the Price of Climate Risk," Working Paper, University of Wisconsin–Madison 2026.

Cookson, J Anthony, Emily Gallagher, and Philip Mulder, "Coverage Neglect in Homeowners Insurance," *Available at SSRN 5057551*, 2024.

Cruz, José-Luis and Esteban Rossi-Hansberg, "The economic geography of global warming," *The Review of Economic Studies*, 2024, *91* (2), 899–939.

Deryugina, Tatyana, Laura Kawano, and Steven Levitt, "The economic impact of Hurricane Katrina on its victims: Evidence from individual tax returns," *American Economic Journal: Applied Economics*, 2018, *10* (2), 202–233.

Desmet, Klaus and Esteban Rossi-Hansberg, "On the spatial economic impact of global warming," *Journal of Urban Economics*, 2015, *88*, 16–37.

_ , Robert E Kopp, Scott A Kulp, Dávid Krisztián Nagy, Michael Oppenheimer, Esteban Rossi-Hansberg, and Benjamin H Strauss, "Evaluating the economic cost of coastal flooding," *American Economic Journal: Macroeconomics*, 2021, *13* (2), 444–486.

Di Maggio, Marco, Amir Kermani, Benjamin J Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao, "Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging," *American Economic Review*, 2017, *107* (11), 3550–3588.

Dice, Jacob, Mallick Hossain, and David Rodziewicz, "Flood Risk Exposures and Mortgage-Backed Security Asset Performance and Risk Sharing," *Available at SSRN 4854082*, 2024.

DiPasquale, Denise and Edward L Glaeser, "Incentives and social capital: Are homeowners better citizens?," *Journal of Urban Economics*, 1999, *45* (2), 354–384.

Duarte, Victor and Julia Fonseca, "Self-Fulfilling Credit Scores," Technical Report 35508, National Bureau of Economic Research 2026.

Efron, Bradley, "Increasing properties of Pólya frequency functions," *The Annals of Mathematical Statistics*, 1965, *36* (1), 272–279.

Fonseca, Julia and Lu Liu, "Mortgage Lock-In, Mobility, and Labor Reallocation," *The Journal of Finance*, 2024, *79* (6), 3729–3772.

_ , _ , and Pierre Mabille, "Unlocking Mortgage Lock-In: Equilibrium Effects in a Spatial Housing Ladder Model," Technical Report 35237, National Bureau of Economic Research 2026. Conditionally accepted, *Review of Financial Studies*.

Fuster, Andreas and Paul S. Willen, "Payment Size, Negative Equity, and Mortgage Default," *American Economic Journal: Economic Policy*, 2017, *9* (4), 167–191.

Gallagher, Emily A, Radhakrishnan Gopalan, and Michal Grinstein-Weiss, "The effect of health insurance on home payment delinquency: Evidence from ACA Marketplace subsidies," *Journal of Public Economics*, 2019, *172*, 67–83.

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

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

_ and _ , "Why do borrowers default on mortgages?," *The Quarterly Journal of Economics*, 2023, *138* (2), 1001–1065.

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

_ , Ammon Lam, and Ryan Lewis, "The Effect of Insurance Premiums on the Housing Market and Climate Risk Pricing," *Available at SSRN 4192699*, 2024.

_ , Ankit Kalda, Varun Sharma, and Vikas Soni, "Uneven Adaptation: Insurance Pricing and Household Climate Resilience," Technical Report 109, Hutchins Center on Fiscal & Monetary Policy at Brookings 2026a.

_ , Ryan Lewis, and Dean Parker, "Does Disaster Aid Cause Moral Hazard in the Housing Market?," *Working Paper*, 2026b.

_ , Stephanie Johnson, Nitzan Tzur-Ilan, and Derek Wenning, "Race to the Bottom: How Climate Risk and Price Competition Erode Insurer Quality," *Working Paper*, 2026c.

Georgic, Will and H Allen Klaiber, "Stocks, flows, and flood insurance: A nationwide analysis of the capitalized impact of annual premium discounts on housing values," *Journal of Environmental Economics and Management*, 2022, *111*, 102567.

Gete, Pedro, Athena Tsouderou, and Susan M Wachter, "Climate risk in mortgage markets: Evidence from Hurricanes Harvey and Irma," *Real Estate Economics*, 2024, *52* (3), 660–686.

Gibson, Matthew and Jamie T Mullins, "Climate risk and beliefs in New York floodplains," *Journal of the Association of Environmental and Resource Economists*, 2020, *7* (6), 1069–1111.

Giglio, Stefano, Matteo Maggiori, Krishna Rao, Johannes Stroebel, and Andreas Weber, "Climate change and long-run discount rates: Evidence from real estate," *The Review of Financial Studies*, 2021, *34* (8), 3527–3571.

Gomes, Francisco, Michael Haliassos, and Tarun Ramadorai, "Household finance," *Journal of Economic Literature*, 2021, *59* (3), 919–1000.

Gruber, Jon and Adam Solomon, "Optimal Flood Insurance in a Second-Best World: Fiscal Spillovers, Reclassification Risk and Moral Hazard," Working Paper, New York University 2026. NYU Working Paper.

Gupta, Arpit, "Foreclosure contagion and the neighborhood spillover effects of mortgage

defaults," *The Journal of Finance*, 2019, *74* (5), 2249–2301.

Indarte, Sasha, "Moral hazard versus liquidity in household bankruptcy," *The Journal of Finance*, 2023, *78* (5), 2421–2464.

Issler, Paulo, Richard Stanton, Carles Vergara-Alert, and Nancy Wallace, "Mortgage Markets with Climate-Change Risk: Evidence from Wildfires in California," *Working Paper*, 2024.

Jotikasthira, Chotibhak, Anastasia Kartasheva, Christian T Lundblad, and Tarun Ramadorai, "Strategic Claim Payment Delays: Evidence from Property and Casualty Insurance," 2025.

Jung, Hyeyoon, Robert F Engle, Shan Ge, and Xuran Zeng, "Physical Climate Risk Factors and an Application to Measuring Insurers' Climate Risk Exposure," Technical Report, Staff Report 2023.

Kalda, Ankit, Varun Sharma, Vikas Soni, and Derek Wenning, "Beyond the Storm: Climate Risk and Homeowners' Insurance," *Available at SSRN 5138641*, 2025.

Keys, Benjamin J, "The Credit Market Consequences of Job Displacement," *The Review of Economics and Statistics*, 2018, *100* (3), 405–415.

_ and Philip Mulder, "Neglected no more: Housing markets, mortgage lending, and sea level rise," Technical Report, National Bureau of Economic Research 2020.

_ and _ , "Property Insurance and Disaster Risk: New Evidence from Mortgage Escrow Data," Technical Report, National Bureau of Economic Research 2024.

Kleiner, Kristoph, Noah Stoffman, and Scott E Yonker, "Friends with bankruptcy protection benefits," *Journal of Financial Economics*, 2021, *139* (2), 578–605.

Koijen, Ralph SJ and Motohiro Yogo, "The cost of financial frictions for life insurers," *American Economic Review*, 2015, *105* (1), 445–475.

Kousky, Carolyn, Mark Palim, and Ying Pan, "Flood Damage and Mortgage Credit Risk: A Case Study of Hurricane Harvey," *Journal of Housing Research*, 2020, *29* (S1), S86–S120.

Liao, Yanjun and Philip Mulder, "What's at Stake? Understanding the Role of Home Equity in Flood Insurance Demand," 2021.

Liu, Lu, "Non-Salient Fees in the Mortgage Market," Technical Report 819, Bank of England Staff Working Paper 2019.

Lopez, Luis A and Nitzan Tzur-Ilan, "Air pollution and rent prices: Evidence from wildfire smoke," *Available at SSRN 4537395*, 2023.

Marcoux, Kendra and Katherine RH Wagner, "Fifty years of US natural disaster insurance policy," 2023.

Mello, Steven, "Fines and financial wellbeing," *Review of Economic Studies*, 2024, p. rdae111.

Mulder, Philip, "Mismeasuring risk: The welfare effects of flood risk information," Working Paper, University of Pennsylvania 2021.

_ and Carolyn Kousky, "Risk rating without information provision," in "AEA Papers and Proceedings," Vol. 113 American Economic Association 2023, pp. 299–303.

Murfin, Justin and Matthew Spiegel, "Is the Risk of Sea Level Rise Capitalized in Residential Real Estate?," *The Review of Financial Studies*, 2020, *33* (3), 1217–1255.

Oh, Sangmin, Ishita Sen, and Ana-Maria Tenekedjieva, "Pricing of climate risk insurance: Regulation and cross-subsidies," *Available at SSRN 3762235*, 2022.

Sastry, Parinitha, "Who Bears Flood Risk? Evidence from Mortgage Markets in Florida," *The Review of Financial Studies*, 2026, *39* (8).

_ and Ishita Sen, "Perspectives on Insurance and Climate Risk," *Working Paper*, 2025.

_ , _ , Ana-Maria Tenekedjieva, and Therese C Scharlemann, "Climate Risk and the U.S. Insurance Gap: Measurement, Drivers and Implications," *Available at SSRN 4909444*, 2026.

_ , _ , and _ , "When insurers exit: Climate losses, fragile insurers, and mortgage markets," *Working Paper*, 2023.

Sheldon, Tamara L and Crystal Zhan, "The impact of hurricanes and floods on domestic migration," *Journal of Environmental Economics and Management*, 2022, *115*, 102726.

Sodini, Paolo, Stijn Van Nieuwerburgh, Roine Vestman, and Ulf von Lilienfeld-Toal, "Identifying the benefits from homeownership: A Swedish experiment," *American Economic Review*, 2023, *113* (12), 3173–3212.

Solomon, Adam, "When Insurance Markets Fail: Catastrophe-Risk Frictions and Public Reinsurance," Working Paper, New York University 2026.

Van Binsbergen, Jules H., Joao Cocco, Marco Grotteria, and S. Lakshmi Naaraayanan, "The impact of carcinogenic risk exposure on housing values: Estimates from chemical reclassifications," *Review of Financial Studies*, 2025. Forthcoming.

Wagner, Katherine RH, "Adaptation and adverse selection in markets for natural disaster insurance," *American Economic Journal: Economic Policy*, 2022a, *14* (3), 380–421.

_ , "Designing insurance for climate change," *Nature Climate Change*, 2022b, *12* (12), 1070–1072.

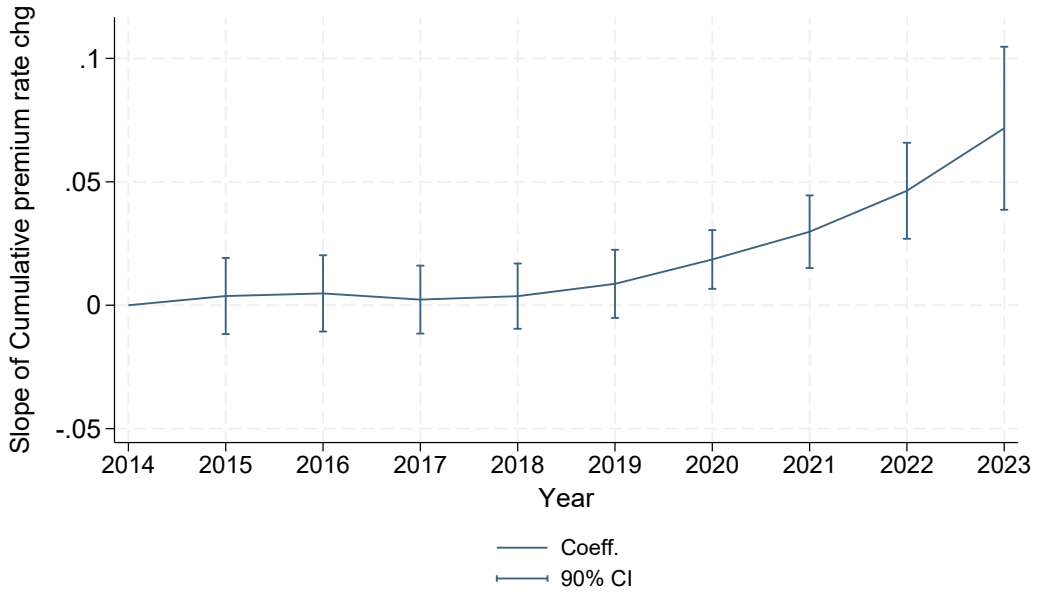
Weill, Joakim, "Perilous flood risk assessments," *Available at SSRN 4143914*, 2022.

Wong, Francis, "Taxing Homeowners Who Won't Borrow," Technical Report, CESifo Working Paper 2024.

Wylie, David, Natee Amornsiripanitch, John Heilbron, and Kevin Zhao, "JUE Insights: Who bears climate-related physical risk?," *Journal of Urban Economics*, 2025, *149*, 103791.

FIGURE 1

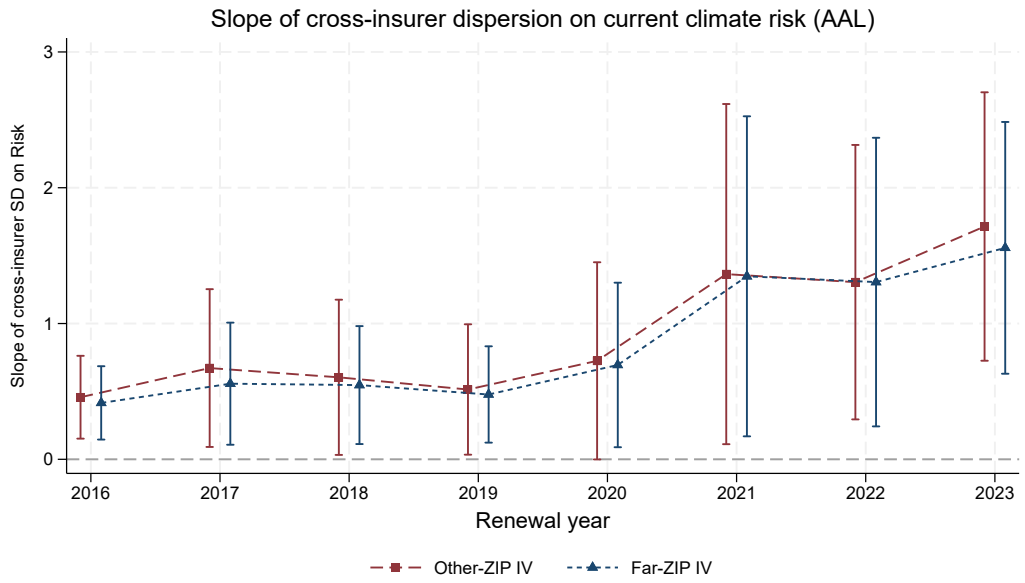
PREMIUMS HAVE INCREASED MORE IN HIGH-CLIMATE-RISK AREAS



NOTES: Sources: ICE McDash; CoreLogic Climate Risk Analytics; 2010 Decennial Census. We estimate the following regression at the ZIP-by-year level. $\Delta(Premium\ Rate)_{z,2013\ to\ t} = \gamma_z + \delta_{c,t} + \sum_{t=2015}^{2023} \beta_t ClimateRisk_z + \alpha_t Controls_z + \epsilon_{z,t}$, where z denotes ZIP code, c denotes the county, and t the year. The dependent variable is the cumulative change in the premium rate (premiums relative to the coverage limit, in percentage points), averaged at the ZIP-year level and chained from a zero anchor in 2013, across the owner-occupied policies renewing in the ZIP code that year (investor-owned properties, second homes, and policies with unrecorded occupancy are excluded). The estimation sample runs from 2014 to 2023 and the 2014 climate-risk interaction is omitted, so the plotted slopes are normalized to zero in 2014. The main independent variable is CoreLogic's average annual loss (AAL) measure of climate risk at the ZIP level, expressed in percent of the replacement cost; it is the median of the property-level AAL across the primary structures in the ZIP code.

FIGURE 2

CROSS-INSURER DISPERSION OF PREMIUM INCREASES AND CLIMATE RISK, BY RENEWAL YEAR

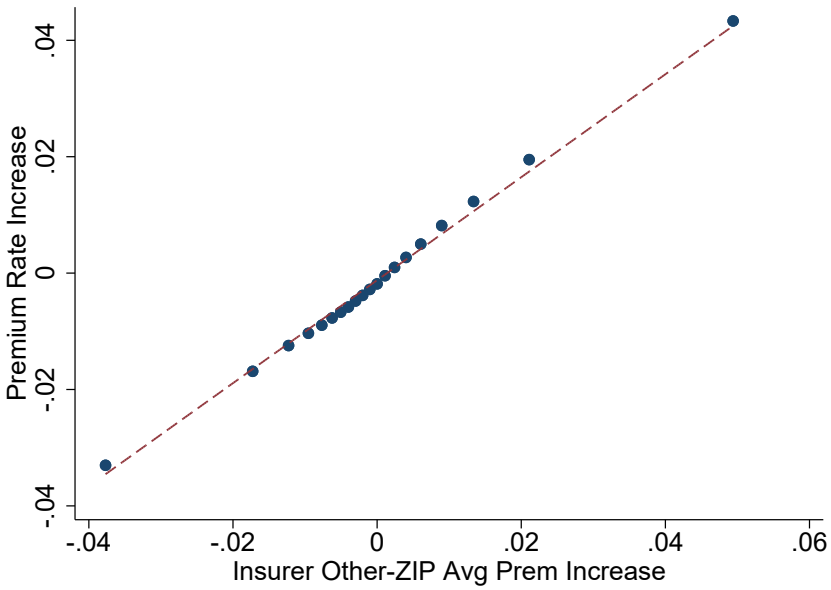


NOTES: Sources: ICE McDash; CoreLogic Climate Risk Analytics. This figure plots, by insurance renewal year, the slope from regressing a within-ZIP-by-renewal-month measure of cross-insurer dispersion in the premium-increase instrument on ZIP-level climate risk, together with its 90% confidence interval. Climate risk is CoreLogic's current average annual loss (AAL) at the ZIP level, measured as the median of property-level AAL within the ZIP code. Here AAL enters as a fraction of the coverage limit, rather than the percent-of-coverage scale used in Figure 1, so the slopes plotted in this figure are one hundred times the size they would take on that scale. Dispersion is computed in two steps within each ZIP-by-renewal-month cell: we first average the instrument to one observation per insurer, then take the standard deviation across those insurers, keeping cells with at least three insurers. The two series use the Other-ZIP and Far-ZIP instruments. All specifications include renewal-year fixed effects, and standard errors are clustered by state $\times$ year. The series span renewal years 2016–2023.

