Economics of Property Insurance

Hyeyoon Jung* Kyle (Jaehoon) Jung†

August 11, 2026

Abstract

We study the economics of homeowners' property insurance by examining how contract design balances the trade-off between incentive alignment and risk sharing. Using granular *contract-level* property insurance data merged with property-level disaster risk for millions of U.S. households, we develop and structurally estimate a model in which insurers optimally determine contract terms given property risk and household risk preferences. The estimates provide, to our knowledge, the first large-scale contract-level structural measures of risk aversion, risk premia, and the cost of moral hazard, allowing us to quantify how disaster risk is allocated between insurers and households. We find that the cost of moral hazard is small, yet the very mechanism used to mitigate it substantially increases households' exposure to disaster risk: contract design leaves policyholders exposed to roughly 29 percent of total expected losses. This residual exposure is most pronounced for low-FICO households and for properties with the greatest tail risk. Counterfactuals indicate that mandating full insurance would lead to substantial market exit, increasing household vulnerability. We further show that insurers' financial constraints are systematically correlated with the riskiness of underwritten properties and with household characteristics.

*Federal Reserve Bank of New York. (hyeyoon.jung@ny.frb.org)

†New York University, Stern School of Business. (jj1419@stern.nyu.edu)

We are grateful for the recognition of this work with the Best Paper Award at the Mid-Atlantic Research Conference in Finance. We thank Neil Kapoor and Katharina Dowlin for their excellent research assistance and Dennis Kongoli and Henry Dyer for invaluable technical help. We also thank Motohiro Yogo, Sasha Indarte, Quirin Fleckenstein, Ai He, Chotibhak (Pab) Jotikasthira, Parinitha Sastry, Ana-Maria Tenekedjieva, Fabrice Tourre, Mallick Hossain, Robert Engle, Shan Ge, and participants of NBER Insurance meeting, SFS Cavalcade, Finance Theory Group conference, Beachside Banking seminar, the Federal Reserve Bank of New York seminar, NYU Stern Finance QFE seminar, Seoul National University seminar, Korea University seminar, Mid-Atlantic research conference, Baruch-NUS Real Estate conference, Risk Theory Society seminar, Federal Reserve Bank of Philadelphia conference on consumer finance and macroeconomics in finance, Bocconi-CEPR finance conference, ESADE Spring Workshop, Pre-WFA Early Career Women in Finance conference for helpful comments. The views expressed in this paper are those of the author(s) and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. All errors are our own.

1 Introduction

Housing is the largest component of household wealth in the US, totaling more than $45 trillion.[1] This suggests that, from households' perspective, property risk is primary, not incidental. Losses from natural disasters are rare but lumpy. Realized damages typically far exceed the liquid savings of many households.[2]

Insurance is the primary means of hedging these losses. Homeowners' property insurance is one of the most widely held financial contracts in the household portfolio, comparable in prevalence to checking accounts. US households paid about $150 billion per year in homeowners insurance premiums in 2023.[3]

However, what households actually receive in return for these premiums depends crucially on the structure of the insurance contract. Property insurance in practice is not a simple promise to cover the entire loss. Instead, contracts include a deductible, which is the amount the household must pay before the insurer compensates anything, and a coverage limit, which caps the insurer's total payout. These terms determine how much of a disaster loss is ultimately borne by the household. Since many households have limited liquid savings, the extent to which a disaster is financially smoothed through insurance versus passed through to the households' balance sheet is a critical matter for household financial stability. Contract design therefore plays a central role in the financial consequences of disasters, both at the household level, in shaping risk exposure and recovery, and at the macro level, by influencing how local shocks transmit into housing markets, credit outcomes, and aggregate economic activity.

These contract features arise from a fundamental risk-sharing versus incentive trade-off. In a frictionless setting, insurers, who are diversified and thus effectively risk-neutral, would provide full insurance to risk-averse households against property losses. In practice,

[1] Financial Accounts of the United States (Board of Governers of the Federal Reserve System).

[2] Paid claim amounts, averages about $19,000 in low-risk areas and $24,000 in high-risk areas (Federal Insurance Office, 2025). By contrast, Federal Reserve surveys indicate that 48% of US adults report being unable to cover unexpected expense of $2,000 or more using liquid savings (Federal Reserve Board, 2025).

[3] Woleben (2025), S&P Global Market Intelligence.

however, insurers cannot observe the effort households take to maintain or protect the property once insured. This classic moral hazard problem can be mitigated by exposing policyholders to residual risk through deductibles and coverage limits.

In this paper, we study three key questions about the economics of homeowners' property insurance contracts. First, how are insurance contracts designed in practice—specifically, how do premiums, deductibles, and coverage limits relate to underlying property risk and insurer pricing? Second, how large is moral hazard in homeowners insurance, and to what extent do cost-sharing terms mitigate it? Third, how much risk is ultimately retained by households as a result of contract design, and how does this retained risk vary across households and insurers?