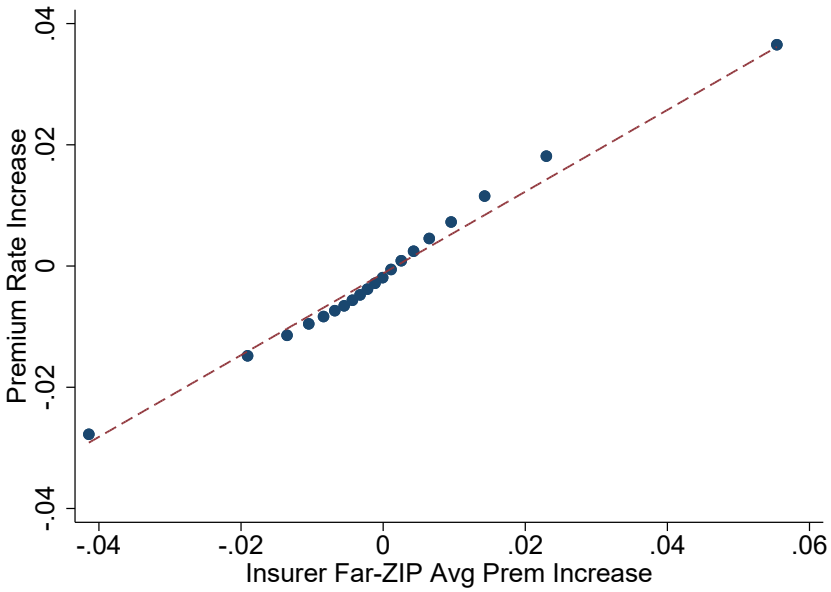
FIGURE 3

PREMIUM RATE INCREASE VS. THE INSTRUMENTS

Panel A: Other-ZIP Instrument



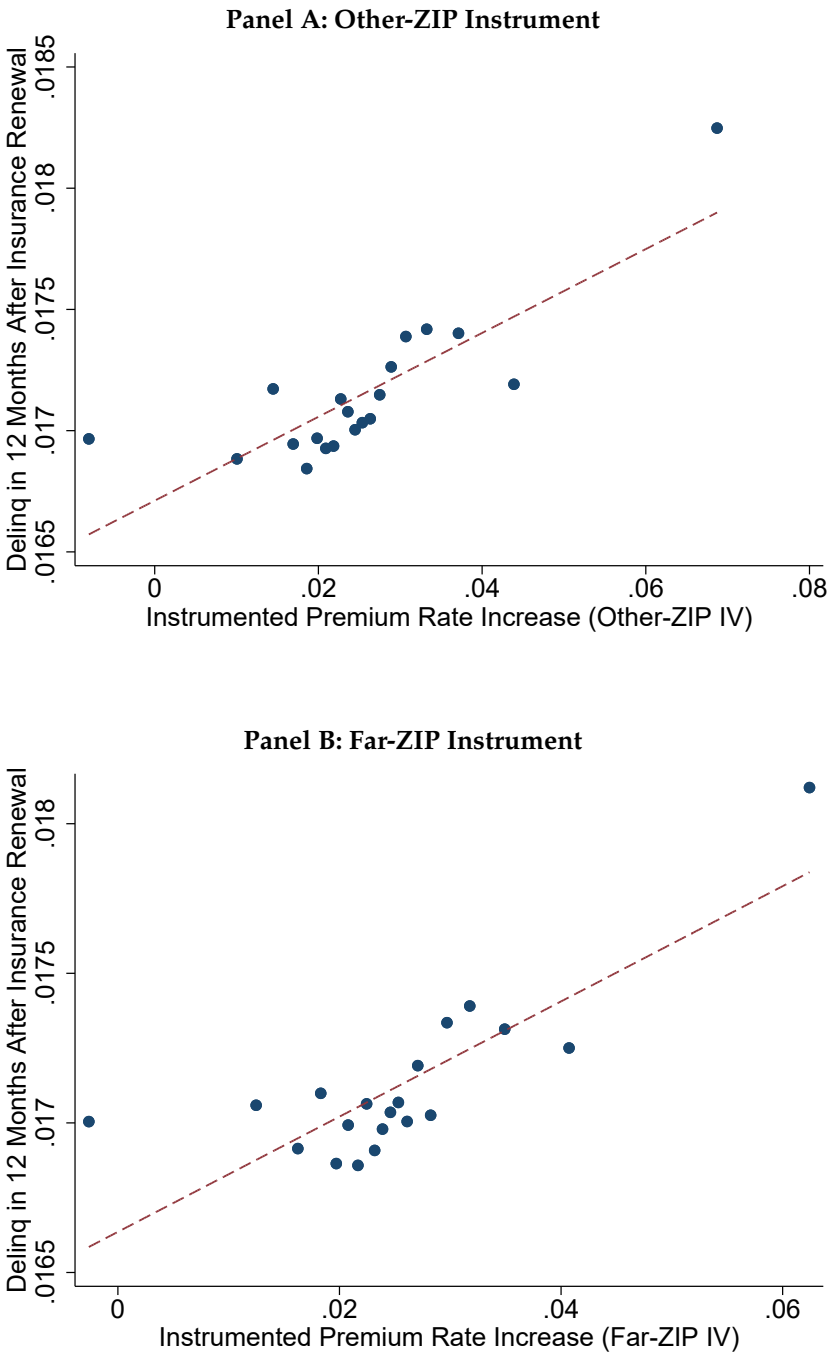
Panel B: Far-ZIP Instrument



NOTES: Sources: ICE McDash; CoreLogic Climate Risk Analytics. Bin-scatter plots of the household-level premium rate increase against each residualized leave-out instrument: the average increase among policies of the same insurer renewing in the same month in other ZIP codes within the state (Panel A), and that average restricted to ZIP codes farther away than the state's median ZIP-code-pair distance (Panel B). Both axes are residualized on the controls and fixed effects of Table 2.

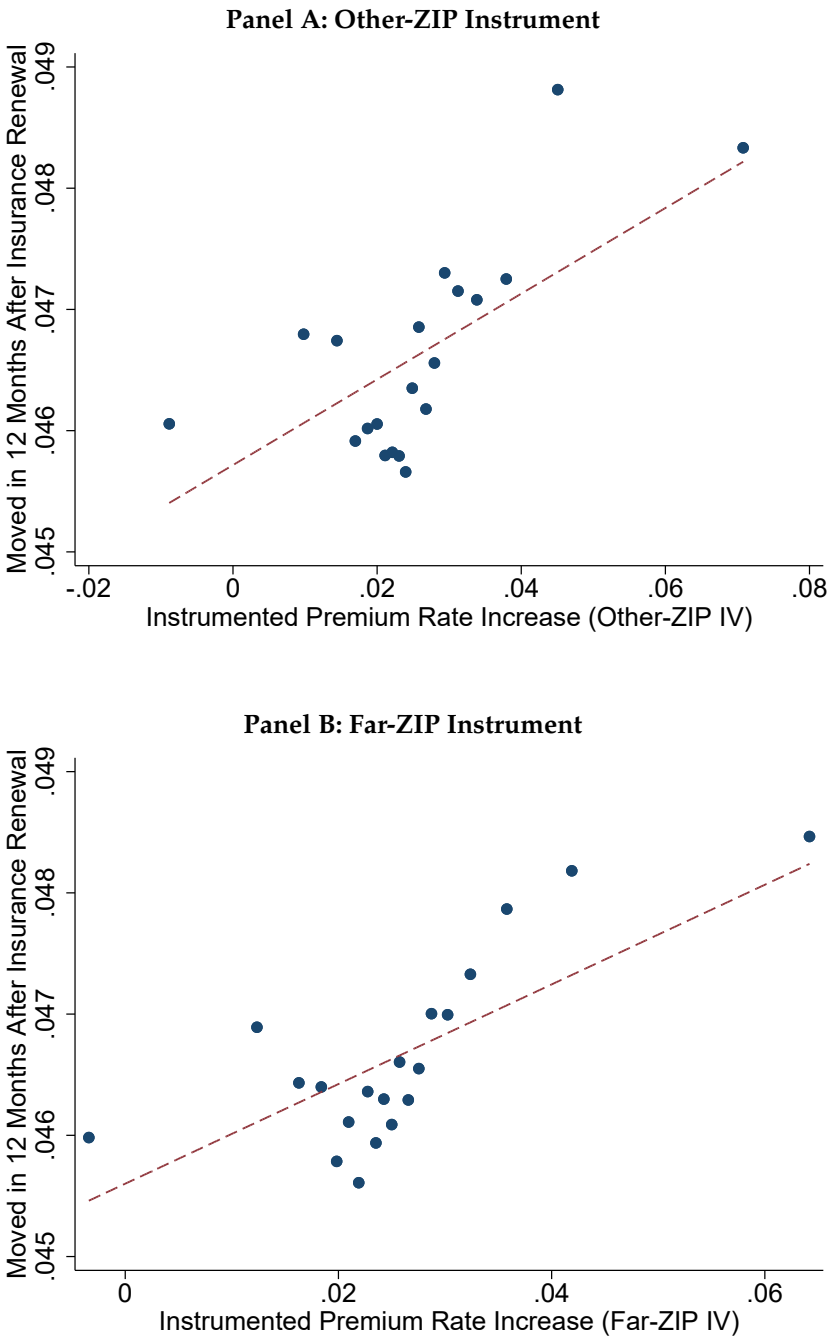
FIGURE 4

MORTGAGE DELINQUENCY AND PREMIUM RATE INCREASES



NOTES: Sources: ICE McDash; CoreLogic Climate Risk Analytics. Bin-scatter plots of 60-day mortgage delinquency against the instrumented premium rate increase, in percentage points of the coverage limit: fitted by the Other-ZIP first stage of Column (2) of Table 2 (Panel A) and by the Far-ZIP first stage of Column (4) (Panel B). Each panel uses the sample of its own column; both axes are residualized on that table's controls and fixed effects.

FIGURE 5
MOVING AND PREMIUM RATE INCREASES



NOTES: Sources: Credit Risks Insight Servicing McDash (CRISM); ICE McDash; CoreLogic deeds and owner-transfer records; CoreLogic Climate Risk Analytics. Bin-scatter plots of moving probability against the instrumented premium rate increase, in percentage points of the coverage limit: fitted by the Other-ZIP first stage of Column (2) of Table 4 (Panel A) and by the Far-ZIP first stage of Column (4) (Panel B). Each panel uses the sample of its own column and that table's moving definition; both axes are residualized on its controls and fixed effects.

TABLE 1

SUMMARY STATISTICS

	Mean	SD	25 Pctl	Median	75 Pctl
30-Day Delinquent	0.05	0.22	0.00	0.00	0.00
60-Day Delinquent	0.02	0.13	0.00	0.00	0.00
90-Day Delinquent	0.01	0.10	0.00	0.00	0.00
ZIP Had Recent Disaster	0.24	0.43	0.00	0.00	0.00
Premium Rate Change (pp×100)	0.91	9.02	-0.95	0.42	2.92
Premium Rate Increase (pp×100)	2.75	6.00	0.00	0.42	2.92
Insurer Other-ZIP Avg Prem Increase (pp×100)	2.70	3.13	0.88	1.74	3.24
Insurer Far-ZIP Avg Prem Increase (pp×100)	2.76	3.26	0.87	1.76	3.36
Premium Rate (%)	0.48	0.31	0.28	0.39	0.57
CredScore Before Renewal (McDash)	757.03	71.09	720.00	777.00	808.00
Climate Risk (%)	0.15	0.12	0.08	0.11	0.18
Climate Risk (std)	-0.00	1.00	-0.60	-0.35	0.23
DTI	33.27	11.13	25.00	34.00	42.00
CredScore (Origination)	738.81	66.99	702.00	753.00	787.00
LTV	73.45	19.54	61.74	77.35	89.34
Premium $ Changes	98.04	482.29	12.00	60.00	156.00
Annual Premium $	1,477.13	1,306.77	816.00	1,176.00	1,728.00
Coverage Limit	382,925.08	942,759.74	207,100.00	289,000.00	414,800.00
Deductible/Coverage Limit (%)	0.51	0.42	0.24	0.39	0.64
Insurer Premium Share (%)	0.46	1.59	0.01	0.05	0.22
Insurer HHI	592.65	26.48	576.75	590.36	611.19
Insurer Premium Share (State, %)	1.48	3.42	0.01	0.21	1.07
Insurer HHI (State)	904.59	263.93	756.19	879.09	1,024.93
CredScore Before Renewal (CRISM)	768.15	67.91	738.00	791.00	820.00
Orig Mtg Rate - Current Avg	-0.22	1.64	-0.82	0.05	0.78
Moved	0.05	0.21	0.00	0.00	0.00
Δ Credit Utilization	11.13	16.31	1.21	6.17	15.81
Δ Premium Rate to New ZIP Avg	0.02	0.20	-0.05	0.04	0.12

NOTES: Sources: ICE McDash; First Street Foundation; CoreLogic Climate Risk Analytics; Credit Risks Insight Servicing McDash (CRISM); FRED. Statistics use the delinquency sample of Table 2, except for variables available only in matched CRISM data, which use the relocation sample of Table 4.

TABLE 2
EFFECT OF PREMIUM INCREASES ON MORTGAGE DELINQUENCIES

	(1) 60-Day Delinquent	(2) Premium Rate Increase	(3) 60-Day Delinquent	(4) Premium Rate Increase	(5) 60-Day Delinquent
Premium Rate Increase	0.035***		0.017***		0.019***
	(14.01)		(3.84)		(3.86)
Insurer Other-ZIP Avg Prem Increase (residual)		0.885***			
		(88.24)			
Insurer Far-ZIP Avg Prem Increase (residual)				0.674***	
				(38.68)	
Lag Premium Rate	0.003***	0.022***	0.003***	0.022***	0.003***
	(4.55)	(10.85)	(3.77)	(10.61)	(3.78)
CredScore/1k before Renewal	-0.371***	-0.041***	-0.362***	-0.041***	-0.362***
	(-40.20)	(-21.66)	(-34.91)	(-21.79)	(-35.61)
Climate Risk (std)/1k	0.109	0.429***	0.090	0.461***	0.077
	(1.30)	(3.42)	(0.89)	(3.52)	(0.77)
OLS/2SLS Stage	OLS	1st Stage	2nd Stage	1st Stage	2nd Stage
Sample	All	All	All	All	All
ZIP-by-Renew Month FE	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y
Y Mean	0.017	0.025	0.017	0.025	0.017
X Mean	0.028	-0.001	0.025	-0.001	0.025
X SD	0.060	0.023	0.055	0.024	0.055
N	54,154,105	35,904,523	35,904,523	35,619,020	35,619,020
Kleibergen-Paap Wald F stat			7,787		1,496

NOTES: Sources: ICE McDash; CoreLogic Climate Risk Analytics. The dependent variable is an indicator of whether the household is delinquent on the mortgage for at least 60 days in Columns (1), (3) and (5), and *Premium Rate Increase* in (2) and (4). *Premium Rate Increase* is set to zero when the premium rate decreases; the OLS/2SLS Stage row reports which stage each column estimates. Columns (2) and (3) use the Other-ZIP instrument and Columns (4) and (5) the Far-ZIP instrument, both defined in Section 6. The instrumented columns start in renewal year 2016, the OLS column in 2014. We control for the borrower's credit risk score from ICE McDash in the month prior to the insurance renewal, the premium rate before the renewal, and the standardized climate risk of the households' residence, the property-level average annual loss standardized and divided by 1,000. We also control for the following fixed effects (FE): 5-digit ZIP by insurance policy renewal month FE, climate risk bin (tercile) by insurance policy renewal month FE, mortgage origination year FE, LTV bin FE, DTI bin FE, and credit score bin FE. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE 3
EFFECTS OF PREMIUM INCREASES ON OTHER MORTGAGE AND CREDIT OUTCOMES

	(1) 30-Day Delinquent	(2) 90-Day Delinquent	(3) Mortgage Foreclosure (2yr)	(4) Δ Credit Utilization (%)	(5) Δ Credit Card Balance ($)
Premium Rate Increase	0.035***	0.015***	0.0027***	0.977***	
	(4.80)	(3.80)	(2.89)	(3.06)	
Premium Chg ($)					0.076**
					(2.12)
Lag Premium Rate	0.014***	0.001*	0.002***	0.156***	299.063***
	(10.48)	(1.78)	(10.96)	(3.39)	(13.05)
CredScore/1k before Renewal	-0.868***	-0.220***	-0.064***	7.738***	11888.745***
	(-63.59)	(-25.36)	(-24.31)	(21.27)	(72.75)
Climate Risk (std)/1k	0.233	0.107	-0.029	27.554***	-4303.620
	(1.40)	(1.26)	(-1.23)	(4.00)	(-1.35)
OLS/2SLS Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage
Sample	All	All	All	All	All
ZIP-by-Renew Month FE	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y
Y Mean	0.050	0.011	0.0024	10.964	255.744
X Mean	0.025	0.025	0.025	0.025	99.245
X SD	0.055	0.055	0.055	0.054	368.687
N	35,619,020	35,619,020	35,619,020	30,337,580	29,362,656
Kleibergen-Paap Wald F stat	1,496	1,496	1,496	1,464	498

NOTES: Sources: ICE McDash; Credit Risks Insight Servicing McDash (CRISM); CoreLogic Climate Risk Analytics. All columns report the second stage of instrumental variable regressions. The dependent variables are: (1) an indicator for the mortgage being at least 30 days delinquent in the 12 months after insurance renewal; (2) an indicator for it being at least 90 days delinquent over the same window; (3) an indicator for the start of a mortgage foreclosure in the 24 months after insurance renewal; (4) the change in credit card utilization, measured as peak utilization percentage in 12 months post-renewal minus pre-renewal utilization, where utilization is the balance on the borrower's credit cards as a percentage of their credit limit and values above 110 are set to missing; (5) the change in credit card balance, in dollars, over the year from the month before renewal to the twelfth month of the policy year. The independent variable is *Premium Rate Increase* in Columns (1)-(4) and *Premium Chg ($)*, i.e. the change in insurance premium at policy renewal relative to the previous policy year, in Column (5). In every column the independent variable is instrumented with the Far-ZIP instrument, as in Column (5) of Table 2 and Table 4. Columns (1)–(3) use the ICE McDash sample; Columns (4) and (5) use the smaller matched CRISM sample, which is why the sample sizes differ. Controls and fixed effects are as in Table 2, with the borrower's credit score taken from ICE McDash in Columns (1)–(3) and from CRISM in Columns (4) and (5). Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE 4
EFFECTS OF PREMIUM INCREASES ON MOVING PROBABILITIES

	(1)	(2)	(3)	(4)	(5)
	Moved	Prem Rate Increase	Moved	Prem Rate Increase	Moved
Premium Rate Increase	0.035***		0.035***		0.041***
	(13.95)		(5.22)		(5.78)
Insurer Other-ZIP Avg Prem Increase (residual)		0.913*** (48.96)			
Insurer Far-ZIP Avg Prem Increase (residual)				0.697*** (32.64)	
Lag Premium Rate	-0.006***	0.024***	-0.006***	0.025***	-0.006***
	(-8.66)	(10.40)	(-6.50)	(10.40)	(-6.82)
CredScore/1k before Renewal	-0.014***	-0.039***	-0.016***	-0.039***	-0.015***
	(-5.17)	(-21.80)	(-4.81)	(-21.90)	(-4.83)
Climate Risk (std)/1k	-1.072***	0.278***	-1.036***	0.306***	-1.032***
	(-3.62)	(2.97)	(-2.78)	(3.17)	(-2.81)
OLS/2SLS Stage	OLS	1st Stage	2nd Stage	1st Stage	2nd Stage
Sample	Sep-Jan	Sep-Jan	Sep-Jan	Sep-Jan	Sep-Jan
ZIP-by-Renew Month FE	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y
Y Mean	0.048	0.026	0.047	0.026	0.047
X Mean	0.027	-0.001	0.026	-0.001	0.026
X SD	0.059	0.023	0.055	0.025	0.055
N	16,352,082	10,475,160	10,475,160	10,385,294	10,385,294
Kleibergen-Paap Wald F stat			2,397		1,066

NOTES: Sources: Credit Risks Insight Servicing McDash (CRISM); ICE McDash; CoreLogic deeds and owner-transfer records; CoreLogic Climate Risk Analytics. The dependent variable is an indicator for whether the household relocates in Columns (1), (3) and (5), defined as a mortgage termination within twelve months of renewal accompanied by a change in the borrower's credit-record address, a replacement purchase loan, or a matched sale of the insured property, and *premium rate increase* in (2) and (4). Column (1) presents OLS estimates. Columns (2) and (4) present the first stage of instrumental variable regressions, while Columns (3) and (5) present the corresponding second stage. Columns (2) and (3) use the Other-ZIP instrument and Columns (4) and (5) the Far-ZIP instrument, as in Table 2. Controls and fixed effects are as in that table, with the borrower's credit score taken from CRISM. We restrict the sample to borrowers whose home insurance renews between September and January. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE 5

HETEROGENEOUS EFFECTS OF PREMIUM INCREASES ON DELINQUENCY AND RELOCATION

	(1) Delinq.	(2) Moved	(3) Delinq.	(4) Moved	(5) Delinq.	(6) Moved
Premium Rate Increase	0.085***	0.021***	0.012***	0.040***	0.015***	0.042***
	(12.03)	(8.68)	(11.61)	(15.37)	(8.63)	(14.26)
Premium Rate Increase × Mid CredScore	-0.060***	0.016***				
	(-10.83)	(5.06)				
Premium Rate Increase × High CredScore	-0.080***	0.028***				
	(-9.24)	(8.20)				
Premium Rate Increase × Mid LTV			0.010***	-0.000		
			(6.09)	(-0.09)		
Premium Rate Increase × High LTV			0.053***	-0.012***		
			(10.60)	(-2.93)		
Premium Rate Increase × Mid DTI					0.025***	-0.008***
					(6.92)	(-3.43)
Premium Rate Increase × High DTI					0.043***	-0.014***
					(6.04)	(-4.24)
Controls	Y	Y	Y	Y	Y	Y
ZIP-by-Renew Month FE	Y	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y	Y
Sample	All	Sep–Jan	All	Sep–Jan	All	Sep–Jan
Y Mean	0.017	0.048	0.017	0.048	0.017	0.048
N	54,154,105	16,352,082	54,154,105	16,352,082	54,154,105	16,352,082

NOTES: Sources: ICE McDash; Credit Risks Insight Servicing McDash (CRISM); CoreLogic deeds and owner-transfer records; CoreLogic Climate Risk Analytics. Households are sorted into three bins of their credit score (Columns 1–2), loan-to-value (LTV) ratio (Columns 3–4), and debt-to-income (DTI) ratio (Columns 5–6); within each pair, the odd-numbered column reports an indicator for 60-day mortgage delinquency and the even-numbered column reports an indicator for moving. In each column we interact *Premium Rate Increase* with indicators for the second and third bins of the sort variable; their main effects are absorbed by the corresponding bin fixed effects and are not separately reported. The bin cutoffs are fixed values applied to each household's origination characteristic: credit score at 680 and 780 (below 680; 680 to below 780; 780 and above), LTV at 75 and 85 (below 75; 75 to below 85; 85 and above), and DTI at 30 and 45 (below 30; 30 to below 45; 45 and above); the lowest group is the omitted base, and the "Mid" and "High" interactions correspond to the second and third bins, respectively. We place the cutoffs at these values so that each bin coincides with the boundaries of the corresponding fixed-effect bins (20-point credit score bins and 5-point LTV and DTI bins) and no bin straddles more than one fixed-effect bin. These bins are therefore not exact terciles; the results are robust to instead using strict, equal-sized terciles. The delinquency columns use the full sample and the moving columns restrict to borrowers whose insurance renews between September and January. All columns report OLS estimates. Controls and fixed effects are as in Table 2, with the borrower's credit score taken from ICE McDash in the delinquency columns and from CRISM in the moving columns. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE 6
EFFECT OF PREMIUM INCREASES ON MOVING DEPENDS ON MORTGAGE RATE LOCK-IN

	(1) Moved	(2) Moved	(3) Moved
Premium Rate Increase	0.037***	0.041***	0.050***
	(14.44)	(4.41)	(4.98)
Premium Rate Increase × (Orig Mtg Rate - Current Avg)/100	0.316***	0.450	0.647**
	(3.61)	(1.53)	(2.00)
Instrument: Insurer Avg Prem Increase		Other-ZIP	Far-ZIP
OLS/2SLS Stage	OLS	2nd Stage	2nd Stage
Sample	Sep–Jan	Sep–Jan	Sep–Jan
ZIP-by-Renew Month FE	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y
Y Mean	0.048	0.047	0.047
X Mean	0.027	0.026	0.026
X SD	0.059	0.055	0.055
N	16,351,915	10,475,035	10,385,169
Kleibergen-Paap Wald F stat		1,361	340

NOTES: Sources: Credit Risks Insight Servicing McDash (CRISM); ICE McDash; CoreLogic deeds and owner-transfer records; FRED; CoreLogic Climate Risk Analytics. The independent variables are the change in premium as a percentage of the coverage limit and an additional interaction term between *Premium Rate Increase* and the difference between the interest rate the household is paying on its mortgage and the average 30-year mortgage rate in the month of insurance renewal, divided by 100. Column (1) presents OLS estimates. Columns (2) and (3) present the second stage of instrumental variable regressions, in which both *Premium Rate Increase* and its interaction with the rate gap are instrumented, using the instrument and its interaction with the rate gap. Column (2) uses the Other-ZIP instrument and Column (3) the Far-ZIP instrument, as in Table 2. Controls and fixed effects are as in that table, with the borrower's credit score taken from CRISM. We restrict the sample to borrowers whose home insurance renews between September and January. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE 7

THE EFFECT OF INSURANCE PREMIUM INCREASES ON MOVERS' PREMIUM CHANGES AND OTHER MOVER OUTCOMES

	(1) Δ Premium Rate (subsequent policy)	(2) Δ Climate Risk (pp)	(3) Δ Prem Rate to New ZIP Avg	(4) Δ to Ins-Other-ZIP Avg	(5) Δ Insurer HHI/1k
Premium Rate Increase	0.280***	-0.125***	-1.224***	0.009	0.098
	(4.22)	(-2.76)	(-21.51)	(0.10)	(1.42)
Premium Rate Increase × Moved	-1.184***				
	(-9.40)				
Moved	-0.012***				
	(-3.02)				
Lag Premium Rate	-0.050***	0.037***	-0.687***	-0.028***	0.005
	(-15.88)	(5.53)	(-78.28)	(-9.14)	(0.66)
CredScore/1k before Renewal	-0.051***	0.014***	-0.053***	-0.017**	-0.080***
	(-13.95)	(3.18)	(-7.67)	(-2.06)	(-5.56)
Climate Risk (std)/1k	0.689***	-139.131***	15.485***	-0.310	-11.155***
	(4.04)	(-36.92)	(3.75)	(-0.47)	(-4.18)
OLS/2SLS Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage
Sample	Sep-Jan	Sep-Jan	Sep-Jan	Sep-Jan	Sep-Jan
Subsample	Movers,Stayers	Movers	Movers	Movers	Movers
Old ZIP-by-Renew Yr FE	Y				
Old State-by-Renew Yr FE		Y	Y	Y	Y
Climate Risk Bin-by-Renew Yr FE	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y
Y Mean	0.010	-0.005	0.023	0.004	0.013
X Mean	0.023	0.028	0.028	0.025	0.028
X SD	0.050	0.058	0.058	0.051	0.058
N	7,799,676	442,098	467,111	99,370	468,754
Kleibergen-Paap Wald F stat	266	592	596	441	600

NOTES: Sources: Credit Risks Insight Servicing McDash (CRISM); ICE McDash; CoreLogic deeds and owner-transfer records; CoreLogic Climate Risk Analytics. Every column reports the second stage, instrumenting *Premium Rate Increase* with the Far-ZIP instrument. Column (1) reports the change in premium rate from policy year t to $t+1$ on the full September-to-January sample of movers and stayers, interacting *Premium Rate Increase* with an indicator for moving within 12 months after renewal. Columns (2)–(5) restrict the sample to movers. The dependent variable is the change in climate risk, destination minus origin expected annual loss, property-level where the CoreLogic match succeeds and ZIP-median otherwise, in percentage points of replacement cost, in Column (2); the change from the mover's pre-move premium rate to the average premium rate in the new ZIP code in Column (3), the ICE McDash mean there in the year following renewal, computed only for ZIP-years with more than ten policies; the change in the mover's insurer's average premium rate in other ZIP codes within the same state, the destination state for across-state movers, in Column (4); and the change in the insurer Herfindahl-Hirschman index from the old to the new ZIP in Column (5), built from insurer shares of annual premium dollars written in the ZIP code. Controls are as in Table 2, with the borrower's credit score taken from CRISM and climate risk measured at the old residence. Fixed effects are as in Table 2 but by renewal year rather than renewal month, with 5-digit old ZIP by renewal year in Column (1) and old state by renewal year in Columns (2)–(5). Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE 8
THE EFFECT OF INSURANCE PREMIUM INCREASES ON INVESTOR PURCHASE

	(1)	(2)	(3)	(4)	(5)
		Prem Rate		Prem Rate	
	Investor	Increase	Investor	Increase	Investor
Premium Rate Increase	0.027***		0.052**		0.045*
	(7.20)		(1.97)		(1.65)
Insurer Other-ZIP Avg Prem Increase (residual)		0.885***			
		(41.32)			
Insurer Far-ZIP Avg Prem Increase (residual)				0.708***	
				(35.80)	
Lag Premium Rate	0.025***	0.029***	0.024***	0.029***	0.025***
	(15.38)	(11.50)	(10.18)	(11.20)	(10.22)
CredScore/1k before Renewal	-0.164***	-0.072***	-0.145***	-0.073***	-0.146***
	(-24.46)	(-23.91)	(-20.08)	(-23.79)	(-20.03)
Climate Risk (std)/1k	0.779*	0.463**	0.607	0.529***	0.542
	(1.87)	(2.49)	(1.29)	(2.69)	(1.15)
OLS/2SLS Stage	OLS	1st Stage	2nd Stage	1st Stage	2nd Stage
Sample	Transactions	Transactions	Transactions	Transactions	Transactions
5-Digit ZIP FE	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Yr FE	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y
Y Mean	0.041	0.033	0.041	0.033	0.041
X Mean	0.035	0.001	0.033	0.001	0.033
X SD	0.072	0.025	0.067	0.027	0.067
N	1,365,925	925,426	925,426	918,386	918,386
Kleibergen-Paap Wald F stat			1,707		1,282

NOTES: Sources: CoreLogic deeds and owner-transfer records; ICE McDash; CoreLogic Climate Risk Analytics. The dependent variable is an indicator that equals one when the buyer is classified as a market investor *and* the purchase is financed outside the standard mortgage channel, meaning either an all-cash purchase or a mortgage from a private-party lender; it is zero for all other observed purchases. We classify a buyer as a market investor by applying a name-pattern classification to the grantee name in CoreLogic's owner-transfer records: a buyer is institutional if its name identifies a business entity, and is a market investor if it is institutional and is not a government agency, nonprofit, trust or estate, religious organization, or homeowner association. Requiring the financing condition in addition restricts the indicator to a more sophisticated subset of investor purchases. It also excludes owner-occupant households that purchase through a limited liability company or a similar entity and finance the purchase with a standard mortgage, which the name-pattern classification alone would count as investors. The sample is the set of home purchases matched to a vacated property in the buyer analysis sample, so the unit of observation is a transaction. Column (1) presents OLS estimates. Columns (2) and (3) present the first and second stages of the instrumental variable regression using the residualized Other-ZIP instrument, and Columns (4) and (5) repeat the first and second stages using the residualized Far-ZIP instrument. Controls are as in Table 2, measured for the previous owner and at the old residence. The fixed effects differ: 5-digit ZIP FE (not interacted with renewal month), climate risk bin (tercile) by insurance policy renewal year FE, mortgage origination year FE, LTV bin FE, DTI bin FE, and credit score bin FE. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE 9
THE EFFECT OF INVESTORS ON INSURANCE OUTCOMES

	(1) Log Premium	(2) Premium Rate	(3) Premium Rate Increase	(4) Log Cov	(5) Deductible ×100/Cov
Investor	-0.148***	-0.040***	-0.006***	-0.072***	0.108***
	(-30.56)	(-15.57)	(-12.36)	(-29.35)	(24.58)
Investor × Block-Yr Avg Prem Increase (std)	-0.032***	-0.026***	-0.005***	-0.013***	0.010**
	(-5.71)	(-9.40)	(-4.80)	(-4.00)	(2.01)
Climate Risk (std)/1k	0.908***	3.196***	0.204***	-1.545***	2.808***
	(4.57)	(4.78)	(4.70)	(-4.69)	(10.79)
CredScore before Renewal/1k	-0.497***	-0.348***	-0.055***	0.085***	-0.134***
	(-28.90)	(-24.92)	(-21.67)	(29.64)	(-24.71)
House Age ≤ 5	-0.194***	-0.072***	-0.007***	-0.007***	-0.017***
	(-60.25)	(-19.69)	(-9.40)	(-3.16)	(-8.90)
Log(Sqft)	0.450***	-0.064***	-0.002***	0.596***	-0.139***
	(24.01)	(-17.09)	(-8.05)	(25.75)	(-20.89)
Log(Assessed Value$)	0.189***	-0.041***	-0.001	0.253***	-0.034***
	(8.46)	(-5.57)	(-1.38)	(7.87)	(-2.79)
House Age Bin FE	Y	Y	Y	Y	Y
Block Group-by-Renew Yr FE	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Yr FE	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y
Y Mean	7.180	0.490	0.032	12.678	0.506
Investor Mean	0.055	0.055	0.055	0.055	0.055
N	42,096,367	42,096,367	42,008,487	42,386,402	40,396,342

NOTES: Sources: ICE McDash; CoreLogic property records; CoreLogic Climate Risk Analytics. The sample consists of annual homeowners' policies in ICE McDash matched to CoreLogic property records. The unit of observation is a policy renewal. The main independent variables are an indicator for investor-owned properties, as designated by McDash (a different concept from the buyer classification in Table 8, which comes from CoreLogic buyer names), and its interaction with the standardized Census block-group-by-renewal-year average of premium rate increases relative to the prior year. The dependent variables are: log total insurance premiums (Column 1), the premium rate (Column 2), the premium rate increase from the prior renewal (Column 3), log coverage limit (Column 4), and the deductible as a share of the coverage limit multiplied by 100 (Column 5). Additional controls include climate risk, credit score from ICE McDash prior to renewal (divided by 1,000), an indicator for properties under five years old, log square footage, and log assessed property value. All specifications include fixed effects for house age (10-year) bins, Census block group-by-renewal year, climate risk bin-by-renewal year, and origination year, credit score bin, LTV bin, and DTI bin. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

Appendix

TABLE A.1
HOMEOWNERS' INSURANCE PREMIUMS AND PREMIUM-RATE CHANGES BY STATE

	Annual Premium (Dollars), 2023		Premium Rate Change (pp ×100), 2023		Cumulative Premium-Rate % Change, 2013–2023
State	Mean	SD	Mean	SD	
United States	1,974	1,314	3.3	12.2	26.0
Florida	3,413	1,923	8.5	23.0	43.3
Louisiana	3,222	1,813	11.7	23.6	27.5
Colorado	2,942	1,464	7.0	13.8	65.5
Oklahoma	2,755	1,375	6.4	18.9	29.8
Texas	2,713	1,482	8.2	18.0	40.2
Nebraska	2,557	1,114	6.0	14.1	52.3
Minnesota	2,419	1,144	4.3	10.7	21.8
Kansas	2,387	1,095	2.1	13.2	24.7
Rhode Island	2,329	1,196	2.7	9.0	31.9
Mississippi	2,315	1,180	1.6	16.3	21.0
Connecticut	2,303	1,411	1.5	8.8	28.1
South Dakota	2,233	1,118	5.3	13.2	41.5
Massachusetts	2,181	1,206	2.0	6.8	23.3
Alabama	2,091	1,074	2.8	13.3	16.5
Missouri	2,075	1,040	2.5	12.3	20.4
Montana	2,068	1,117	1.9	11.7	36.6
Wyoming	2,061	1,224	4.6	14.8	39.8
North Dakota	2,027	942	2.1	8.9	6.6
Arkansas	2,018	1,012	3.7	13.2	21.2
Georgia	1,995	1,064	1.4	9.9	31.0
South Carolina	1,928	1,280	2.0	11.1	17.0
New York	1,923	1,389	1.5	6.5	9.8
Illinois	1,871	1,005	3.6	10.1	30.8
Tennessee	1,849	1,046	2.3	11.2	15.1
District of Columbia	1,805	1,237	1.1	6.6	5.8
Maryland	1,791	1,023	2.7	7.8	33.9
New Jersey	1,786	1,065	0.9	6.2	12.9
Iowa	1,762	800	4.6	10.6	26.6
Kentucky	1,732	901	3.1	10.8	12.9
California	1,701	1,207	1.0	7.8	14.5
New Mexico	1,662	1,028	4.2	12.1	34.7
Virginia	1,626	998	2.4	8.8	24.5

Continued on next page

Table A.1 (continued)

State	Annual Premium (Dollars), 2023		Premium Rate Change (pp ×100), 2023		Cumulative Premium-Rate % Change, 2013–2023
	Mean	SD	Mean	SD	
Indiana	1,617	832	2.7	9.9	13.1
North Carolina	1,616	994	2.6	10.3	31.2
Washington	1,467	846	1.9	8.0	9.7
New Hampshire	1,460	818	0.9	7.4	10.3
Arizona	1,457	938	4.9	11.3	37.8
Michigan	1,421	823	1.2	9.2	5.8
Vermont	1,410	912	0.6	7.7	8.6
Ohio	1,404	745	2.5	9.2	14.6
Pennsylvania	1,401	871	1.3	7.0	12.8
Idaho	1,359	870	2.3	8.1	27.4
Wisconsin	1,346	728	1.8	7.5	10.1
Delaware	1,328	850	1.1	6.8	25.7
West Virginia	1,309	750	0.9	8.5	11.5
Utah	1,305	828	3.2	7.9	31.1
Oregon	1,302	832	2.7	7.6	23.5
Maine	1,301	825	0.9	8.2	13.5
Nevada	1,158	728	0.8	6.5	20.0

NOTES: Source: ICE McDash. This table uses the full panel of annual McDash policies rather than the regression estimation sample, so the sample restrictions described in Table 2 are not applied here. The table covers the continental United States; Alaska and Hawaii are excluded and the District of Columbia is retained. The first row, labeled United States, pools the entire continental sample and is not an average of the state rows beneath it; the state rows follow in descending order of the mean annual premium. *Annual Premium* is twelve times the monthly premium, in dollars, reported as a cross-section of 2023 renewals. *Premium Rate Change* is the year-over-year change in the premium rate, where the premium rate is the annual premium expressed as a percentage of the coverage limit. It is reported for the 2023 cross-section and scaled by 100, the same scale as the corresponding row of Table 1. This is the uncensored change and can therefore be negative; it is not the increase-only variable used in the regressions, which is set to zero whenever the premium rate falls. The *Cumulative Premium-Rate % Change* column follows the construction used for Figure 1, not an endpoint comparison of levels: within each state we average the within-loan, same-coverage premium-rate change across the loans renewing in each year from 2014 through 2023 and chain those annual averages from a zero anchor in 2013. Because each link is a year-over-year difference, the chain telescopes to the 2013–2023 change while requiring only that a loan be observed in consecutive years rather than across the whole window. Same coverage means the deductible and the replacement-cost provision are both unchanged at renewal, so the measure reflects movements in price rather than changes in what is being purchased. The column divides that chained total by the state's mean premium rate in 2013, expressing the cumulative change relative to the 2013 level. Because base rates differ across states, this ordering is not the ordering of absolute increases: Colorado's absolute increase is smaller than Florida's, but measured against its much lower 2013 base it is the larger increase, 65.5% against 43.3%.