To our knowledge, this paper provides the first contract-level structural estimates of policyholders' risk preferences and moral hazard in property insurance, allowing us to construct contract-specific counterfactuals and characterize how incentive and risk-sharing primitives vary across households, properties, and insurers. We use this framework to show that insurance contract design plays a central role in allocating disaster risk between insurers and households. We then document that the resulting residual risk has meaningful financial consequences, with liquidity-constrained households retaining the greatest exposure precisely in high-risk areas. Finally, we show that insurer financial constraints are systematically correlated with the riskiness of underwritten properties and with the characteristics of households selecting into these contracts, highlighting the importance of accounting for endogenous selection when analyzing contract design and risk allocation.

The research questions we examine are important because contract design determines how disaster losses are allocated between insurers and households, with implications for household financial resilience and for how shocks propagate more broadly. However, surprisingly little is known about property insurance contracts, both about facts and mechanisms. A primary challenge is data availability: granular data on contract terms,

premiums, and underlying property-level risk have been rarely observed. Second, even with such a dataset, reduced-form approaches face identification challenges, as premiums, deductibles, and coverage limits are joint equilibrium outcomes shaped by both household preferences and insurer pricing and risk management.

We address the first challenge above by exploiting granular *contract-level* insurance data covering premiums and contract terms, combined with *property-level* disaster risk and recovery value, as well as characteristics of policyholders and insurers. The fully merged dataset covers about 8.7 million contracts per year, allowing us to exploit the rich cross-sections. The second challenge is addressed by writing a structural model of insurance contracting. The model provides a connection between the observable contract features, such as deductible and coverage, and the unobservable primitives of policyholders, such as their risk preference and the effect of their negligence on property risk. The model not only allows for the estimation of these primitives, but also counterfactual analysis to examine potential effects of regulatory interventions, which are prevalent in the insurance market.

Our main finding is that while insurance contracts effectively mitigate moral hazard, they do so by shifting a substantial amount of risk back to households. To quantify the tradeoff, we develop and structurally estimate a model in which insurers optimally design contract terms given the property risk and household risk preference. The estimates yield *contract-level* measures of risk aversion, risk premium, cost of moral hazard, and the resulting increase in uninsured exposure. We find that the cost of moral hazard is quantitatively small, about 0.7% of the risk premium, indicating the ease of incentivizing policyholder's effort through risk exposure. However, the very mechanisms that mitigate moral hazard lead policyholders to retain considerable risk: contract design leaves policyholders exposed to roughly 29% of total expected losses. This residual risk is large especially for properties subject to severe tail risk. In the cross-section, these effects are most pronounced among low-FICO households. Among insurers, more constrained insurers underwrite

riskier properties and therefore charge higher premiums, but these higher premiums reflect higher expected losses rather than higher risk premia. Taken together, the results suggest that moral-hazard mitigation comes at a distributional cost; financially constrained households retain the most risk precisely where disaster exposure is greatest.

Using our structural model and the estimated parameters, we analyze the counterfactual in which insurers are required to provide full coverage. Under this mandate, a large share, 46%, of policyholders would lose access to insurance altogether, implying that removing insurers' ability to design incentive-compatible contracts could trigger market failure and ultimately expose households to even greater risk.

We make use of state of the art *contract-level* data, from Intercontinental Exchange ("ICE") McDash, covering the details of homeowners' property insurance and rich characteristics of the insurance policyholder. We merge this dataset with property-level disaster risk metrics from CoreLogic. The merge is done at the loan-level, meaning that, for each insurance policyholder, we observe the contract terms of insurance policy, including premium, coverage limit, deductible, insurer identity, and the policyholder's property characteristics, including location, expected damage, and tail damage values. We further merge this dataset with per-square-foot-structural value estimates provided by the National Structure Inventory at the zip code level. For our main cross-sectional analysis, we focus on the year when the structural value estimates are readily available at the most granular level.

For this one year, we observe almost 8.7 million contracts spanning all states of the US, including over 200 insurance companies. This granular dataset allows us to estimate key parameters, such as policyholder-level risk aversion, property-level damage distribution, contract-level risk premium, and cost of moral hazard. The rich cross-section of policyholders, properties, and insurance companies, in turn, allows us to make progress on understanding mechanisms through cross-sectional analyses.

Based on the constructed dataset, we begin by documenting key stylized facts about the US property insurance contracts. We document that coverage limits are rarely binding. The median ratio of coverage to property recovery value far exceeds the 99th-percentile loss rate, meaning that insurers would pay the same amount in nearly all realized loss scenarios even if coverage were higher. Thus, coverage plays a limited role in limiting insurer payouts, implying that deductibles, rather than coverage, are likely the primary mechanism used to mitigate moral hazard. Second, deductibles are small relative to property value, but large relative to expected loss. This suggests that, because a large share of losses occur in the region where the deductible applies, the deductible plays a particularly important role in determining household risk retention and mitigating moral hazard. Third, premiums substantially exceed expected losses. This is consistent with equilibrium in a market with risk-averse policyholders: households are willing to pay more than the expected claim payout to transfer disaster risk off their balance sheets, generating a risk premium. Fourth, damage risk is low on average but highly skewed. These patterns are robust to alternative constructions of insured exposure, reflecting that flood inundation is typically excluded while non-hazard claims such as liability and theft are covered under standard homeowners policies.

To analyze these endogenous contract terms, we develop a model in which premiums, deductibles, and coverage limits are designed optimally by insurers given the households' risks and preferences. The model links the distribution of damage risk and household risk aversion to the observed structure of insurance contracts. This framework allows us to quantify how much risk is transferred to insurers versus retained by households, how much risk premium is forgone due to moral hazard, and how these components vary across households, insurers, and regions.

We model the insurance design problem based on the canonical moral hazard model by Holmström (1979). The policyholder, the agent, can make an unobservable action choice that determines the risk to the property. Specifically, the policyholder can make a personally

costly effort to reduce the risk to the property. Given the information asymmetry regarding the agent's hidden action, the insurer, the principal, faces the tradeoff between risk-sharing and incentives. On one hand, the risk-neutral insurer profits from taking on the risk-averse policyholder's risk. On the other hand, the insurer also gains from incentivizing the policyholder to properly manage their property (effort). This tradeoff endogenously gives rise to insurance contracts that offer only partial coverage; deductibles and coverage limits limit the range of protection that the insurer provides to the policyholder. By exposing the policyholder to residual risk, the contract can incentivize the policyholder to exert effort to reduce their own risk. The region of protection for the insurer, defined below by the deductible and above by the coverage, is determined by the informativeness of the realized damage regarding the policyholder's hidden action. In particular, zero or low damage would be likely under effort and very high damage would be likely under negligence (no effort). As a result, the incentive regions are concentrated in the tails, which is consistent with the contract design we observe in practice: the policyholder is exposed to risk for damages below the deductible and above the coverage, while enjoying full protection for damages in between.

The model allows us to uncover the unobservable primitives of policyholders' risk preferences and hidden actions. Specifically, the key objects we estimate are: the policyholder's risk aversion, the cost of policyholder's unobservable effort, and the counterfactual outcome distribution had the policyholder chosen negligence. First, the risk aversion is identified by the participation condition: the policyholder should prefer to be insured. Second, the cost of effort is identified by the incentive compatibility condition: if the policyholder participates in the insurance contract, she should prefer exerting effort to negligence. Third, the counterfactual damage distribution under negligence comes from the insight in Holmström (1979) that the optimal payoff is a function of the likelihood ratio, which is how likely an outcome is under effort vs. negligence. The counterfactual

distribution can therefore be backed out from the payoff function given the distribution under effort that we observe in equilibrium.

For the estimation, we use the property's risk characteristics, such as the expected loss rate and the recovery value, as well as contract characteristics, including the deductible and the coverage. With these data as inputs, we estimate the model in steps. First, we estimate the loss rate distribution from the property risk characteristics. Second, we decompose the premium into the expected claim payout and the risk premium. Third, we estimate the risk aversion from the participation condition. Fourth, we estimate the cost of moral hazard by constructing the counterfactual optimal contract in the absence of moral hazard. Fifth, we estimate the cost of the policyholder's effort, using moment conditions including the incentive compatibility condition. Sixth and finally, we recover the counterfactual loss rate distribution had the policyholder neglected her property.

The results of the estimation can be summarized as the following. We estimate contract-level risk aversion ranging from 15 to 170, with a mean of 79. Our estimates are moderate relative to those of prior literature on other types of insurance. Unlike pooled estimates in the literature, our approach yields policyholder-level estimates, allowing us to characterize rich heterogeneity across contracts. Consistent with economic intuition, we find that more risk-averse policyholders tend to own safer properties, lending validity to our estimates. On average, the estimated risk premium accounts for 73% of the total premium, and it increases with both risk aversion and the underlying level of risk, confirming the internal consistency of our model. The estimated cost of moral hazard is low, only about $7, or 0.7% of the risk premium, suggesting that it is cost-efficient to incentivize due care by policyholders. However, this comes at the expense of higher uninsured exposure: deductibles and coverage limits that reduce moral hazard expose policyholders to 29% of total expected losses. As expected, the exposure due to moral hazard is larger for contracts covering properties with greater tail risk, since deductibles and coverage caps truncate both tails of the loss distribution.

In the cross-section, we find that low-FICO policyholders are more risk averse, consistent with the financial constraint channel, and face higher costs of moral hazard, reflecting greater uncertainty about their unobservable actions. They also pay higher risk premiums and experience larger increases in uninsured exposure due to contract features designed to limit moral hazard, underscoring the distributional implications of these frictions. On the insurer side, we find that more financially constrained insurers, proxied by lower risk-based capital ratios, tend to underwrite riskier properties and charge higher risk premiums, although not higher risk premiums relative to total premiums. Together, these indicate that financial constraints shape both sides of the market, with constrained households bearing more residual risk and constrained insurers taking on riskier exposures.

To evaluate the economic validity of our cost of moral hazard estimates, we examine whether they behave consistently with economic intuition in two cross-sectional settings. First, we relate the estimated cost of moral hazard to the loan-to-value (LTV) ratio at origination. Because a lower LTV implies greater homeowner equity and hence stronger incentives to maintain the property and avoid excessive damages, we expect moral hazard to decline as ownership increases. Importantly, our moral hazard estimates are derived solely from insurance contract characteristics—premiums, deductibles, and coverage—along with the underlying risk distribution, without incorporating borrower or financing variables such as credit scores or mortgage terms. The relationship with LTV therefore provides an out-of-sample validation of the measure. Restricting the sample to loans originated within one year, we find that moral hazard indeed increases significantly with LTV. As a second validation, we exploit cross-state variation in insurance regulations, such as requiring pre-binding property condition verification (e.g., roof age, condition, or wind mitigation features). We expect such regulations to mitigate information asymmetry and thereby reduce moral hazard. Consistent with this hypothesis, we find that states with inspection requirements exhibit significantly lower estimated costs of moral hazard.