TABLE A.2
MORTGAGE DELINQUENCY EFFECT ACROSS INVESTOR TYPES

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				60-Day Mortgage Delinquent			
Premium Rate Increase	0.022***	0.016***	0.011***	0.029	0.019**	0.014***	0.019***
	(3.13)	(2.91)	(2.82)	(0.75)	(2.48)	(3.20)	(3.54)
Lag Premium Rate	0.006***	0.000	0.001	0.010***	0.006***	0.000	0.003***
	(4.53)	(0.54)	(1.18)	(3.83)	(6.25)	(0.54)	(4.09)
CredScore/1k before Renewal	-0.526***	-0.246***	-0.231***	-0.364***	-0.327***	-0.265***	-0.297***
	(-54.56)	(-23.95)	(-26.73)	(-13.27)	(-19.41)	(-26.64)	(-28.62)
Climate Risk (std)/1k	0.202	0.026	0.057	1.332	-0.242	0.161	0.076
	(0.87)	(0.19)	(0.57)	(1.52)	(-1.36)	(1.31)	(0.69)
OLS/2SLS Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage
Sample	GNMA	FNMA	FHLMC	Private Securitized	Portfolio	Cnvtnl w/ PMI	Cnvtnl w/o PMI
ZIP-by-Renew Month FE	Y	Y	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y	Y	Y
Y Mean	0.038	0.012	0.011	0.026	0.013	0.013	0.013
X Mean	0.032	0.024	0.024	0.022	0.023	0.026	0.023
X SD	0.063	0.053	0.052	0.054	0.053	0.056	0.051
N	5,861,946	13,432,394	10,794,285	382,855	3,613,964	10,551,959	18,044,323
Kleibergen-Paap Wald F stat	1,547	1,261	1,564	411	981	1,718	1,062

NOTES: Sources: ICE McDash; CoreLogic Climate Risk Analytics. "Investor" here is the mortgage investor or guarantor recorded by ICE McDash, not a property investor. This table presents the effect of premium increases on mortgage delinquency across investor types. Every column reports the second stage, instrumenting with the Far-ZIP instrument. The dependent variable is an indicator of whether the household is delinquent on the mortgage for at least 60 days within the 12 months after insurance policy renewal. The main independent variable is the change in premium as a percentage of the coverage limit. We use Ginnie Mae loans in Column (1), Fannie Mae loans in (2), Freddie Mac loans in (3), private mortgages that are securitized in (4), and those that remain in banks' portfolios in (5). The assignment of loans to these investor types uses lagged loan information, that is, the investor type and loan type recorded at the household's previous policy renewal, so that the subsample is predetermined with respect to the current premium shock. In Columns (6) and (7), we further split conventional mortgages by whether they have private mortgage insurance or not. Controls and fixed effects are as in Table 2. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE A.3
CONTROLLING FOR INSURER-YEAR FIXED EFFECTS

	(1)	(2)	(3)	(4)	(5)	(6)
	60-Day Delinquent			Moved		
Premium Rate Increase	0.036***	0.008**	0.010**	0.034***	0.017***	0.022***
	(14.68)	(2.50)	(2.49)	(13.58)	(2.69)	(2.90)
Lag Premium Rate	0.002***	0.003***	0.003***	-0.006***	-0.006***	-0.006***
	(3.49)	(3.96)	(3.97)	(-10.16)	(-7.41)	(-7.82)
CredScore/1k before Renewal	-0.369***	-0.361***	-0.361***	-0.014***	-0.017***	-0.017***
	(-40.23)	(-34.90)	(-35.58)	(-5.54)	(-5.45)	(-5.48)
Climate Risk (std)/1k	0.095	0.088	0.078	-1.059***	-1.006***	-1.001***
	(1.15)	(0.88)	(0.78)	(-3.58)	(-2.71)	(-2.73)
OLS/2SLS Stage	OLS	2nd Stage	2nd Stage	OLS	2nd Stage	2nd Stage
Instrument		Other-ZIP	Far-ZIP		Other-ZIP	Far-ZIP
Sample	All	All	All	Sep-Jan	Sep-Jan	Sep-Jan
ZIP-by-Renew Month FE	Y	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y	Y
Insurer-Year FE	Y	Y	Y	Y	Y	Y
Y Mean	0.017	0.017	0.017	0.048	0.047	0.047
X Mean	0.028	0.025	0.025	0.027	0.026	0.026
X SD	0.060	0.055	0.055	0.059	0.055	0.055
N	54,149,927	35,904,506	35,619,003	16,349,369	10,475,146	10,385,274
Kleibergen-Paap Wald F stat		8,487	967		7,623	662

NOTES: Sources: ICE McDash; Credit Risks Insight Servicing McDash (CRISM); CoreLogic deeds and owner-transfer records; CoreLogic Climate Risk Analytics. This table adds insurer by insurance policy renewal year fixed effects to the baseline specification, where the insurer is the one that wrote the household's previous policy, which is also the insurer we use to construct the instrument. The dependent variable is an indicator of whether the household is delinquent on the mortgage for at least 60 days within the 12 months after insurance policy renewal in Columns (1)–(3), and an indicator for whether the household relocates in Columns (4)–(6). We use the full sample in Columns (1)–(3) and restrict the sample to borrowers whose home insurance renews between September and January in Columns (4)–(6). Columns (1) and (4) report OLS estimates, while the remaining columns report the second stage of two-stage least-squares regressions using the Other-ZIP instrument (Columns 2 and 5) or the Far-ZIP instrument (Columns 3 and 6). Controls and fixed effects are as in Table 2, with the borrower's credit score taken from ICE McDash in Columns (1)–(3) and from CRISM in Columns (4)–(6), plus the insurer by renewal year fixed effects. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE A.4
DELINQUENCY AND RELOCATION EFFECTS AND RECENT DISASTERS

	(1)	(2)	(3)	(4)	(5)	(6)
	60-Day Mortgage Delinquent			Moved		
Premium Rate Increase	0.024***	0.016***	0.016***	0.029***	0.047***	0.047***
	(3.01)	(3.01)	(3.03)	(3.23)	(5.09)	(5.13)
Premium Rate Increase × Disaster			0.010			-0.018
			(1.09)			(-1.46)
Lag Premium Rate	0.000	0.005***	0.003***	-0.006***	-0.006***	-0.006***
	(0.31)	(5.05)	(3.76)	(-5.32)	(-5.43)	(-6.81)
CredScore/1k before Renewal	-0.382***	-0.354***	-0.362***	-0.012***	-0.016***	-0.015***
	(-24.29)	(-31.25)	(-35.60)	(-3.26)	(-4.27)	(-4.81)
Climate Risk (std)/1k	0.064	0.076	0.075	-0.674*	-1.170**	-1.030***
	(0.28)	(0.69)	(0.74)	(-1.96)	(-2.48)	(-2.80)
OLS/2SLS Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage	2nd Stage
Sample	All	All	All	Sep-Jan	Sep-Jan	Sep-Jan
Subsample	Disaster	No Disaster	All	Disaster	No Disaster	All
ZIP-by-Renew Month FE	Y	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y	Y
Y Mean	0.020	0.016	0.017	0.048	0.046	0.047
X Mean	0.029	0.024	0.025	0.030	0.024	0.026
X SD	0.063	0.052	0.055	0.063	0.052	0.055
N	9,101,343	26,513,949	35,615,293	2,563,867	7,820,266	10,384,134
Kleibergen-Paap Wald F stat	217	2,945	108	161	2,167	78

NOTES: Sources: ICE McDash; Credit Risks Insight Servicing McDash (CRISM); CoreLogic deeds and owner-transfer records; First Street Foundation; CoreLogic Climate Risk Analytics. The dependent variable is an indicator of whether the household is delinquent on the mortgage for at least 60 days within the 12 months after insurance policy renewal in Columns (1)–(3), and an indicator for whether the household relocates in Columns (4)–(6). All columns report the second stage, instrumenting with the Far-ZIP instrument. Within each outcome, the first column uses households in ZIP codes that experienced natural disasters within the current calendar year or the preceding two years according to First Street Foundation's Climate Data (Columns 1 and 4); the second uses households in the remaining non-disaster-affected ZIP codes (Columns 2 and 5); and the third uses the full sample, adding an indicator for recently disaster-affected ZIP codes and its interaction with *Premium Rate Increase* (Columns 3 and 6). The delinquency columns use the full sample; the relocation columns restrict to borrowers whose home insurance renews between September and January. Controls and fixed effects are as in Table 2, with the borrower's credit score taken from ICE McDash in Columns (1)–(3) and from CRISM in Columns (4)–(6). Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE A.5
DELINQUENCY EFFECT WHEN HOUSEHOLDS ARE ATTENTIVE TO PREMIUM CHANGES

	(1) 60-Day Delinquent	(2) 60-Day Delinquent	(3) 60-Day Delinquent
Premium Rate Increase	0.019***	0.017***	0.020***
	(8.03)	(3.33)	(3.46)
Lag Premium Rate	0.003***	0.003***	0.003***
	(4.31)	(3.14)	(3.33)
CredScore/1k before Renewal	-0.331***	-0.317***	-0.317***
	(-31.36)	(-27.12)	(-29.91)
Climate Risk (std)/1k	0.167	0.094	0.077
	(1.21)	(0.56)	(0.51)
OLS/2SLS Stage	OLS	2nd Stage	2nd Stage
Instrument		Other-ZIP	Far-ZIP
Sample	Attentive	Attentive	Attentive
Subsample	Non-Switcher	Non-Switcher	Non-Switcher
ZIP-by-Renew Month FE	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y
Y Mean	0.015	0.015	0.015
X Mean	0.029	0.026	0.026
X SD	0.066	0.059	0.059
N	5,080,629	3,460,308	3,430,460
Kleibergen-Paap Wald F stat		1,528	733

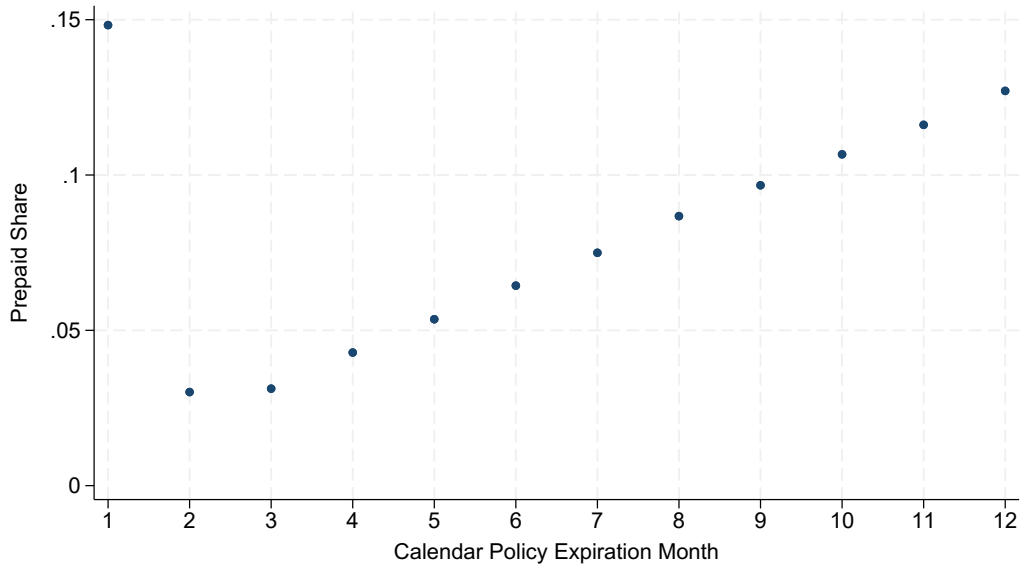
NOTES: Sources: ICE McDash; CoreLogic Climate Risk Analytics. This table presents results testing whether the effect of premium increases on mortgage delinquency exists among households that were likely attentive to premium changes but that did not switch insurers, so that the instrumented premium increase is one they continued to bear. Attentive households are those that adjusted their deductible percentages or changed between Replacement Cost Value and Actual Cash Value coverage at the renewal, relative to the previous policy year. Column (1) reports OLS estimates. Columns (2) and (3) report the second stage, using the Other-ZIP instrument (Column 2) and the Far-ZIP instrument (Column 3). Controls and fixed effects are as in Table 2. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

Internet Appendix

For Online Publication

FIGURE IA.1

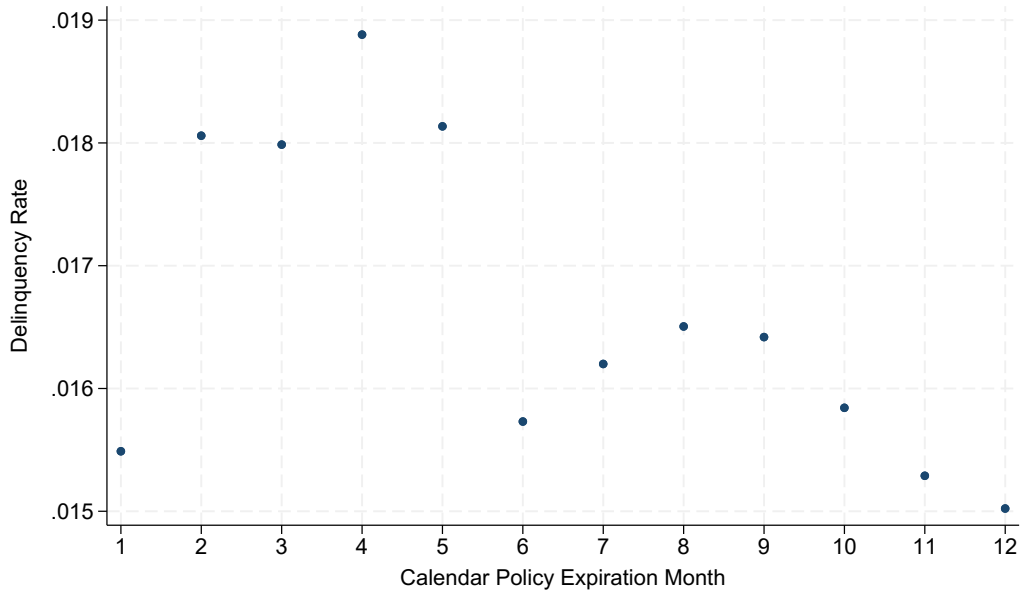
PROBABILITIES OF MORTGAGE PREPAYMENT REPORTED BY INSURANCE POLICY EXPIRATION MONTH



NOTES: Source: ICE McDash. This figure plots the share of mortgages that are prepaid within 12 months of renewal, by the calendar month of the insurance policy expiration date, on the delinquency estimation sample.

FIGURE IA.2

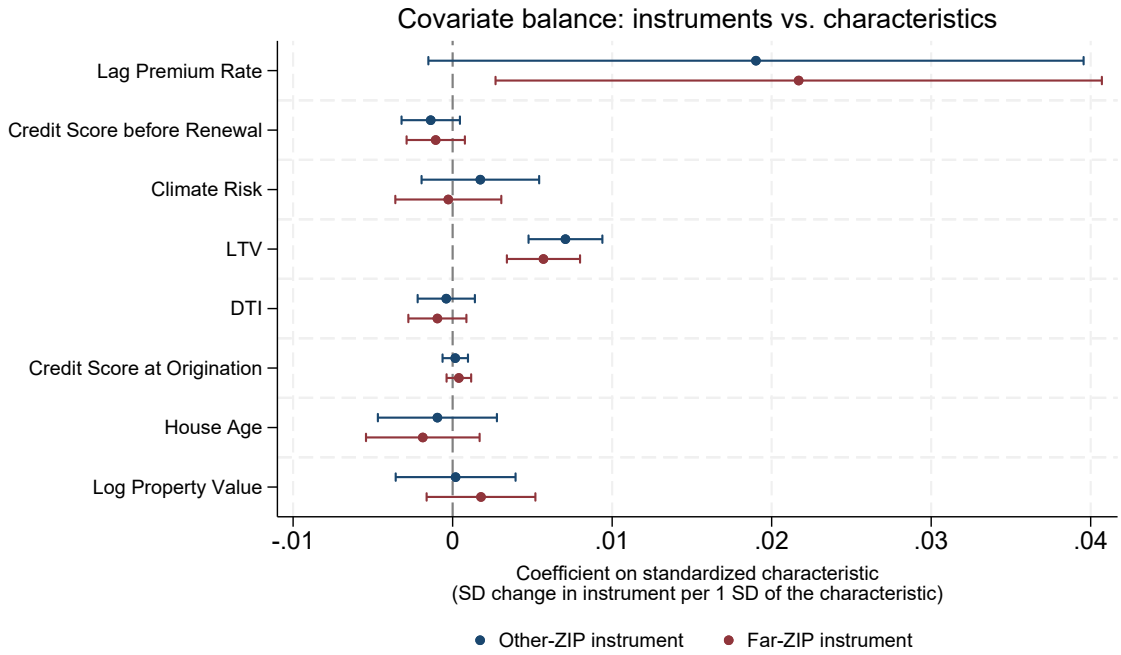
SHARE OF DELINQUENT MORTGAGES REPORTED BY INSURANCE POLICY EXPIRATION MONTH



NOTES: Source: ICE McDash. This figure plots the share of mortgages that are 60 or more days delinquent in the 12 months after renewal, by the calendar month of the insurance policy expiration date, on the delinquency estimation sample.

FIGURE IA.3

COVARIATE BALANCE: INSTRUMENTS AND PREDETERMINED CHARACTERISTICS

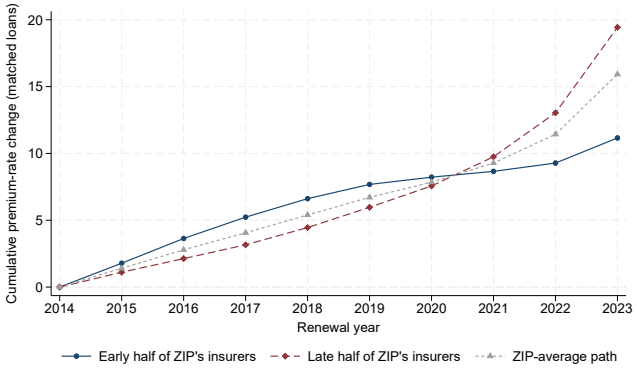


NOTES: Sources: ICE McDash; CoreLogic property records; CoreLogic Climate Risk Analytics. This figure reports covariate-balance regressions of our instruments on predetermined household, house, and loan characteristics, estimated on the full McDash sample (the mortgage-delinquency sample). Each row corresponds to one predetermined household, house, or loan characteristic. For each instrument, the marker is the coefficient from a regression of the standardized residualized leave-out instrument on all of the standardized characteristics jointly, and the horizontal cap spans its 90% confidence interval. Because every variable (the instrument and each characteristic) is standardized to mean zero and unit standard deviation within the estimation sample, each coefficient is the standard-deviation change in the instrument associated with a one-standard-deviation change in the characteristic; the dashed vertical line at zero marks perfect balance. Navy markers denote the Insurer Other-ZIP instrument and maroon markers the Insurer Far-ZIP instrument. The regressions include the fixed effects of Table 2, and standard errors are clustered at the state-by-year level. The estimated coefficients are small: the largest, on the lagged premium rate, is about two hundredths of a standard deviation, and every other characteristic is well below one hundredth. The instruments are therefore close to orthogonal to predetermined characteristics conditional on the fixed effects.