With the parameters and the counterfactual distribution, we examine the potential effect of mandating insurers to provide complete insurance. Under this hypothetical regulation, we find that 46% of policyholders will lose insurance, exposing them fully to property risk. The result suggests that taking away the insurers' means of incentivizing the policyholders' effort to manage their properties can result in market failure, and thus greater risk to households.

Contribution to literature. This paper contributes to several strands of literature. First, this paper contributes to the studies identifying and estimating agency frictions with a structural model. In the context of executive compensation, works such as Margiotta and Miller (2000), Gayle and Miller (2009), Gayle and Miller (2015), Gayle et al. (2022), Bertomeu et al. (2025b), and Jung (2025b) use structural model building on Holmström (1979) to identify and estimate moral hazard. In the insurance literature, Einav et al. (2010) estimate the demand for insurance and the welfare implications of adverse selection. Einav et al. (2013) estimate the impact of moral hazard on the policyholder's selection of medical insurance. We contribute to this literature by providing, to our knowledge, the first large-scale contract-level structural estimates of policyholders' risk preferences and moral hazard in property insurance. Our framework also constructs the counterfactuals for each insurance contract, enabling rich characterization of how risk aversion and incentive effects vary across households, properties, and insurers.

Second, we contribute to the growing literature on climate finance and insurance. Keys and Mulder (2024) link rising premiums to climate risk; Jung et al. (2025) measure property insurers' exposure to physical climate risk; Blonz et al. (2024) study the role of policyholders' credit risk in insurance pricing; and Sastry et al. (2025) analyze property insurance demand using an IO framework. Cookson et al. (2024) document the demand side friction in the property insurance market. We complement this work by showing that insurance contract design is central to how climate risk is allocated between insurers and households.

While contract terms effectively curb moral hazard, they systematically determine how much of the increasing tail risk associated with climate exposure is transferred back to policyholders.

Third, this paper contributes to the real effects of insurance and hedging. On households, Ge et al. (2024) show that rising insurance premiums increase mortgage delinquency, and Jotikasthira et al. (2025) document delays in claim payouts after insurers experiencing adverse events. On firms, Aunon-Nerin and Ehling (2008), Perez-Gonzalez and Yun (2010), and Jung (2025a) document real effects of insurance and derivative hedging. We add to this literature by showing that the contract features that mitigate moral hazard also generate significant household exposure to disaster risk, and that financially constrained households retain the most risk precisely in regions where disaster exposure is greatest. This highlights a previously underexplored channel through which insurance contract design affects household financial resilience.

Fourth, we contribute to the literature on financial constraints faced by intermediaries. Koijen and Yogo (2014) and Koijen and Yogo (2016) document the role of regulatory frictions in pricing by life insurers; Ge (2021) shows how insurers' capital constraints affect premiums; and Oh et al. (2025) examine regulatory frictions in P&C pricing. We extend this line of work by documenting that insurer financial constraints are systematically correlated with the riskiness of the properties they underwrite and with the characteristics of the households that select into these contracts, including their risk preferences and financial constraints. These patterns highlight the need to account for endogenous selection when analyzing insurance contract design and risk allocation.

Outline of the Paper. The rest of the paper is organized as the following. In Section 2, we describe the data sources and how we construct the sample dataset. Section 3 explains the main components of property insurance contracts and documents stylized facts about the key contract terms. Section 4 covers the model and Section 5 describes

the structural estimation approach. Section 6 studies the estimated results and Section 7 analyzes counterfactuals and discusses policy implications. Section 8 concludes.

2 Data

We use a number of large datasets on insurance policy, mortgage loans, property-, borrower-, and insurer characteristics from the following data sources.

Property Insurance Data. ICE McDash insurance module provides data on the US homeowners' property insurance contracts. The data contains detailed characteristics of the contract, including the premium, the deductible, the coverage limit, and the identity of the insurer. Since this database comes from residential mortgage servicers, McDash also provides a mortgage loan module, which provides data on the mortgage loan associated with each insurance policy.[4] Therefore, we observe rich information on borrowers (e.g., credit score), mortgage contract details (e.g., loan amount, origination date), and property characteristics (e.g., zip code, appraisal value). The dataset covers approximately two-thirds of installment-type loans in the residential mortgage servicing market.

Disaster Risk Data. We complement the insurance data with property-level climate risk metrics from CoreLogic. CoreLogic constructs peril-specific disaster risk measures (e.g., earthquake, wildfire, inland flood, severe convective storm, winter storm, hurricane storm surge, and hurricane wind) using proprietary models that combine location-specific hazard intensities, such as water depth and velocity, with structure characteristics including foundation and construction type. Simulated hazard realizations generate a distribution

[4]Mortgage servicers are responsible for monitoring insurance coverage on the properties securing mortgages. Furthermore, they are required to enforce various rules related to insurance coverage. If there are lapses in insurance coverage the servicer is required to force-place insurance on the property. If insurance coverage is not maintained, the servicer is liable for any damage that may result (e.g., disaster strikes while the property is uninsured). Because of these legal responsibilities and liability risks, servicers maintain detailed data on insurance.

of potential losses, from which the average annual loss (AAL) and aggregate exceedance-probability loss (AEP) are reported as shares of recovery value. AEP loss rate values are provided at 1% and 2% exceedance probabilities; AEP 1% loss rate and AEP 2% loss rate correspond to the loss levels that have a 1% and 2% chance of being exceeded in a given year, respectively.[5] CoreLogic also provides real estate deed and property modules, allowing the disaster risk metrics to be merged with mortgage loan information and property characteristics at the property level. This dataset covers approximately 192 million properties.

Among the nine perils, we exclude earthquake-related measures because standard homeowners policies typically exclude coverage for losses caused by earth movement. While homeowners policies generally exclude losses caused by surface-water inundation (covered by FEMA's National Flood Insurance Program), the exclusion applies to inundation as the cause of loss rather than to all damage occurring under flooding conditions. Flood conditions may still generate covered losses, for example, through structural failure, debris impact, or water intrusion classified under covered causes. Therefore, we retain flood in the risk metric for the baseline analysis and consider excluding it as a robustness check.

Structure Value Data. AAL measure from CoreLogic is computed as a loss rate in terms of recovery value, and therefore computing the expected loss in dollars requires data on recovery value. As this is not available from either McDash or CoreLogic, we impute it from the structural value estimates provided by the National Structure Inventory (NSI). NSI base layer was created and is maintained by the US Army Corps of Engineers (USACE), and the data has been used in various applications by USACE, FEMA and other agencies. The data provides structure-level characteristics, such as type (e.g., residential or commercial),

[5] As a validation exercise, we compare CoreLogic's hurricane wind risk estimates with zip-level estimates from First Street. We focus on hurricane wind risk because, unlike flood risk, it is less sensitive to within-zip variation in elevation, floodplain coverage, and other highly localized geographic features. The correlation is high (0.82). Comparable measures for other hazards are not available. This cross-vendor comparison supports the reliability of the CoreLogic measure.

structure value, content value, and year of construction. We use these variables to compute the zip-level recovery value. Since this particular dataset is most readily available for the year of 2021 (released in 2022), we focus on that year. We confirm that the patterns of contract characteristics are consistent with those of other years.

Insurer Data. We use insurer regulatory filings collected by the National Association of Insurance Commissioners (NAIC) to obtain insurers' financial information.

To construct a panel of insurance policy characteristics and expected losses, we take the following steps. First, we merge the insurance data with the climate risk data by key mortgage characteristics: loan amount, origination date, and zip code. We rarely find multiple loans with the same amount originated on the same day within a zip code. By using the three variables, we are able to match about 70% of the insurance dataset. For those observations that are not merged, we assign the average climate risk metric of that zipcode originated in that year for that loan size bin.

Second, we merge the insurance-climate risk-merged dataset with the structural value data at the zip-year level. For this, we converted each longitude-latitude combination to a corresponding zip code. Since the NSI structural value variable is adjusted for depreciation, we undo this adjustment for it based on the formula provided by the NSI. We then compute a per-square-feet recovery value by dividing the gross recovery value by the structure size in square feet provided by the NSI, and obtain a mean value for each zip code. We multiply this zip-level per-square-feet recovery value by the property size (in square feet) in CoreLogic side.[6]

Third, we merge the above with insurer financial information at the insurer-year level. For the main cross-sectional analysis, we focus on the year when the structural value is readily available. Since the estimation requires all insurance characteristics (premium,

[6]One may be concerned about differential selection bias between the NSI dataset and the CoreLogic. We compare the zip-code level statistics of the two datasets and confirm that the mean ratio between them is 0.98, alleviating the concern.

coverage, and deductible) as well as loss distribution parameters (annual average loss and aggregate exceedance probabilities), we keep contracts with all those variables non-missing. Moreover, we focus on non-expired contracts on non-condo properties because condos typically have building-level policies purchased by the condo association. We focus on escrowed accounts because they are mainly monitored by the data provider, mortgage servicers. This leaves us with almost 8.7 million contracts. Table 1 shows summary statistics of the key contract-level characteristics.

	Mean	St.Dev.	P10	Median	P90	N
Premium (Annual)	1,515	816	720	1,308	2,604	8,699,672
Coverage	340,709	174,743	177,900	300,676	546,000	8,699,672
Deductible	1,521	1,006	736	1,000	2,500	8,699,672
AAL	.0011	.00095	.00016	.00092	.0022	8,699,672
AEP 2% Loss	.0099	.011	.00065	.008	.02	8,699,672
AEP 1% Loss	.019	.022	.0011	.014	.04	8,699,672
Recovery Value	445,196	222,809	224,469	391,763	733,906	8,699,672

Table 1: Summary Statistics of Insurance Contract Characteristics Annualized premium, coverage, deductible, and recovery value are in dollar amount. AAL represents annual average loss rate (as a share of recovery value), AEP 2% Loss corresponds to the loss rate with aggregate exceedance probability equal to 2%, and AEP 1% Loss corresponds to the loss rate with aggregate exceedance probability equal to 1%.