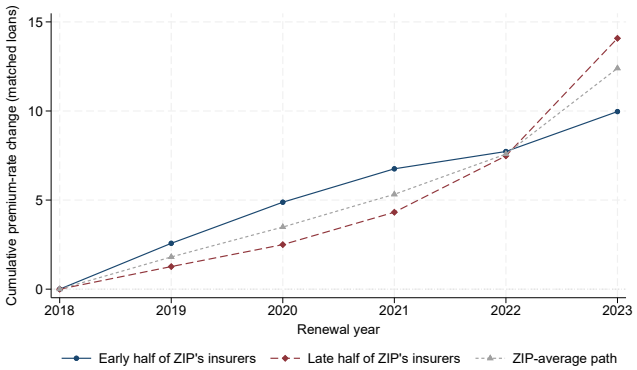
FIGURE IA.4

STAGGERED INSURER REPRICING WITHIN A ZIP CODE

Panel A. 2014–2023 Window



Panel B. 2018–2023 Window



NOTES: Source: ICE McDash. This figure plots the cumulative premium-rate change of each insurer within a ZIP code, by renewal year. Cumulative change is measured as a matched loan-level long difference: for each loan we take its premium rate as a percentage of coverage limit at its renewal in year y minus its rate at its renewal in the window's first year, keeping only loans whose insurer is unchanged between the two dates, and average these differences within insurer-ZIP-year cells. The vertical axis reports this cumulative change in percentage points of the coverage limit multiplied by 100, so a value of 10 on the axis is a cumulative increase of 0.10 percentage points. Because the cumulative change is a single level difference for a fixed set of loans rather than a sum of period-to-period changes, no chaining is involved; and no value is interpolated or carried forward for an insurer-ZIP-year with no matched loan, so the plotted series show a marker only at years with a genuine underlying observation. The sample is restricted to insurer-ZIP pairs with at least two matched loan-years per window-year, in ZIP codes with at least four such insurers, and to insurer-ZIP pairs whose cumulative change as of their last observed year is positive (net premium increases). Within each ZIP, we rank insurers by their *half-life year*, the first year by which the insurer has implemented half of its eventual cumulative increase, and split them into an early-moving and a late-moving half. Each panel plots the ZIP-year-average cumulative change together with the average for the early- and late-moving halves (each expressed as the ZIP-average path plus the within-ZIP deviation of that group), weighted by the number of matched loan-years per insurer-ZIP pair. Panel A uses a 2014–2023 window and Panel B a 2018–2023 window; the later window trades sample length for a larger set of loans that can be matched to the window's first year. The early- and late-moving paths cross between 2020 and 2021 in Panel A and in 2022 in Panel B. By 2023 the late-moving half has raised premium rates by roughly 0.19 and 0.14 percentage points of the coverage limit in the two windows, against roughly 0.11 and 0.10 percentage points for the early-moving half.

TABLE IA.1
RESIDUALIZING THE INSTRUMENTAL VARIABLE

	(1) Insurer Other-ZIP Avg Prem Rate Increase	(2) Insurer Far-ZIP Avg Prem Rate Increase
12-Month Lag Insurer Other-ZIP Avg Prem Rate Increase	0.469***	
	(9.97)	
24-Month Lag Insurer Other-ZIP Avg Prem Rate Increase	0.183***	
	(3.73)	
12-Month Lag Insurer Far-ZIP Avg Prem Rate Increase		0.441***
		(10.66)
24-Month Lag Insurer Far-ZIP Avg Prem Rate Increase		0.192***
		(4.45)
Renew Month FE	Yes	Yes
Y Mean	0.036	0.037
N	17,871,227	17,462,596
R-squared	0.407	0.377

NOTES: Source: ICE McDash. This table presents the regression used to residualize the initial values of the instrumental variable. The unit of observation is a cell of the leave-out aggregate, defined by the prior-year insurer, the ZIP code, and the renewal year-month, not a household. The number of cells is therefore not comparable with the household-renewal counts reported in the regression tables. The dependent variable is *Insurer Other-ZIP Avg Prem Rate Increase* in Column (1) and *Insurer Far-ZIP Avg Prem Rate Increase* in Column (2), both expressed as a percentage of the coverage limit. The independent variables are its 12-month and 24-month lagged values. We also include renewal year-month fixed effects. Because we require two lags, the instruments are available from renewal year 2016 onward. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE IA.2

EFFECTS OF PREMIUM INCREASES ON SWITCHING INSURERS AT THE NEXT RENEWAL

	(1) 1(Switch Insurer at Next Renewal)	(2) Premium Rate Increase	(3) 1(Switch Insurer at Next Renewal)	(4) Premium Rate Increase	(5) 1(Switch Insurer at Next Renewal)
Premium Rate Increase	0.460***		0.888***		0.950***
	(26.58)		(10.77)		(11.54)
Insurer Other-ZIP Avg Prem Increase (residual)		0.883***			
		(87.52)			
Insurer Far-ZIP Avg Prem Increase (residual)				0.676***	
				(38.50)	
Lag Premium Rate	0.009***	0.022***	0.004*	0.022***	0.002
	(4.70)	(11.04)	(1.84)	(10.78)	(1.16)
CredScore/1k before Renewal	-0.033***	-0.039***	-0.001	-0.039***	0.002
	(-9.44)	(-21.37)	(-0.15)	(-21.55)	(0.37)
Climate Risk (std)/1k	-0.893***	0.385***	-1.155***	0.414***	-1.201***
	(-3.90)	(3.20)	(-4.46)	(3.30)	(-4.60)
OLS/2SLS Stage	OLS	1st Stage	2nd Stage	1st Stage	2nd Stage
Sample	All	All	All	All	All
Insurer Switch Measured At	Next Renewal	Next Renewal	Next Renewal	Next Renewal	Next Renewal
ZIP-by-Renew Month FE	Y	Y	Y	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y	Y	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y	Y	Y	Y
Y Mean	0.075	0.025	0.062	0.025	0.062
X Mean	0.027	-0.001	0.025	-0.001	0.025
X SD	0.059	0.023	0.054	0.024	0.054
N	45,072,018	29,889,519	29,889,519	29,651,554	29,651,554
Kleibergen-Paap Wald F stat			7,659		1,482

NOTES: Sources: ICE McDash; CoreLogic Climate Risk Analytics. This table is reproduced from Ge et al. (2026c). The dependent variable in Columns (1), (3) and (5) is an indicator for switching insurers within 12 months of renewal, equal to one when the insurer recorded at the household's next annual renewal differs from the insurer at the current one; in the first-stage Columns (2) and (4) it is *Premium Rate Increase*. Table IA.3 instead measures switching at the current renewal, against the insurer on the previous policy. Column (1) is OLS; Columns (2) and (4) report the first stage (with *Premium Rate Increase* as the dependent variable) and Columns (3) and (5) the corresponding second stage, instrumented with the Other-ZIP instrument (Columns 2–3) and the Far-ZIP instrument (Columns 4–5). Controls and fixed effects are as in Table 2. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

TABLE IA.3
INSTRUMENTS AND INSURER SWITCHING AT THE CURRENT RENEWAL

	(1)	(2)
	1(Switch Insurer at Current Renewal)	
Insurer Other-ZIP Avg Prem Increase (residual)	0.659***	
	(14.15)	
Insurer Far-ZIP Avg Prem Increase (residual)		0.525***
		(14.55)
Lag Premium Rate	0.041***	0.041***
	(22.00)	(21.43)
CredScore/1k before Renewal	0.005	0.004
	(1.31)	(1.27)
Climate Risk (std)/1k	-0.815***	-0.805***
	(-3.75)	(-3.78)
OLS/2SLS Stage	OLS	OLS
Sample	All	All
Insurer Switch Measured At	Current Renewal	Current Renewal
ZIP-by-Renew Month FE	Y	Y
Climate Risk Bin-by-Renew Month FE	Y	Y
Orig: Yr, CredScore, LTV, DTI FE	Y	Y
Y Mean	0.053	0.053
X Mean	-0.001	-0.001
X SD	0.023	0.024
N	35,854,295	35,571,286
R-squared	0.047	0.046

NOTES: Sources: ICE McDash; CoreLogic Climate Risk Analytics. This table presents results testing whether households are more likely to switch insurers when their previous insurers increase premiums. We use two instrumental variables as the independent variable separately. The dependent variable is an indicator of whether the switch-insurer occurs in policy year t, measured at the same time as the independent variable, that is, whether the insurer on the current policy differs from the insurer on the previous one. The main independent variables are the instruments constructed in policy year t. Both columns report OLS estimates on the full sample. This table therefore differs from Table IA.2 in two respects: the switch is measured at the current renewal rather than at the following one, and the right-hand side carries the instruments themselves rather than the instrumented premium rate increase. Controls and fixed effects are as in Table 2. Standard errors are corrected for clustering at the state-by-year level. ***p<0.01, **p<0.05, *p<0.1.

Internet Appendix IA.A The Magnitude of the First-Stage Coefficients

The first-stage coefficients in Table 2 are below one, 0.885 for the Other-ZIP instrument and 0.674 for the Far-ZIP instrument. This section explains why both fall below one, and why the Far-ZIP coefficient is the smaller of the two.

We restrict the leave-out averages to renewals with an unchanged insurer, deductible, and coverage type, which has two implications for the first-stage coefficient. The first concerns the focal household, whose own premium change carries no such restriction: it is the change in the rate the household pays, whichever insurer writes the policy and whatever terms it holds. Notice of a premium increase raises the probability that a household changes insurer at renewal in search of a lower rate (Table IA.3), raises its deductible, or moves from Replacement Cost Value to Actual Cash Value coverage, and the rate it then pays reflects those choices, so the increase a household realizes is smaller than the increase its prior insurer implemented elsewhere in the state.

The second concerns the policies entering the average. Policyholders elsewhere in the state face the same choice, and the restriction excludes those who exercise it, which is what makes each contributing observation a measure of the insurer's repricing rather than of another household's response to it. The policies that remain are those that retained their terms, which are disproportionately those that received smaller increases, so the average understates the insurer's repricing, and understates it by more where the repricing is larger. This raises the first-stage coefficient, in the opposite direction from the household's own response, and changes the units in which the instrument is measured. It leaves the second stage unaffected, because the second stage is the ratio of the instrument's effect on the outcome to its effect on the premium, and a rescaling of the instrument moves both terms by the same factor. To threaten identification, the selection would have to correlate with the focal household's outcome conditional on the fixed effects, which requires insurer-specific differences in adjustment elsewhere in the state; a shock to affordability common to the insurers serving a ZIP code in a given renewal month is absorbed. Figure IA.3 tests the instruments as constructed, this selection included, and finds them orthogonal to predetermined characteristics.

A further reason the first-stage coefficient falls below one is that the leave-out average is a noisy proxy of insurers' repricing action. We observe an insurer's premium changes only for the policies in our data, which cover mortgaged single-family properties in the servicing portfolios described in Section 3.1. Each average is therefore computed from a subsample of the policies an insurer renewed in that state and month, and measures the insurer's repricing with error. Classical error in a right-hand-side variable attenuates the estimated coefficient toward zero. What the second stage estimates is the ratio of the instrument's effect on the outcome to its effect on the premium, and noise in the instrument attenuates both by the same factor. The second-stage coefficient is therefore unaffected.

The Far-ZIP first-stage coefficient of 0.674 is below the 0.885 of the Other-ZIP instrument for two reasons. First, an insurer's rate changes are more similar in nearby ZIP codes, which share catastrophe exposure and loss experience and often fall within the

same rating territory of a rate filing, and the Far-ZIP average excludes those ZIP codes. Second, the Far-ZIP average is computed from roughly half of the ZIP code pairs available to the Other-ZIP average, so it uses fewer policies and carries more of the estimation error. Both reasons also imply the fall in the Kleibergen-Paap Wald F statistic, from 7,787 to 1,496. The second-stage estimates remain close between the two instruments.

Internet Appendix IA.B A Model of the Move-versus-Switch Decision

A household facing a premium increase has two remedies. It can switch insurers, at a cost of a few hours of search, or it can relocate, at a cost of several percent of the value of the house. Because the inexpensive remedy is available to every household, an account of the relocation response documented in this paper must explain why a premium increase ever induces the costly action. This appendix presents a model that does so.

The model rests on three observations. First, insurers quote different premiums for the same property, and only the difference between a household's own premium and the best locally available quote can be undone by switching. Second, the price of insuring a location is charged by every insurer serving it, so it can be escaped permanently only by relocating; insurers incorporate a change in that price at different dates, however, so at any renewal date part of it is temporarily avoidable by switching to an insurer that has not yet repriced. Third, households accumulate non-insurance reasons to move, so an insurance shock need not be large enough to justify relocation on its own in order to induce it.

A fourth observation organizes the link between the model and the empirical design. A household observes the premium it is quoted, not the distribution of premiums quoted to its neighbors, so it cannot separate the part of an increase that its own insurer has implemented ahead of others from the part that reflects a durable change in the price of insuring its location. Our fixed effects perform that separation for the econometrician. They do not perform it for the household, and Section IA.B.3 shows that the distinction determines what the design can recover.

Section IA.B.1 states the decision problem in terms of two state variables and characterizes the household's choice. Section IA.B.2 maps premium shocks into those states. Section IA.B.3 derives comparative statics with respect to the underlying shocks and with respect to the increase the household observes. Section IA.B.4 introduces financial constraints, and Section IA.B.5 collects the proofs.

IA.B.1 The decision problem

A household occupies a property, lives for $t = 0, 1, \ldots, T$, and discounts at $\beta \in (0, 1]$. At date 0 it chooses among remaining with its current insurer and property, switching insurers, and relocating.

Housing mismatch

Households accumulate a dollar-denominated mismatch $m \geq 0$ between the current house and the house they would choose conditional on moving. It evolves as

$$m_t \;=\; \rho_m \, m_{t-1} + \eta_t^m, \qquad \eta_t^m = \begin{cases} \sim \Gamma(k, \theta) & \text{with probability } \pi, \\ 0 & \text{with probability } 1 - \pi, \end{cases} \tag{2}$$

with $m = 0$ at the date of purchase and $\rho_m \in [0, 1]$ permitting the relevance of past events to decay. Innovations are infrequent and highly dispersed: in most years no event occurs,

and when one does, whether a change of employment, a change in family composition, or a health event, its magnitude is drawn from a skewed distribution, so that occasional realizations are large enough to induce relocation irrespective of insurance costs. Moving realizes m; switching insurers does not affect it. Throughout, m denotes the date-0 stock.

Costs

Switching costs $\kappa^s \geq 0$ represent search and administrative effort. Relocation costs

$$c^m \;=\; \kappa^m - \delta L^{mtg}, \qquad \delta \;\equiv\; \text{origination rate} - \text{current market rate}, \tag{3}$$

where $\kappa^m > 0$ denotes transaction and relocation costs and $L^{mtg} > 0$ a scale factor for the household's mortgage. The second term is the mortgage consequence of relocation: a household that moves surrenders its existing loan and originates a new one at prevailing rates, which is advantageous when the origination rate exceeds the current rate ($\delta > 0$) and costly otherwise. The latter case is mortgage lock-in (Fonseca and Liu, 2024).[23] We maintain $c^m > \kappa^s$: relocation is the costlier action. Comparative statics in δ are taken over the range on which this ordering holds.

Assumption 1 (Re-shopping accompanies relocation)**.** A household that relocates takes out a new policy on the new property and selects the best available insurer there. The incumbent insurer's advantage at the origin has two parts. The part that is specific to the origin, namely its position in the sequence in which insurers reprice that location, does not transfer to the new property. The part that attaches to the household, such as multi-policy discounts, underwriting classification, or claims history, may transfer, and the incremental cost to a mover of selecting a non-incumbent insurer is κ^s_M, with $0 \leq \kappa^s_M < \kappa^s$. The baseline sets $\kappa^s_M = 0$.

A mortgage closing requires a new policy on the new property, priced as new business, and the search costs of obtaining it are shared with the other searches that a housing transaction entails. Relocation therefore accomplishes what switching accomplishes. Consistent with $\kappa^s_M < \kappa^s$, thirty-eight percent of movers in our sample change insurer at the move. That the remainder do not is consistent with a household-specific component of the incumbent's advantage that transfers across properties, which corresponds to $\kappa^s_M > 0$.

The two state variables

Two objects summarize the household's position. The first is the present value of the amount by which its own insurer prices above the best locally available quote, which switching removes. We write it as $x \geq 0$ and construct it in Section IA.B.2. The second

[23] The product δL^{mtg} is a reduced form and not a valuation. A rate differential applied to an outstanding principal balance has the units of one year's interest, whereas the object required here is the present value of exchanging the existing loan for one at current rates, which depends on the remaining maturity, the amortization schedule, the prepayment option, and the discount rate. We therefore read L^{mtg} as a duration-weighted balance, equivalently δ as a present-value-equivalent spread, chosen so that the product carries those units; we do not derive it. The results below use only $\partial c^m / \partial \delta = -L^{mtg} < 0$, so the lock-in comparative static does not depend on the normalization, and no dollar magnitude is claimed for the term.

is the present value of what only relocation delivers, namely the accumulated housing mismatch together with the eventual difference in the price of insurance between the origin and the destination:

$$y \;\equiv\; m + A_L\, G. \tag{4}$$

Here G is the eventual location price gap and A_L its present-value loading; both are defined in Section IA.B.2, which also shows how premium shocks move them. The two components of y are perfect substitutes in the decision: a dollar of housing mismatch and a dollar of location price gap have the same effect on the choice.

Payoffs are measured net of remaining in place:

$$v(\varnothing) = 0, \qquad v(S) = x - \kappa^s, \qquad v(M) = y + x - c^m. \tag{5}$$

Switching removes the wedge at cost κ^s; it changes neither the local price of insurance nor the match between household and house. Relocation realizes the mismatch, exchanges the origin's price of insurance for the destination's, and, by Assumption 1, also removes the wedge, at cost c^m.

Proposition 1 (Decision rule)**.** *Let $c^m > \kappa^s \geq 0$. The optimal choice is*

$$\begin{aligned} \text{move} &\iff y > c^m - \min\{x, \kappa^s\}, \\ \text{switch} &\iff x > \kappa^s \;\; \textit{and} \;\; y \leq c^m - \kappa^s, \\ \text{stay} &\iff x \leq \kappa^s \;\; \textit{and} \;\; y \leq c^m - x. \end{aligned} \tag{6}$$

The relocation frontier $y^*(x) = c^m - \min\{x, \kappa^s\}$ is decreasing in x on $[0, \kappa^s)$, with slope -1, and constant at $c^m - \kappa^s$ for $x \geq \kappa^s$; it has a kink at $(\kappa^s,\, c^m - \kappa^s)$.

The household relocates when the value of doing so, comprising the relocation gap y and whatever wedge the move removes, exceeds the cost of relocation. The term $\min\{x, \kappa^s\}$ reduces the relocation threshold, but by at most κ^s: beyond that point the wedge can be removed by the inexpensive action, and it therefore ceases to constitute a reason to move. The household switches when the wedge alone justifies the search cost but the relocation gap does not justify the incremental cost of moving. Two properties of the solution merit emphasis.

First, where switching is optimal absent relocation considerations, that is, for $x > \kappa^s$, the comparison between the two remedies is

$$v(M) - v(S) \;=\; y - (c^m - \kappa^s) \;=\; m + A_L G - (c^m - \kappa^s), \tag{7}$$

which does not depend on the wedge. Both remedies remove it, so it cannot determine the choice between them. The household relocates rather than switches when its accumulated mismatch, together with the eventual price gap between origin and destination, exceeds the incremental cost of relocating over switching.

Second, the availability of the inexpensive remedy caps the contribution of the wedge to the relocation decision at κ^s. For $x < \kappa^s$ a household that would not switch on its own account nonetheless removes its wedge without additional cost when it moves, so the wedge reduces the effective cost of relocation dollar for dollar. For $x > \kappa^s$ additional

wedge induces switching but does not provide a motivation for relocation. No corresponding bound applies to y. This asymmetry is the reason the empirical design requires the treatment in Section IA.B.3: it places a ceiling on what any variation that moves only the wedge can do to the relocation margin.

IA.B.2 Mapping premium shocks into the decision states

This section defines the wedge, the eventual location gap, and their present-value loadings, and derives how the shocks move them. The distinction that organizes it is between a shock to the price of insuring a location and the portion of that shock already reflected in insurers' quotes at a given renewal date.

Quotes, the local price floor, and the wedge

The insurers j serving the property quote per-year premiums P_j that differ across insurers. Let P_c denote the current insurer's quote, let a index the insurer offering the cheapest alternative quote $P^a \equiv \min_{j \neq c} P_j$, and define

$$W \equiv \max\{0,\ P_c - P^a\}, \qquad L \equiv \min\{P_c, P^a\}. \tag{8}$$

We refer to W as the *wedge* and to L as the *local price floor*. Every local quote can be written as the local price floor plus a nonnegative insurer wedge. Switching can remove the wedge but does not change L: it cannot produce a quote below the current floor. The wedge is zero for a household whose current insurer is the cheapest available, so its distribution has a mass point at zero. Throughout we hold the identity of a fixed over the horizon, so the floor compares the incumbent with a single alternative. Allowing a third insurer to undercut both at a later date would lower the floor further and enlarge the wedge; none of the signs below depend on this.

Shocks and staggered pass-through

Two shocks move quotes at date t. The *location shock* $\varepsilon_L \geq 0$ is an increase in the price of insuring the location, arising from risk reassessment, disaster-driven repricing, or regulatory pass-through; the repricings the paper studies are upward revisions, and we restrict the shock accordingly. Every insurer serving the location comes to charge it, but they do not incorporate it to the same extent at the same date. The *insurer shock* ε_I is the current insurer's own movement, arising from its loss experience, expense loadings, or pricing-model updates, and is independent of ε_L.

Assumption 2 (Staggered pass-through and eroding insurer components)**.** Let $\lambda_{j,t} \in [0, 1]$ denote the cumulative share of the location shock that is reflected in insurer j's quote as of date t, with $\lambda_{j,t-1} = 0$ because the shock occurs at t. The $\lambda_{j,t}$ are dispersed across insurers at any date and converge to unity as repricing is completed. Let $\kappa_{t+h} \in [0, 1]$ denote the share of a date-t insurer-specific movement remaining in the quote at date $t + h$, with $\kappa_t = 1$ and κ_{t+h} declining toward zero as that component erodes through competition, regulatory review, and the insurer's own mean reversion. Writing $\Delta\lambda_{j,t} \equiv \lambda_{j,t} - \lambda_{j,t-1}$,

the change in insurer j's quote at date t is

$$\Delta P_{j,t} \;=\; \Delta\lambda_{j,t}\,\varepsilon_L + \varepsilon_{I,j}, \tag{9}$$

where $\varepsilon_{I,j}$ is the innovation to insurer j's own component, mean zero and independent of ε_L; for the current insurer we write $\varepsilon_{I,c} \equiv \varepsilon_I$. No further insurer-specific innovations arrive after date t, but the date-t component decays along κ_{t+h}, so for $h \geq 1$

$$\Delta P_{j,t+h} \;=\; \Delta\lambda_{j,t+h}\,\varepsilon_L + \big(\kappa_{t+h} - \kappa_{t+h-1}\big)\,\varepsilon_{I,j}. \tag{10}$$

Insurers update loss models, file rate requests, and obtain regulatory approvals on different schedules, so at a given renewal date insurers serving the same location differ in how much of a common repricing they have implemented.

We set the inherited wedge W_{t-1}, which summarizes the pass-through differentials left over from earlier location shocks together with accumulated insurer-specific differences, to zero, so that the two quotes are equal at $t-1$. Because the floor in (8) is a minimum taken over the incumbent and the alternative, quantities are stated below as levels relative to that common starting point rather than as increments. This keeps the expressions exact whether or not the incumbent is the cheaper of the two at any given date, and it absorbs a crossing of the two quotes into the minimum rather than assuming it away. Writing $\varepsilon_{I,a}$ for the alternative's own movement, which we normalize to zero, Equations (9) and (10) give

$$\begin{aligned} P_{c,t+h} - P_{c,t-1} &= \lambda_{c,t+h}\,\varepsilon_L + \kappa_{t+h}\,\varepsilon_I, \\ L_{t+h} - L_{t-1} &= \min\big\{\lambda_{c,t+h}\,\varepsilon_L + \kappa_{t+h}\,\varepsilon_I,\ \lambda_{a,t+h}\,\varepsilon_L\big\}, \end{aligned} \tag{11}$$

by summation of the quote changes in Assumption 2 from t to $t+h$. Differencing the two lines gives the wedge, exactly and on the whole space,

$$W_{t+h} \;=\; \big(P_{c,t+h} - P_{c,t-1}\big) - \big(L_{t+h} - L_{t-1}\big) \;=\; \max\big\{0,\ (\lambda_{c,t+h} - \lambda_{a,t+h})\,\varepsilon_L + \kappa_{t+h}\,\varepsilon_I\big\}, \tag{12}$$

since $u - \min\{u, v\} = \max\{0, u - v\}$. The mass point at $W = 0$ is the event that the incumbent is the cheaper of the two, and it is carried by the maximum rather than excluded by assumption.

Three features of (11) and (12) matter for what follows. The observed premium change at date t is the current insurer's own movement plus the portion of the location shock that this insurer has so far passed through; at later dates it reflects further pass-through of the location shock together with the decay of the insurer component. The location-price component of the cost of remaining in place rises along the lower of the two paths, so it is governed at each date by the pass-through of whichever insurer is behind; the remainder of what the household pays is the wedge, which switching can remove. And a location shock enters the switchable wedge whenever the household's insurer has passed through more of it than the cheapest alternative. Because $\lambda_{c,t+h}$ and $\lambda_{a,t+h}$ both converge to unity as repricing completes, their difference vanishes, and because κ_{t+h} declines to zero, so does the insurer-specific part. The wedge is transitory in both of its components.

The location gap and the two present-value loadings

Let $\varepsilon_{L'}$ denote the location shock at the destination and u_L the standing gap between origin and destination that predates the current shock. Because the household's relocation decision turns on what it will pay over its horizon rather than on this date's quotes, we define the *eventual location gap*

$$G \;\equiv\; (\varepsilon_L - \varepsilon_{L'}) + u_L, \tag{13}$$

the difference between the origin's and the destination's price of insurance once repricing is complete. As a normalization, the benchmark sets $u_L = 0$ and $\varepsilon_{L'} = 0$; Proposition 5 below evaluates a comparative static at that point. G is negative for a household that would relinquish an inexpensive location by moving.

The loadings convert these annual amounts into present values, and they differ because the objects have different paths. The wedge is transitory in both of its parts: the pass-through differential disappears as the lagging insurer catches up, and the insurer-specific difference erodes along κ_{t+h}. The location gap is durable, but it too reaches the household only as repricing proceeds, since a household that remains at the origin pays the alternative's path once it has switched. Every loading therefore carries a pass-through path explicitly. Accordingly, define

$$A_\Lambda \equiv \sum_{h=0}^{T} \beta^h \,\mathbb{E}_t\big[\lambda_{c,t+h} - \lambda_{a,t+h}\big], \qquad A_I \equiv \sum_{h=0}^{T} \beta^h \,\mathbb{E}_t\big[\kappa_{t+h}\big], \qquad A_L \equiv \sum_{h=0}^{T} \beta^h \,\mathbb{E}_t\big[\lambda_{a,t+h}\big], \tag{14}$$

with $A_\Lambda \geq 0$ and $A_I, A_L > 0$. The restriction $A_\Lambda \geq 0$ reflects the selection that defines a: the cheapest alternative at date t is inexpensive in part because it lags, so the incumbent's expected cumulative pass-through weakly exceeds that of the insurer quoting the floor. The first state variable values the wedge over the horizon,

$$x \;\equiv\; \max\big\{0,\; A_\Lambda\,\varepsilon_L + A_I\,\varepsilon_I\big\} \;\geq\; 0. \tag{15}$$

Each of the first two loadings is a discounted expectation of a path rather than an annuity factor derived from a separate persistence parameter, and none of the three divides by a quantity that the results below allow to be zero.[24]

The loading on G is the alternative's path rather than a plain discount sum, and that choice is what makes the two states add up. On the region where the wedge is positive, a household that does nothing pays $\sum_h \beta^h \mathbb{E}_t[\lambda_{c,t+h}]\varepsilon_L + A_I\varepsilon_I$ by (11); one that switches once pays the alternative's path, so it saves $A_\Lambda\varepsilon_L + A_I\varepsilon_I$, which is exactly x; and one that relocates pays the destination's price, so it saves the whole of $\sum_h \beta^h \mathbb{E}_t[\lambda_{c,t+h}]\varepsilon_L + A_I\varepsilon_I$. Since $A_\Lambda = \sum_h \beta^h \mathbb{E}_t[\lambda_{c,t+h}] - A_L$ under (14), the mover's saving is $A_L G + x$ net of mismatch, which is $y + x$. The decomposition is therefore an identity rather than an approximation, and no component of the shock is credited twice.

[24] The present value of the wedge path is $\sum_h \beta^h \mathbb{E}_t[W_{t+h}]$, with the censoring in (12) applied period by period; (15) applies it to the discounted index instead, and understates that sum when the index changes sign along the path. The two coincide when the pass-through differential and κ_{t+h} decay proportionally. The results below use only that $x \geq 0$ and that x is nondecreasing in each shock.

The identity is confined to that region. Where the incumbent is itself the cheaper quote the wedge is censored at zero, and the censoring condition $A_\wedge \varepsilon_L + A_I \varepsilon_I < 0$ says precisely that the household's own path lies below the alternative's. Such a household escapes its own insurer's path by relocating, which is worth $\sum_h \beta^h \mathbb{E}_t[\lambda_{c,t+h}]\varepsilon_L + A_I \varepsilon_I$, whereas (5) credits it with $A_L G$ alone. The model therefore overstates the gain from relocating for households whose current insurer is unusually inexpensive, by exactly $|A_\wedge \varepsilon_L + A_I \varepsilon_I|$. Two things limit the consequence. The direction is signed rather than unknown, and the households affected are those for which $x = 0$, so the switching margin is inactive and the wedge channel of Proposition 3 is absent for them.[25]

Loading G at the plain discount sum $\sum_h \beta^h$ would break this. It values the location gap as though it were already fully in prices, which leaves no room for the pass-through differential $A_\wedge \varepsilon_L$ that the wedge already carries, so the mover would be credited twice with the part of the location shock that repricing has not yet delivered. The over-credit would be exactly $\sum_h \beta^h \mathbb{E}_t[1 - \lambda_{a,t+h}]\, \varepsilon_L$, the discounted mass of unfinished repricing, and nothing in Assumption 2 bounds it, since that assumption requires the shares to converge but imposes no rate and no completion date.[26]

The present value of movers' premium savings reported in Section 8.5 is a related but distinct calculation: it capitalizes a constant estimated annual saving over the expected tenure, whereas the saving in the model phases in as repricing completes. The two agree when repricing is fast relative to the horizon and otherwise differ by the same discounted mass of unfinished repricing. We treat the reported figure as a constant-flow approximation to $A_L G$ rather than as the same object.

Consistent with the transitory wedge, Ge et al. (2026c) find that the premium reduction obtained by switching insurers is partial and partly temporary, which is what Assumption 2 implies: a household that switches moves to an insurer that is inexpensive partly because it has not yet repriced, and that advantage erodes as the new insurer catches up.[27]

IA.B.3 Comparative statics

We state the comparative statics within a location and conditional on the insurers' pass-through profile, which is the comparison the empirical design makes: our fixed effects compare households in the same ZIP code and renewal month. Within such a cell the

[25]The discrepancy could be removed by valuing G at the path the household would itself follow if it stayed, namely the lower of the incumbent's and the alternative's, which amounts to making the loading on G household-specific rather than common within a location. We do not do so here. It would buy exactness on a region where the switching margin is inactive, at the cost of carrying a household-specific loading through comparative statics that treat the loadings in (14) as constants within a cell, and through the independence structure assumed at the head of Section IA.B.3.

[26]Because κ_{t+h} declines to zero, the eventual price of insuring the origin is the location price common to all insurers serving it, and carries no insurer-specific term. This is what allows G in (13) to be a difference between two locations rather than between two insurer-location pairs.

[27]There is no double counting in (5), as the accounting after (14) shows. The wedge term values only the advantage of leaving the current insurer for the alternative, and the location term values only what remains after that switch, namely the alternative's own path toward the eventual location price. Because the two loadings partition $\sum_h \beta^h \mathbb{E}_t[\lambda_{c,t+h}]$ between them, the pass-through differential is counted once, in the wedge. The identity holds where the wedge is positive; the censored region, where it does not, is treated after (14).

location shock ε_L is common to all households and the loadings in (14) are constants. Households differ in the insurer shock ε_I, which generates the cross-sectional distribution F_x of the wedge state, and in accumulated mismatch m, which generates F_y. We assume m is independent of the insurer-side variables, so x and y are independent within the cell; that F_y has a continuous density f_y of full support; and that F_x has a continuous density f_x on $(0, \infty)$. The mass point at $x = 0$ noted after (8) contributes no separate term to the derivatives below: the relocation frontier at $x = 0$ equals the threshold c^m that applies at the atom itself, so mass crossing the censoring boundary moves between regions with identical payoffs. Population responses aggregate these cell-level responses, which preserves their signs. Proofs are in Section IA.B.5.

Proposition 2 (Effects of the two shocks)**.** *(i)* Insurer shock. *Both margins respond positively:*

$$\frac{\partial \Pr(S)}{\partial \varepsilon_I} = A_I f_x(\kappa^s) F_y(c^m - \kappa^s) > 0, \qquad \frac{\partial \Pr(M)}{\partial \varepsilon_I} = A_I \int_{x<\kappa^s} f_y(c^m - x) f_x(x)\, dx \;\geq\; 0. \tag{16}$$

(ii) Location shock. *Relocation responds positively,*

$$\frac{\partial \Pr(M)}{\partial \varepsilon_L} = A_L\, \mathbb{E}_x\big[f_y(y^*(x))\big] + A_\Lambda \int_{x<\kappa^s} f_y(c^m - x) f_x(x)\, dx \;>\; 0, \tag{17}$$

and switching responds through two opposing channels,

$$\frac{\partial \Pr(S)}{\partial \varepsilon_L} = \underbrace{A_\Lambda\, f_x(\kappa^s) F_y(c^m - \kappa^s)}_{\textit{pass-through differential}} - \underbrace{A_L \big[1 - F_x(\kappa^s)\big] f_y(c^m - \kappa^s)}_{\textit{substitution toward relocation}}, \tag{18}$$

so that switching rises with a location shock if and only if

$$A_\Lambda \;>\; \frac{A_L \big[1 - F_x(\kappa^s)\big] f_y(c^m - \kappa^s)}{f_x(\kappa^s)\, F_y(c^m - \kappa^s)}. \tag{19}$$

At $A_\Lambda = 0$, that is, when insurers reprice simultaneously at every horizon, the first term vanishes and switching alone is strictly less likely after a location shock.

Each shock reaches the two margins through the state it moves. The insurer shock enters only the wedge, and therefore induces switching directly and relocation only among households whose wedge is below the switching threshold, for which the move removes the wedge without additional cost. The location shock enters the eventual location gap in full and the wedge temporarily, to the extent that the household's insurer has repriced ahead of the cheapest alternative. Relocation responds unambiguously to both shocks.

The response of switching to a location shock combines two channels. The shock raises the wedge of households whose insurer has repriced ahead of the cheapest alternative, which induces switching; it simultaneously raises the value of relocating, which induces some households that would have switched to move instead. Condition (19) determines

which channel dominates. We state it as a condition and do not assert its sign, because its two sides are not separately identified by the moments we observe.

One element of it does map to an observable. Since relocation is infrequent, $F_y(y^*(x))$ is close to one at every x, so $F_y(c^m - \kappa^s) \approx 1$ and $\Pr(S) \approx 1 - F_x(\kappa^s)$: the share of households whose wedge exceeds the switching threshold is approximately the switching rate itself, which is a minority outcome. The remaining elements are not read off the data in the same way. A low relocation rate restricts the mass of y above the frontier but does not by itself bound the density f_y at it, since a distribution can place little mass above a point while remaining dense at it; the argument for a small f_y at the frontier is the one given for Proposition 5, which rests on the mismatch process in (2) rather than on the relocation rate. And the switching response in Proposition 2 identifies the product $A_I f_x(\kappa^s)$, not $f_x(\kappa^s)$ alone, so a strong switching response bounds that density only relative to the loading. The ratio of A_L to the remaining loadings works against the condition holding, and (14) places no bound on that ratio: the loadings weight different paths, and which is larger turns on how quickly the lagging insurer catches up and the insurer-specific component erodes. Signing (19) therefore requires calibration of the two densities and the loadings, which we do not attempt; the condition is recorded so that the two channels are visible, not as a prediction.

What the empirical design moves

Proposition 2 states responses to the two structural shocks. Neither is observed, by us or by the household, and the two enter the design differently. Our specifications compare households in the same ZIP code and renewal month, so the component of ε_L common to that cell is absorbed. The variation that identifies our estimates is therefore variation in $\lambda_{c,t}$, the share of the location shock that a household's own insurer has implemented by its renewal date.

That variation moves x and does not move y. By the second property of Proposition 1, its effect on relocation is bounded. The next remark states the bound, and the proposition that follows states the channel through which the design nonetheless recovers a relocation response that is not so bounded.

Two departures from the preceding subsection are needed and are maintained for the remainder of this one. Households within a cell now differ in $\lambda_{c,t}$, which is the variation the instruments supply, rather than being conditioned on a common profile. And the household is no longer assumed to know the shocks it faces, so its decision is stated in terms of what it can observe.

Remark 1 (The bound under full information)**.** Suppose the household knows ε_L. Then variation in $\lambda_{c,t}$ holding ε_L fixed changes x and leaves y unchanged, and by Proposition 1 the induced change in $\Pr(M)$ operates only through the region $x < \kappa^s$. The relocation response to such variation is bounded above by the response to an increase of κ^s in the wedge, however large the underlying location shock. Under this reading the design would recover a relocation response of the order of the switching cost.

The remark is the reason the informational structure matters, and it must be stated as an assumption rather than asserted, because the decision rule of Proposition 1 was derived for a household that knows its own y.