Although the three moments of the damage distribution - mean, 98th percentile (AEP 2%), and 99th percentile (AEP 1%) - are correlated, depending on the nature of hazard, a region with a high mean can have low tail risk. See Figure A.1 for comparing the three moments spatially. For instance, counties in California with high AAL (value above 75 percentile) can have medium AEP 1% value (falling between the 25 and 75 percentile).

3 Anatomy of Property Insurance Contracts

A property insurance contract transfers the financial consequences of property loss from the policyholder to the insurer in exchange for the premium.[7] The main economic components of property insurance contracts are the premium, the coverage limit, and the deductible.

The premium is the price of transferring risk to the insurer. Economically, it reflects the expected loss and risk premium. The expected loss would depend on the exposure to natural hazards as well as policyholder characteristics. The exposure to natural hazards would depend on the property location as well as property characteristics such as building age, level, and materials. The observable policyholder characteristics are relevant because it could be informative about the extent of moral hazard, i.e. unobservable actions by policyholders that affect the riskiness of the property. Both policyholder characteristics and insurers' ability to manage risk (e.g., through reinsurance) subject to constraints (e.g., regulation and other frictions) would affect the risk premium.

The coverage limit specifies the limit of protection. Coverage caps the insurer's liability and controls exposure to tail risks. Broader coverage provides greater consumption smoothing from the policyholder's side and therefore comes with a higher premium. A contract with a lower coverage increases the residual risk born by the policyholder and therefore comes with a lower premium. From the insurer's perspective, lower coverage offers protection from moral hazard because it incentivizes policyholders to put more efforts into maintaining their homes.

The deductible is the amount of loss the policyholder must bear before the insurer starts to cover any damages. It plays an important role in aligning incentives. By exposing the policyholder to the initial portion of loss, deductibles discourage negligent behavior and mitigate moral hazard. An optimal level of deductible depends on the relative likeliness of small damages under policyholder's effort vs. negligence. If small damages

[7] For a comprehensive overview of public and private natural disaster insurance programs in the United States, see Wagner and Marcoux (2024).

are informative about the policyholder's actions, the insurer has the incentive to expose the policyholder to that risk via deductibles.

3.1 Stylized Facts

Based on the constructed dataset of 8.7 million contracts, we examine the relationships among premium, coverage limit, deductible, as well as the damage risk distribution, and document stylized facts about US property insurance contracts.

1. **Coverage limits are rarely binding.** The median coverage-to-recovery value ratio is 77%, while the median 99th-percentile damage-to-recovery value ratio is only 1.4%. This implies that the realized losses rarely approaches the coverage limit, so reducing coverage limit would not meaningfully reduce expected insurer payouts. In other words, coverage plays a limited role in constraining payouts from the insurers' perspective. This suggests that deductibles are relatively more important than coverage limits for determining household risk exposure and incentive provision.

   This also implies that measures such as premium divided by coverage are not informative about the effective price of insurance. Expected payouts are highly nonlinear in the coverage limit, and the deductible determines whether any payout occurs at all. The relevant price of insurance is therefore determined by the deductible and the distribution of losses relative to it, rather than by the nominal coverage limit.

   One might be concerned that our damage measure overstates insured exposure because it includes flood risk, whereas direct inundation losses are typically excluded from standard homeowners policies. Excluding flood, however, makes little difference: the median tail loss rate (99th-percentile damage-to-recovery value ratio) remains 1.4%, and the mean damage-to-recovery value declines by only 0.1 percentage point.

Another natural question is whether the measure understates effective insured exposure because homeowners insurance also covers non-climate-related hazards such as theft and liability. As a back-of-the-envelope adjustment, removing the average share of claims attributable to theft, liability, and other non-property categories reported in industry data increases the median tail loss rate to 1.6%.[8] Even under a conservative benchmark in which the CoreLogic measures are interpreted as capturing only catastrophe risk, the tail loss rate remains a very small fraction of the coverage limit. This benchmark is conservative because the CoreLogic estimates are based on the full distribution of modeled losses, whereas catastrophe losses are defined only as realized losses exceeding a large threshold. Applying a catastrophe share of total insured losses based on industry data and our calculations yields a median tail loss rate of 2.2%, still small relative to the median coverage-to-recovery value ratio of 77%, indicating that coverage limits rarely bind.[9]

2. **Deductibles are small relative to property value, but large relative to expected loss.** Deductibles are typically offered in discrete, round-number increments (e.g., \$500, \$1,000, \$1,500, \$2,000, \$2,500, \$5,000), with \$1,000 as the modal choice. Scaled by property recovery value, the median deductible is only 0.3%. However, the median annual expected loss is even smaller, at 0.09%, implying that households frequently absorb the initial portion of damages. Excluding the flood component reduces the median expected loss slightly to 0.087%, while adjusting for theft and liability shares increases it to 0.10%. Under the conservative benchmark in which the CoreLogic measure captures only catastrophe risk, the expected loss rises to 0.14%, which remains well below the median deductible of 0.3%.