Assumption 3 (Household information and inference)**.** The household observes its own quote change $d \equiv \Delta P_{c,t} = \lambda_{c,t}\varepsilon_L + \varepsilon_I$, the competing quote, and the valuation x of (15), together with its own m and u_L. It does not observe $\lambda_{c,t}$, $\lambda_{a,t}$, ε_L or ε_I separately, and it does not observe the quotes received by other households at its location. It holds a prior over $(\lambda_{c,t}, \lambda_{a,t}, \varepsilon_L, \varepsilon_I)$ under which the pass-through shares are independent of the shocks, and forms the posterior mean

$$\mu(d, x) \;\equiv\; \mathbb{E}\big[\varepsilon_L \;\big|\; d,\; x\big]. \tag{20}$$

We assume μ is nondecreasing in each of its arguments.

The monotonicity is the formal content of the statement that a household cannot tell the two components apart. Both observables load positively on the location shock: d is increasing in ε_L for every $\lambda_{c,t} > 0$, and by (15) so is x whenever $A_\wedge > 0$, since a household whose insurer has passed through more of a location shock than the cheapest alternative carries a larger wedge on that account. A larger realization of either is therefore evidence of a larger location component under any prior in which the pass-through shares are independent of the shocks. The property is one of the posterior rather than a consequence of the earlier assumptions: with $\lambda_{c,t}$ known it follows from log-concavity by the argument used in Proposition 6, but with $\lambda_{c,t}$ unknown the observed increase involves a product of unknowns, and log-concavity is not preserved under products, so the property is imposed. Only the direction is used below; nothing turns on the strength of either response, and the case in which x carries no information, so that μ is constant in its second argument, is permitted.

Remark 2 (Decision under incomplete information)**.** Under Assumption 3 the household maximizes expected payoff. Because x and κ^s are observed and the payoffs (5) are linear in y, Proposition 1 applies verbatim with y replaced by the assessed relocation gap

$$\hat{y} \;\equiv\; \mathbb{E}\big[y \;\big|\; d, x, m\big] \;=\; m + A_L\big(\mu(d, x) + u_L\big). \tag{21}$$

In particular the household relocates if and only if $\hat{y} > c^m - \min\{x, \kappa^s\}$, and the frontier retains its slope of -1 below κ^s and its kink at κ^s.

The conditioning that our fixed effects perform is available to the econometrician and not to the household. Within a cell every household faces the same ε_L; households whose insurers have repriced further see larger increases and assess the durable component higher. That dispersion in assessments is what the design moves.

Proposition 3 (Relocation response to insurer repricing)**.** *Maintain Assumption 3 and consider the experiment the instruments perform: within a cell, hold $\varepsilon_L \geq 0$ and the distributions of m and ε_I fixed, and compare households whose incumbent has implemented $\lambda'_{c,t}$ of the location shock with those at $\lambda_{c,t} < \lambda'_{c,t}$, where the incumbent that has repriced further at date t is expected to remain weakly further along at every subsequent date, so that $A'_\wedge \geq A_\wedge$. Write d and $d' = d + (\lambda'_{c,t} - \lambda_{c,t})\varepsilon_L$ for the corresponding observed increases, x and $x' \geq x$ for the corresponding wedge states, and*

$$\Delta\hat{y} \;\equiv\; A_L\big[\mu(d', x') - \mu(d, x)\big] \;\geq\; 0 \tag{22}$$

for the induced revision in the assessed relocation gap. Then:

(i) The probability of relocation is weakly higher at $\lambda'_{c,t}$ than at $\lambda_{c,t}$.

(ii) The increase in that probability is at least the mass of households whose assessed gap lies within $\Delta\hat{y}$ below the relocation frontier. Under full information μ is the known ε_L and does not vary with what the household observes, $\Delta\hat{y} = 0$, and the response reduces to the bound of Remark 1,

$$\Pr\big(\hat{y} \in (c^m - \kappa^s,\, c^m]\big). \tag{23}$$

The response is therefore not bounded by (23) *whenever $\Delta\hat{y} > \kappa^s$.*

The proposition is stated in the variation the design supplies rather than in the latent shocks, and in the action the household takes rather than in whether relocation is optimal at the realized shock. Part (i) requires no distributional assumption: raising the incumbent's pass-through share weakly raises both terms on the left of the relocation condition in Remark 2, so the set of relocating households weakly expands. Part (ii) locates the difference between the two informational regimes in a single quantity. Under full information the frontier can fall by at most κ^s, because the only channel is the wedge and the frontier is flat beyond the kink. Under incomplete information it falls by a further $\Delta\hat{y}$, which is the revision in the assessed annual location cost valued over the horizon, and which bears no relation to the search cost. The condition in part (ii) is therefore not demanding if that revision is of the order of the present values reported in Section 8.5, since κ^s is a few hours of search; the magnitude discussion following Proposition 5 states that supposition and what it does and does not license.

Two features of this account are worth stating plainly. Households whose insurers repriced early assess the durable component higher than their neighbors although every household in the cell faces the same ε_L; each assessment is correct given the household's own information, and the dispersion among them is what the instruments move. And the estimated response would be smaller in a setting where households could observe local benchmark prices, since μ would then vary less with what the household observes. We regard this as a property of the environment rather than a limitation of the design.[28]

Remark 3 (An alternative: the current premium as the reference)**.** The cap in Proposition 1 arises because the household nets out what it could obtain by switching. Suppose instead that it evaluates relocation against the premium it currently pays. Then $\min\{x, \kappa^s\}$ in Proposition 1 is replaced by x, the relocation frontier is $y^*(x) = c^m - x$ at every x with no kink, and the integrals over $\{x < \kappa^s\}$ in Proposition 2 extend over the whole support. The bound of Remark 1 does not then arise at all, since the frontier never becomes flat. This mechanism and the informational one are complements rather than substitutes: either is sufficient to place the relocation response above the bound of Remark 1, and the evidence

[28] We state no derivative of the relocation probability with respect to d. Conditioning on d fixes $\varepsilon_I = d - \lambda_{c,t}\varepsilon_L$ when $\lambda_{c,t}$ is known, so the state variables become functions of a single shock and the independence maintained in Proposition 2 fails conditionally; the shape of the conditional law of ε_L also varies with d under log-concavity alone. A closed-form decomposition into a location channel and a wedge channel is available if the shocks are Gaussian, $\lambda_{c,t}$ is known and the censoring in (15) is ignored. We do not use it: the ordering in part (i) and the comparison in part (ii) are what the argument needs.

in Ge et al. (2026c) that households do respond on the switching margin bears on how much weight the reference-point reading can carry.

Measurement

The correlation of disaster risk within a state is what allows the location component to be measured. An insurer's repricing in distant ZIP codes of the same state is informative about how much of the focal location's repricing it has implemented, so the Far-ZIP instrument is, in the terms of the model, a shifter of $\lambda_{c,t}$ rather than of ε_L, together with whatever insurer-specific component is embedded in the statewide rate action. This is exactly the experiment of Proposition 3: it raises $\lambda_{c,t}$ with ε_L held at its cell value, which raises the observed increase d and, through Assumption 3, both the household's assessment of the durable component and its wedge. This is the sense in which the variation used for identification may reflect which insurer has repriced, as discussed in Section 6. What the design recovers is therefore a response to assessed rather than realized location costs, and Proposition 3 is what establishes that this response is not confined to the bound of Remark 1.[29]

Further comparative statics

Proposition 4 (Mortgage lock-in and relocation)**.** *A reduction in the cost of relocation raises the probability of relocation:* $\frac{\partial \Pr(M)}{\partial \delta} > 0$.

Since δ enters only through c^m, a relaxation of mortgage lock-in lowers the relocation frontier at every x, and relocation expands at the expense of both remaining in place and switching. This is the interaction estimated in Table 6.

Proposition 5 (The marginal relocating household)**.** *Evaluated at the benchmark normalization, the relocation response to the standing location gap is*

$$\frac{\partial \Pr(M)}{\partial u_L} \;=\; A_L \, \mathbb{E}_x\big[\, f_y\big(y^*(x)\big)\big], \tag{24}$$

the density of the relocation gap at the frontier, scaled by the loading.

The proposition characterizes a shift in the standing gap, which moves y and leaves the wedge state unchanged. It is not the response our design recovers: by Proposition 3 the instruments move the assessed gap and the wedge together, and that proposition

[29] Two distinct objects should not be conflated here. The household's observed increase d is the empirical counterpart of the *endogenous* regressor, the premium rate increase the household itself faces, which enters the estimation unresidualized. The instruments are the leave-out averages of the same insurer's changes elsewhere in the state, and it is those averages, not d, that Section 6 residualizes against their own 12 and 24 month lags and renewal-month effects. In the terms of the model the residualization isolates the innovation in an insurer's repricing relative to its own recent history, which is what makes the instrument a shifter of the current installment of pass-through rather than of the level already accumulated. It does not separate the location shock from the insurer's cumulative share, so the variation it delivers still moves d without revealing its composition. The model is stated in the observed increase precisely because that separation is unavailable to the household, and Assumption 3 makes that unavailability the mechanism rather than a caveat.

delivers an ordering and a lower bound rather than a derivative. What the two share is that the gap-side channel operates through the density of the relocation gap at the frontier, and that is the feature the rest of this subsection uses. It carries two implications.

The first concerns magnitude, and it can be stated without a derivative. Suppose the revision in the assessed gap induced by a premium increase of the size we study is of the order of the present values reported in Section 8.5, where movers are estimated to realize about \$13,580.[30] Such a revision is a substantial fraction of a relocation cost of the order of several percent of the value of the house, but remains below it, so a premium increase of realistic size does not by itself make relocation worthwhile for a household away from the frontier. It alters the decision of households for which the assessed gap already lies close to the frontier, either because recent events have raised m or because a standing price gap has raised $A_L G$. Were mismatch to accumulate smoothly, so that its stationary distribution were approximately uniform on $[0, c^m]$, then by part (ii) of Proposition 3 the response would be *at least* the ratio of that revision to c^m, which exceeds the estimate in Table 4 by orders of magnitude. Because this is a lower bound, the comparison does not require the response to be characterized exactly. A process with infrequent and highly dispersed innovations, as in (2), is not subject to it: most households have experienced no recent event, and realizations large enough to approach the frontier typically exceed it and induce relocation outright, leaving little mass immediately below it.

The second concerns the identity of the households whose behavior changes. The model implies that households induced to relocate by a premium increase resemble those that were already close to relocating for unrelated reasons, rather than households experiencing a contemporaneous shock to the property itself. This is consistent with the estimates in Table A.4, in which the relocation response is not attributable to households in areas experiencing disasters. It also implies that the estimated semi-elasticity is local: extrapolation to premium increases much larger than those observed requires the distribution of y rather than the estimated slope.

Proposition 6 (Insurer choice among movers)**.** *Let $\kappa_M^s \geq 0$ as in Assumption 1, suppose ε_I and ε_L are independent with log-concave densities and independent of m, maintain the benchmark normalization, and consider households for which $m > c^m$, which relocate for every realization of the shocks. Such a household changes insurer at the move if and only if $x > \kappa_M^s$, and*

$$d \;\longmapsto\; \Pr\big(x > \kappa_M^s \,\big|\, \Delta P_{c,t} = d\big) \tag{25}$$

is nondecreasing: among movers, those that experienced a larger premium increase before moving are more likely to change insurer.

Only the total increase is observed. A larger increase is associated with a larger wedge carried into the move, through the insurer shock directly, and through the pass-through differential, since an insurer that has repriced more is more likely to price above the lo-

[30] This treats the realized saving as a guide to the size of the assessed revision $\Delta\hat{y}$, which the model does not identify. The two are distinct objects: the household decides on $\mu(d, x)$ and realizes G. As an instrumental-variable estimate the figure is a complier average, so it is measured on the marginal households the argument concerns, but it remains an ex post realization rather than an ex ante revision. Nothing below turns on the exact figure, only on its being large relative to κ^s and a substantial fraction of c^m.

cal alternatives. This raises the probability that re-shopping at the destination yields an insurer preferable to the incumbent.[31]

IA.B.4 Financial constraints

The paper documents that the relocation response to premium increases is concentrated among households with higher credit scores, and the delinquency response among households with lower credit scores. A liquidity requirement is consistent with both patterns, although only the relocation margin is modeled here.

Remark 4 (Liquidity)**.** Suppose a household can finance relocation only if its liquid wealth w covers the outlay the move requires, $w \geq \kappa^m$, whereas switching, which represents search effort rather than a payment, is always feasible. Let $G_c \equiv \Pr(w < \kappa^m)$ denote the constrained share, and let w be independent of the states. Then $\Pr(M) = (1 - G_c)\Pr\big(y > y^*(x)\big)$, and three implications follow. Relocation is increasing in wealth. Wealth and the premium shock are complements in relocation, since a premium increase alters the behavior only of households able to respond: the response is scaled by $1 - G_c$, and the cross-partial with respect to wealth is $A_L \mathbb{E}_x[f_y(y^*(x))]\, g_w(\kappa^m - \mu_w) > 0$, where g_w is the density of wealth and μ_w a location shift in it. Conditional on $w < \kappa^m$ the relocation response is zero, and what such a household absorbs in full is the durable component of the increase, namely $A_L G$. Switching remains feasible and removes the wedge, which by (15) includes the transitory pass-through component $A_\Lambda \varepsilon_L$ of the location shock whenever $A_\Lambda > 0$, so the relief available to a constrained household is partial rather than absent.

The last clause is the point of contact with the delinquency evidence: a household that cannot finance the costly remedy and can obtain only partial relief from the inexpensive one bears the residual, and its other obligations absorb it. That residual is the part of the increase that no insurer serving the location will forgo, which is precisely what the constrained household cannot escape. The model does not include a delinquency decision, so this is an interpretation of that evidence rather than a prediction derived here.

We do not use credit score to identify a second friction. Search costs κ^s plausibly also vary with financial sophistication, and Ge et al. (2026c) report that switching responds more strongly to premium increases among households with higher credit scores. Since the same observable would then proxy for two distinct frictions, we treat that channel as an interpretation of the switching evidence rather than as part of the model.

IA.B.5 Proofs

Throughout, the comparative statics are within a location and conditional on the pass-through profile, as stated at the head of Section IA.B.3; x and y are independent with the densities assumed there, and the loadings (14) are constants. By (4) and (15), on the region where the wedge is positive a unit increase in ε_I displaces x by A_I; a unit increase in ε_L displaces x by A_Λ and displaces y by A_L. A location shift μ_G in G displaces y by A_L.

[31] For households whose relocation is itself induced by the premium increase there is an offsetting selection effect: conditional on the increase and on relocating, such a household is more likely to have experienced a large location component and hence a small insurer component. This attenuates the relationship without necessarily reversing it.

Proof of Proposition 1. The household maximizes over the three payoffs $\{0,\ x-\kappa^s,\ y+x-c^m\}$. If $x \le \kappa^s$, switching is dominated by remaining in place, so the household relocates if and only if $y+x-c^m>0$, that is, $y > c^m - x = c^m - \min\{x,\kappa^s\}$. If $x > \kappa^s$, switching dominates remaining in place, so the household relocates if and only if $y+x-c^m > x-\kappa^s$, that is, $y > c^m-\kappa^s = c^m - \min\{x,\kappa^s\}$; the term in x appears on both sides and cancels. On the relocation region with $x > \kappa^s$ we have $v(M) > v(S) > 0$, and on the segment with $x < \kappa^s$ we have $v(M) > 0 \ge v(S)$, which confirms the stated assignment.

Proof of Proposition 2. By independence and Proposition 1,

$$\Pr(M) = \int \big[1 - F_y\big(c^m - \min\{x,\kappa^s\}\big)\big] f_x(x)\,dx, \qquad \Pr(S) = \big[1-F_x(\kappa^s)\big]\,F_y\big(c^m-\kappa^s\big),$$

and the first may be written region by region as

$$\Pr(M) = \int_{x<\kappa^s} \big[1-F_y\big(c^m-x\big)\big] f_x(x)\,dx + \big[1-F_y\big(c^m-\kappa^s\big)\big]\big[1-F_x(\kappa^s)\big].$$

A shift of x by a gives $\partial \Pr(S)/\partial a = f_x(\kappa^s)F_y(c^m-\kappa^s)$ and $\partial \Pr(M)/\partial a = \int_{x<\kappa^s} f_y(c^m - x) f_x(x)\,dx$, the second because $\min\{x,\kappa^s\}$ is constant beyond the kink. A shift of y by b gives $\partial \Pr(S)/\partial b = -\big[1-F_x(\kappa^s)\big] f_y(c^m-\kappa^s)$ and $\partial \Pr(M)/\partial b = \mathbb{E}_x[f_y(y^*(x))]$. Part (i) follows by setting $a = A_I$, $b=0$; part (ii) by setting $a = A_\wedge$, $b = A_L$ and adding. Condition (19) is the rearrangement of $\partial \Pr(S)/\partial \varepsilon_L > 0$.

Proof of Remark 1. If ε_L is known, variation in $\lambda_{c,t}$ changes x through $A_\wedge$ and leaves G, and hence y, unchanged. The bound follows from the range of the frontier and does not require a derivative, which would in any case have to be integrated along the displacement as households pass through and beyond the sloped region. Let x and $x' \ge x$ denote the wedge states before and after any displacement, however large. By Proposition 1 a household relocates if and only if $y > y^*(\cdot) = c^m - \min\{\cdot,\kappa^s\}$, and y^* takes values in $[c^m-\kappa^s,\ c^m]$ for every argument. A household that does not relocate at x and does at x' therefore satisfies $y^*(x') < y \le y^*(x)$, hence $c^m-\kappa^s < y \le c^m$. The set of such households is contained in $\{y \in (c^m-\kappa^s, c^m]\}$, so $\Pr(M)$ increases by at most $\Pr\big(y \in (c^m-\kappa^s,\ c^m]\big)$, the mass the frontier sweeps as it falls from c^m to $c^m-\kappa^s$.

Proof of Proposition 3. By Remark 2 a household relocates if and only if

$$\hat{y} + \min\{x,\kappa^s\} \ > \ c^m, \qquad \hat{y} = m + A_L\big(\mu(d,x) + u_L\big).$$

Fix m, ε_I and u_L, and raise the incumbent's pass-through share from $\lambda_{c,t}$ to $\lambda'_{c,t}$. Since $\varepsilon_L \ge 0$, the observed increase rises from d to $d' = d + (\lambda'_{c,t}-\lambda_{c,t})\varepsilon_L \ge d$. The wedge index $A_\wedge\varepsilon_L + A_I\varepsilon_I$ is weakly larger at the higher share, since $A'_\wedge \ge A_\wedge$ by hypothesis and $\varepsilon_L \ge 0$, so $x' \ge x$ and therefore $\min\{x',\kappa^s\} \ge \min\{x,\kappa^s\}$. By Assumption 3 μ is nondecreasing in each argument, so $\mu(d',x') \ge \mu(d,x)$ and $\hat{y}$ is weakly larger as well.

For part (i), both terms on the left of the relocation condition are therefore weakly larger at $\lambda'_{c,t}$ for every realization of (m,ε_I). The set of households satisfying the condition weakly expands, and integrating over (m,ε_I) gives the stated ordering. No distributional assumption is used.

For part (ii), write $\theta \equiv c^m - \min\{x', \kappa^s\}$ for the frontier at the higher share. Every household with $\theta - \Delta\hat{y} < \hat{y} \leq \theta$ at the lower share satisfies the relocation condition at the higher one, since $\hat{y}$ rises by $\Delta\hat{y}$ as defined in (22), so the increase in the relocation probability is at least the mass of assessed gaps in that interval; the household did not relocate at the lower share, since the frontier there is at least θ. If μ does not vary with what the household observes, which is the full-information case of Remark 1, then $\Delta\hat{y} = 0$ and the only channel is $\min\{x, \kappa^s\}$, which by Proposition 1 cannot rise by more than κ^s in total because the frontier is flat for $x \geq \kappa^s$; the response is then at most (23). Since $\Delta\hat{y}$ is determined by the revision in the assessed location cost and not by κ^s, the response exceeds (23) whenever $\Delta\hat{y} > \kappa^s$ and the density of $\hat{y}$ is positive on the relevant interval.