[8] Insurance Information Institute (2025) reports that the average share of annual realized losses attributable to theft, liability, and other categories is about 12% based on 2019–2023 claims data.

[9] LexisNexis Risk Solutions (2025) reports that catastrophe losses accounted for 64% of total losses in 2024. Using a back-of-the-envelope calculation, we obtain a similar magnitude. Allocating portions of fire and water damage to accidental rather than natural causes and attributing all wind and hail damage to natural hazards (based on the statistics from Insurance Information Institute (2025) aggregated at the peril-year level), we estimate that approximately 65% of total losses are driven by natural hazards.

Consistent with this, we find that a substantial portion (e.g., two-thirds) of realized damage events fall entirely below the deductible, resulting in no insurer payout. Because a large share of losses occur in the region where the deductible applies, the deductible plays a particularly important role in determining household risk retention and mitigating moral hazard.

3. **Premiums substantially exceed expected losses.** The median ratio of annual expected loss (AAL) to premium is 28%, indicating that a large share of the premium is not accounted for by expected claim costs alone.

   Excluding the flood component reduces the ratio to 26%, while adjusting for theft and liability shares increases it to 32%, which is in line with the zipcode-level realized loss ratio (losses paid divided by written premium) disclosed by the Federal Insurance Office. Even under the conservative benchmark based on a catastrophe-share adjustment, the ratio is 43%, suggesting that premiums substantially exceed expected losses. This is consistent with equilibrium in a market with risk-averse policyholders: households are willing to pay substantially more than the expected claim payout to transfer disaster risk off their balance sheets, generating a risk premium.

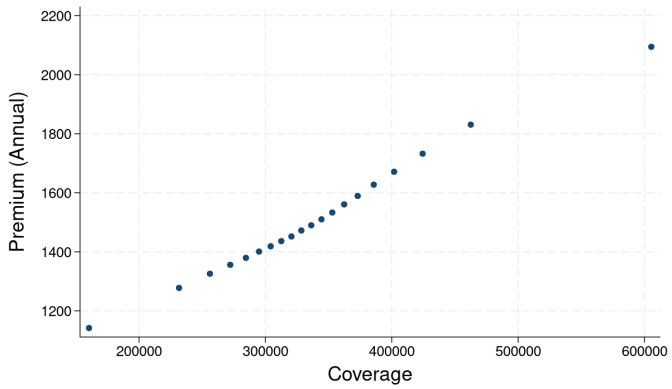
   We term the spread between premiums and expected losses the "policyholders' risk premium" because it captures the price households pay, in equilibrium, to avoid disaster risk. As long as insurers extract the full surplus relative to policyholders' outside option, regardless of the market competition, the spread reflects the policyholders' risk premium. Importantly, this interpretation does not require the spread to be free of insurer operating costs, capital costs, or markups. How this spread is allocated across operating costs, capital buffers, markups, or profits is therefore orthogonal to our definition. Our focus here is to decompose premiums into expected losses and the *policyholders'* risk premium.

Figure 1 summarizes the cross-sectional relationships among premiums, coverage, and deductibles. Panel (a) shows that, controlling for AAL, recovery value, and deductible (and including insurer and state fixed effects), premiums increase with coverage, which is consistent with the fact that higher coverage transfers a larger share of losses to the insurer and is therefore priced higher.[10] Panel (b) suggests that premium decreases with the deductible, for the same coverage, recovery value, and damage risk (AAL) within the same insurer and the state where policy was sold. This is consistent with households bearing more of the initial loss when the deductible is higher. The slope is relatively flat especially for the middle range, though. Panel (c) shows that deductibles increase with coverage after controlling for premium and the same controls.

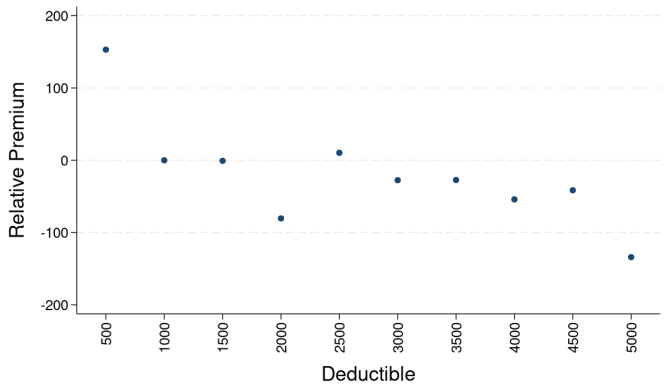
4. **Damage risk is low on average but highly skewed.** The median annual expected loss rate is 0.09% of property recovery value, yet the median 98th-percentile loss rate is 0.8% and the 99th-percentile loss rate is 1.4%. Thus, although policyholders are expected to lose less than 0.1% of property recovery value per year on average, tail losses are an order of magnitude larger: a one-in-one-hundred-years event is roughly 16 times the expected annual loss. Excluding the flood component yields similar magnitudes. This combination of low mean and heavy upper tail implies that the loss distribution is highly skewed with substantial mass near zero, which motivates modeling damage using a zero-inflated beta family of distributions.