Proof of Proposition 4. Since $\partial c^m / \partial \delta = -L^{mtg}$ with $L^{mtg} > 0$,

$$\frac{\partial \Pr(M)}{\partial \delta} = L^{mtg} \int f_y\big(c^m - \min\{x, \kappa^s\}\big) f_x(x)\, dx > 0.$$

Proof of Proposition 5. By (13), $\partial G / \partial u_L = 1$, and by (4), $\partial y / \partial u_L = A_L$. The displacement formula for y in the proof of Proposition 2 gives $A_L \mathbb{E}_x[f_y(y^*(x))]$ with $y^*(x) = c^m - \min\{x, \kappa^s\}$. For a discrete shift of size ΔG the affected mass is $\Pr\big(y \in (y^*(x) - A_L \Delta G,\ y^*(x)]\big)$, so membership requires that y lie within $A_L \Delta G$ of the frontier.

Proof of Proposition 6. For $m > c^m$, the benchmark normalization and $\varepsilon_L \geq 0$ give $y = m + A_L \varepsilon_L \geq m > c^m \geq c^m - \min\{x, \kappa^s\}$ for every realization, so relocation occurs irrespective of the shocks; because m is independent of $(\varepsilon_L, \varepsilon_I)$, restricting attention to these households carries no information about the shocks. Conditional on relocating, Assumption 1 implies that the household changes insurer if and only if the present value of the wedge exceeds the incremental cost, that is, $x > \kappa^s_M$, where x is given by (15) under the normalization $W_{t-1} = 0$. By (11) the observed increase is $d = \lambda_{c,t} \varepsilon_L + \varepsilon_I$, and $x = \max\{0, A_\wedge \varepsilon_L + A_I \varepsilon_I\}$ is nondecreasing in each of ε_L and ε_I, since $A_\wedge \geq 0$, $A_I > 0$, and $\max\{0, \cdot\}$ is nondecreasing. Hence the indicator of $\{x > \kappa^s_M\}$ is a nondecreasing function of the pair. Take $\lambda_{c,t} > 0$, the case $\lambda_{c,t} = 0$ being immediate because $d = \varepsilon_I$ then, and write d as the sum of the summands $\lambda_{c,t} \varepsilon_L$ and ε_I, which are independent with log-concave densities, scaling by a positive constant preserving log-concavity. By the monotonicity theorem for log-concave densities (Efron, 1965), the conditional expectation of any nondecreasing function of independent summands is nondecreasing in their observed sum. Applying it to that indicator gives the stated probability as nondecreasing in d.

Proof of Remark 4. Write $\Pr(M) = [1 - G_w(\kappa^m - \mu_w)] P^M$ with $P^M \equiv \Pr(y > y^*(x))$, which does not depend on μ_w, where G_w is the distribution function of wealth. Then $\partial \Pr(M) / \partial \mu_w = g_w(\kappa^m - \mu_w) P^M > 0$. From the proof of Proposition 2, $\partial P^M / \partial \mu_G = A_L \mathbb{E}_x[f_y(y^*(x))]$, so $\partial \Pr(M) / \partial \mu_G = [1 - G_w(\kappa^m - \mu_w)] A_L \mathbb{E}_x[f_y(y^*(x))]$; differentiating with respect to μ_w gives the stated cross-partial, the wealth parameter entering only the feasibility factor and the premium parameter only the preference factor. If $w < \kappa^m$, the choice set is $\{\varnothing, S\}$, and neither payoff depends on y or m.

Internet Appendix IA.C The Magnitude of the Delinquency Response Relative to the Literature

Section 7 compares our estimated delinquency response with five studies of mortgage payment shocks. Those studies report their effects in different units, over different horizons, for different definitions of delinquency, and at payment changes that differ from the $1,000 we use as a common dose, in both directions. This section records what each study reports, the arithmetic that places the estimates on a common footing, and the resulting comparison in both levels and proportions.

IA.C.1 What Each Study Reports

Table IA.4 collects the reported estimates, together with the hazard ratio at a common $1,000 dose that Section IA.C.2 defines. Four of the five require a step before their estimates can be entered in equations (26) to (28): a conversion to dollars for Fuster and Willen (2017) and Ganong and Noel (2020), an annualization for Abel and Fuster (2021) and Ganong and Noel (2020), and a derived baseline for Di Maggio et al. (2017).

Fuster and Willen (2017) report their effect per percentage point of interest rate rather than per dollar. Their sample is 30-year Alt-A hybrid ARMs with a ten-year interest-only period and fixed-rate periods of three, five, seven or ten years, and every figure we use is from the five-year loans at their first reset. The interest-only feature is what permits the conversion: the balance does not amortize over that period, so the annual payment is the interest rate multiplied by the balance. At the $272,000 mean origination amount of those loans (their Table 2, Panel B), the three-percentage-point reduction that almost two-thirds of them received is $8,160 of annual payment relief. Their Online Appendix F reports that, after loan age 63, about 17 percent of loans become delinquent without the reset and about 8 percent with it, for loans with a combined loan-to-value ratio between 130 and 140.

Di Maggio et al. (2017) report a $277 average reduction in the monthly payment, or $3,324 per year, and a 0.5 percentage point fall in the likelihood of 60-day delinquency after one year. They do not report a baseline delinquency rate directly. They do report that the two-year effect of 1.8 percentage points is about 40 percent of the mean, which implies a two-year baseline of 4.5 percent; treating the annual hazard as constant across the two years gives a one-year baseline of 2.28 percent.

Abel and Fuster (2021) report a monthly rather than an annual probability. Refinancing reduces the monthly probability of 90-day default by about 2.5 basis points against a takeup-weighted monthly rate of 5.9 basis points, and cuts the monthly payment by $173, or $2,076 per year. Compounding the monthly rates over twelve months gives a one-year baseline of 0.71 percent and a treated rate of 0.41 percent.

Ganong and Noel (2020) report their effect per percent of the payment rather than per dollar, and over a two-year rather than a one-year window. Their Table 1, Panel B, gives a mean monthly payment of $1,327, so one percent of the annual payment is $159. Their two-year control mean of 32.1 percent implies a one-year baseline of 17.6 percent under a constant annual hazard.

Wong (2024) is the one study for which no change of baseline or horizon is made. He reports that a $100 increase in annual property taxes raises the 30-day delinquency rate by 0.08 percentage points, which Table IA.4 states as 0.80 percentage points per $1,000, against a mean rate of 7.46 percent. His baseline is a point-in-time monthly rate and his effect is measured over that same month, so his row is not converted to a one-year probability as the other four are. Qualification 4 of Section IA.C.4 states what that leaves.

TABLE IA.4
REPORTED ESTIMATES IN THE PAYMENT-SHOCK LITERATURE

Study	Outcome	Horizon	Payment change ($/yr)	Baseline (%)	Effect at that dose	Hazard ratio φ_{1000}
Fuster and Willen (2017)	60-day	1 year	8,160	17.00	−9.0 pp	0.906
Di Maggio et al. (2017)	60-day	1 year	3,324	2.28	−0.5 pp	0.928
Wong (2024)	30-day	monthly	1,000	7.46	+0.80 pp	1.112
Abel and Fuster (2021)	90-day	1 year	2,076	0.71	−0.30 pp	0.767
Ganong and Noel (2020)	90-day	2 years	159	32.10	−0.38 pp	0.913
This paper	30-day	1 year	1,000	5.00	+0.70 pp	1.144
This paper	60-day	1 year	1,000	1.70	+0.38 pp	1.226
This paper	90-day	1 year	1,000	1.10	+0.30 pp	1.275

NOTES: *Table IA.4.* Baselines and effects are stated at each study's own horizon, so that every row supplies P_0, P_1 and D from a single window. The final column is φ_{1000} from equations (26) and (27) evaluated at those inputs. The one-year baselines at which equation (28) applies these hazard ratios are reported in Table IA.5; they differ from the baselines here only for Ganong and Noel (2020), whose two-year control mean of 32.1 percent annualizes to 17.6 percent. Fuster and Willen (2017): Online Appendix F for the 17 and 8 percent incidences, which their p. 186 restricts to five-year loans with a combined loan-to-value ratio between 130 and 140; Table 2, Panel B, p. 177 for the $272,000 mean origination amount, which is the five-year column; p. 175 for the sample composition, the ten-year interest-only feature, and the three-percentage-point reduction received by almost two-thirds of the five-year loans at first reset; p. 178 for the 60-day definition. Di Maggio et al. (2017): Table 4, Panel A, Column 1, p. 3572 for the monthly reduction in the first year after the reset, a coefficient of −276.57 that we round to $277 before annualizing and that their text on p. 3571 states as about $276; p. 3578 for the 0.5 and 1.8 percentage point effects; p. 3579 for the statement that the two-year effect is about 40 percent of the mean. Every figure is from their conforming five-year ARM sample. The $940 monthly reduction on p. 3564 is from their non-agency sample, for which no delinquency estimate is reported, and is not used here. The 2.28 percent baseline is derived, not reported. Wong (2024): p. 26 for the 0.08 percentage point effect per $100; Figure A5, Panel E for the 7.46 percent mean; p. 15 for the 30-day definition; p. 13 for the restriction to escrow payers. Abel and Fuster (2021): p. 265 for the $173 monthly reduction; p. 276 for the 2.5 basis point effect; Table 3 notes, p. 277 for the 5.9 basis point mean; p. 275 for the 90-day definition. Ganong and Noel (2020): p. 3127 for the 0.38 percentage point effect per percent of payment; p. 3126 for the 32.1 percent control mean; Table 1, Panel B, p. 3113 for the $1,327 mean monthly payment; p. 3125 for the 90-day, two-year definition. Our own estimates are the Far-ZIP second stage of Table 3, Columns (1) and (2), and Table 2, Column (5), each multiplied by 0.2, the premium rate increase corresponding to $1,000 at the $500,000 coverage limit used throughout the paper.

IA.C.2 Placing the Estimates on a Common Footing

Two adjustments are required. The first rescales each estimate from the dose at which it was measured to $1,000. The second expresses every estimate at a one-year horizon.

We treat each reported effect as multiplicative on the default hazard, which is the form in which Fuster and Willen (2017), Abel and Fuster (2021) and Ganong and Noel (2020) estimate their models. Let P_0 and P_1 denote the baseline and treated probabilities over a window, and let D denote the annual payment change. The implied hazard ratio for the full dose is

$$\varphi \;=\; \frac{\ln(1-P_1)}{\ln(1-P_0)}, \tag{26}$$

which does not depend on the length of the window, so φ may be computed over the window a study reports and applied at a one-year baseline, as it is for Ganong and Noel (2020). Because a proportional effect compounds in the dose, the hazard ratio for \$1,000 is

$$\varphi_{1000} \;=\; \varphi^{\,1000/D}, \tag{27}$$

not φ divided by $D/1000$. The level effect at a one-year horizon then follows from applying φ_{1000} at each study's own one-year baseline,

$$\Delta P \;=\; \left|P_0 - \left[1-(1-P_0)^{\varphi_{1000}}\right]\right|, \tag{28}$$

and the proportional effect is $\Delta P/P_0$. Equation (27) matters most where the reported dose is furthest from \$1,000, which is Fuster and Willen (2017) at \$8,160 and Ganong and Noel (2020) at \$159, and it corrects in opposite directions for the two. Scaling linearly would understate the effect of the first \$1,000 for Fuster and Willen (2017), because the default probability has already fallen substantially over the remainder of the dose, and would overstate it for Ganong and Noel (2020), whose coefficient is stated per \$159 and must be compounded over more than six of them.

One row departs from equation (27). Wong (2024) reports his effect per \$100, and Table IA.4 multiplies it by ten rather than compounding it, because his design estimates a linear dose-response and the quadratic specification he reports cannot reject linearity. Compounding instead would raise his hazard ratio from 1.112 to 1.117, his level effect from 0.80 to 0.84 percentage points, and his proportional effect from 10.7 to 11.2 percent, and would lower the final column of Table IA.5 from 1.31 to 1.25. We report the linear figure because it is the one his estimates imply under his own functional form.

Neither adjustment is required for our own estimates: our regressor is linear in the premium change, and our outcomes are already measured over one year. The multiplicative treatment also matters little for them, because our outcome is rare and equation (26) is therefore close to the ratio of probabilities: our 60-day estimate implies a hazard ratio of 1.226, against a ratio of probabilities of 1.224.

IA.C.3 Comparison

Table IA.5 reports the result. Each study is paired with the estimate of ours that shares its definition of delinquency, so that no comparison relies on a mapping between 30-, 60- and 90-day outcomes.

The two metrics rank the studies differently, and the difference is informative rather than a matter of presentation. In levels our estimate is in the middle of the group: smaller than Fuster and Willen (2017), Wong (2024) and Ganong and Noel (2020), larger than Di

TABLE IA.5
COMPARISON AT A COMMON DOSE OF $1,000 AND A ONE-YEAR HORIZON

Study	Outcome	Their estimate			Our estimate		Ratio of
		Baseline (%)	Level (pp)	% of base	Level (pp)	% of base	proportional effects
Fuster and Willen (2017)	60-day	17.00	1.46	8.6	0.38	22.4	2.60
Di Maggio et al. (2017)	60-day	2.28	0.16	7.2	0.38	22.4	3.12
Wong (2024)†	30-day	7.46	0.80	10.7	0.70	14.0	1.31
Abel and Fuster (2021)	90-day	0.71	0.16	23.3	0.30	27.3	1.17
Ganong and Noel (2020)	90-day	17.60	1.40	8.0	0.30	27.3	3.42

NOTES: *Table IA.5.* Level effects are the change in the probability of the stated delinquency outcome caused by a $1,000 change in the annual payment, computed from equations (26) to (28) using the inputs of Table IA.4. Proportional effects divide that level effect by the baseline in the third column, and are computed before rounding, so they need not equal the ratio of the two rounded figures shown. Baselines are one-year rates in every row except the one marked †. The 7.46 percent for Wong (2024) is a point-in-time monthly rate entered without conversion, and the level effect in his row is measured over that same monthly window, so his row compares proportional effects rather than one-year probabilities. The final column divides our proportional effect by theirs. Each study is reported in the direction it estimates: payment relief for all but Wong (2024), whose tax increases share the direction of our premium increases.

Maggio et al. (2017) and Abel and Fuster (2021). In proportional terms it is the larger in every matched pairing. It exceeds Abel and Fuster (2021) by 17 percent, a margin the first qualification below shows the direction of the shock can reverse, and it exceeds the other four by between 30 percent and a factor of 2.4.

The reason is the spread in baseline delinquency rates, which runs from 0.71 percent in Abel and Fuster (2021) to 17.6 percent in Ganong and Noel (2020). The two studies with the largest level effects, Fuster and Willen (2017) and Ganong and Noel (2020), are those with the most distressed samples: loosely documented mortgages originated at the 2005 to 2006 peak and observed through the foreclosure crisis in the first case, and borrowers applying for crisis-era modifications in the second. Both have one-year baselines near 17 percent, an order of magnitude above ours, and in the case of Fuster and Willen (2017) that figure describes loans with a combined loan-to-value ratio between 130 and 140 rather than their sample as a whole. A payment shock that moves the default hazard by a given proportion moves the level of default far more in those samples than in a performing 2014 to 2023 sample. The proportional comparison is far less sensitive to that difference than the level comparison, but it does not remove it. What a proportional hazard model estimates is the hazard ratio itself, and $\Delta P/P_0$ is a transformation of that ratio evaluated at a particular baseline and window, so it retains a dependence on both. The same φ_{1000} for Ganong and Noel (2020) implies 8.0 percent of their one-year baseline and 7.3 percent of their two-year baseline. Because $\Delta P/P_0$ falls further below $\varphi_{1000} - 1$ the larger is P_0, the final column of Table IA.5 overstates the gap for the studies with the most distressed samples. Table IA.4 therefore reports φ_{1000} for every study. Ranking on those hazard ratios instead leaves the ordering unchanged and compresses the final column, whose largest entry falls from 3.42 to 3.15 and whose smallest rises from 1.17 to 1.18. Only three

of the five studies estimate their effect in multiplicative form in any case, as Section IA.C.2 notes; for the other two the directly estimated object is the level effect.

IA.C.4 Qualifications

Four qualifications apply, in descending order of importance.

First, four of the five studies estimate the effect of payment reductions, while we estimate the effect of increases. Table IA.5 reports each in its own direction and does not assume symmetry. Imposing symmetry, so that a \$1,000 increase is assigned the inverse of the hazard ratio estimated for a \$1,000 reduction, raises the proportional effect by about 8 percent for Fuster and Willen (2017), Di Maggio et al. (2017) and Ganong and Noel (2020), whose hazard ratios lie between 0.906 and 0.928, and by 30 percent for Abel and Fuster (2021), whose 0.767 sits much further below one, so that the distance between φ_{1000} and its inverse is correspondingly wider. The ordering is unchanged in every case except Abel and Fuster (2021), whose proportional effect rises to 30.3 percent and slightly exceeds our 27.3 percent. Their estimate is therefore best described as comparable to ours rather than larger or smaller.

Second, Abel and Fuster (2021) estimate the effect of a binary refinancing indicator rather than a dose-response, so no functional form in dollars is identified by their design. Equation (27) imposes one. A linear rescaling would place their level effect at 0.14 rather than 0.16 percentage points, and their proportional effect at 20.4 rather than 23.3 percent. Their row should be read as a range rather than a point.

Third, the linearity that equation (27) assumes is tested only in some of these settings. Fuster and Willen (2017) report hazard ratios at reductions of two, three and four percentage points that imply per-\$1,000 ratios of 0.910, 0.907 and 0.904, so the log hazard is close to linear in dollars across their range. Wong (2024) adds quadratic terms and cannot reject linearity. Neither study identifies the response at \$1,000, and the two extrapolate in opposite directions. The reductions Fuster and Willen (2017) examine are several times \$1,000, so equation (27) extrapolates below the range in which they estimate. Wong (2024) reports his effect per \$100 against a mean increase of \$171 per year, so his row applies a slope estimated over much smaller increases at a dose several times that mean. The mean does not establish how much of his sample lies near \$1,000, and his row therefore extrapolates upward, resting on the linearity his quadratic specification cannot reject. Linearity is not tested at all in the three remaining settings, and the extrapolation is longest for Ganong and Noel (2020), whose coefficient is stated per \$159 and is therefore compounded over more than six units to reach \$1,000. Di Maggio et al. (2017) is scaled in the opposite direction, from \$3,324. Because Ganong and Noel (2020) supply the largest entry in the final column of Table IA.5, the comparison that depends most on equation (27) is also the one that applies it over the longest range.

Fourth, the outcomes are not identical even within a matched pair. Wong (2024) measures an indicator of being 30 or more days past due in a given month, a point-in-time rate, whereas our 30-day outcome is a first transition into delinquency within twelve months among mortgages current at renewal. His sample is also restricted to escrow payers, which selects riskier borrowers and contributes to his higher baseline. Ganong and Noel (2020) measure default among borrowers already applying for a modification,

and Abel and Fuster (2021) recover a local average treatment effect for borrowers who actively chose to refinance, rather than an average effect over the population facing the shock.