Taken together, these stylized facts show that coverage limits rarely bind, while deductibles frequently do, and that premiums reflect both expected loss and the value of transferring tail risk, rather than simply expected payouts. Moreover, losses are highly skewed, with a large mass of small damages and infrequent but substantial tail events. As a result, the deductible is an economically relevant contract margin, governing both risk

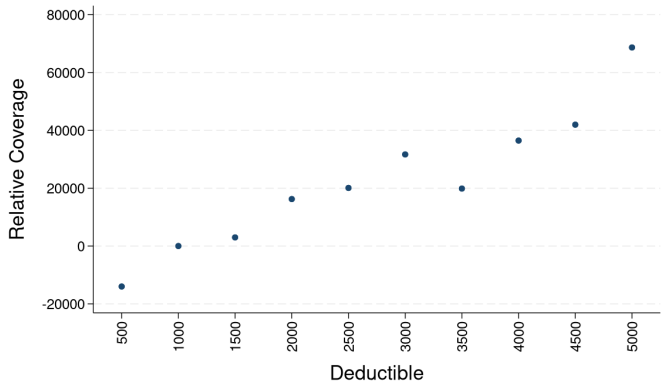
[10] Instead of controlling for the recovery value, one could compare *premium/recovery value* versus *coverage/recovery value*; we confirm that this does not change the result.



(a) Premium vs. Coverage (Controlling for deductible)



(b) Premium vs. Deductible (Controlling for coverage)



(c) Deductible vs. Coverage (Controlling for premium)

Figure 1: Bivariate relationships among premium, coverage, and deductible These figures are binned scatter plots after controlling for AAL, recovery value, insurer fixed effects, and state fixed effects. Panels (b) and (c) plot coefficients on deductible dummies, where the base category is deductible of $1,000.

retention and incentive provision, and must be understood jointly with premiums and coverage.

To analyze these endogenous contract terms, we develop a model in which premiums, deductibles, and coverage limits are designed optimally by insurers, given the risks and preferences of households. The model links the distribution of damage risk and household risk aversion, to the observed structure of insurance contracts. This framework allows us to quantify how much risk is transferred versus retained, how much moral hazard is mitigated through cost-sharing, and how these components vary across households, insurers, and regions.

4 Model

We write a model that is consistent with the stylized facts that we document. With this model, we structurally estimate risk premium, cost of moral hazard, and other latent parameters, such as the cost of effort and the distribution of damages given negligence. This allows us to provide policy implications based on the counterfactual analyses.

4.1 Model Setting

A potential policyholder seeks insurance for stochastic damage x to the property. She is risk averse; she has a CARA utility with risk aversion parameter of ρ.[11] If she takes good care of the property ($a = 1$) at effort cost of ϕ, the damage x is drawn from distribution $f_1(x)$ with a finite support of $[0, H]$, where H denotes the recovery value. If she is negligent ($a = 0$), the damage x is drawn from distribution $f_0(x)$ with the same support of $[0, \bar{x}]$.

[11] CARA utility offers a number of benefits and is a common choice in the theoretical literature on insurance. First, it is highly tractable and allows for analytical inversion of moment conditions. Second, the utility does not depend on the current level of wealth, which tends to be difficult to observe in practice. Third, it allows for dynamic extensions that preserves the dynamics in the static model. The dynamic optimal contract is a series of static optimal contract, as the utility doesn't depend on the level of wealth.

This formulation means that exerting maintenance effort ($a = 1$) stochastically reduces damage. Critically, we do not assume maintenance prevents damage entirely; severe disasters, of course, cause substantial losses regardless. Rather, we assume that maintenance effort shifts the loss distribution downward. This assumption is also consistent with practice. Insurers explicitly deny claims when negligence is documented, revealing that they view effort as materially affecting loss distributions. Notably, it is costly for insurers to verify and differentiate negligence from non-negligence. Despite the cost, insurers still conduct such investigations, suggesting they indeed view policyholders' actions as economically significant.[12]

We make the simplest assumption on the two distributions that the expected damage is lower under effort than under negligence: $E[x|a = 1] < E[x|a = 0]$. Her utility function can therefore be written as:

$$u(w, a) = -\frac{1}{\rho}e^{-\rho(w-a\phi)} \tag{1}$$

, where w denotes her terminal wealth and ϕ denotes the cost of effort. Here, we assume that the welfare implication of action a is additively separable from that of wealth within the exponent.

We assume that, in the absence of insurance, the potential policyholder finds it optimal to exert effort ($a = 1$). In other words, we assume the following:

$$E[u(-x, 1)|a = 1] > E[u(-x, 0)|a = 0] \tag{2}$$

4.2 Optimal Contract

Based on the setting provided in the previous section, we solve for the optimal contract, the net insurance payoff $y(x)$ to the policyholder for a given damage of x. The net insurance

[12]Relatedly, one might be concerned that, in severe disasters, losses cannot be mitigated. However, evidence from Florida suggests otherwise: Czajkowski et al. (2017) find that compliance to Florida Building Code reduces windstorm losses by up to 72 percent.

payoff $y(x)$ to the policyholder can be broken down as follows:

$$y(x) = -p + I(x) \tag{3}$$

where p denotes the premium and $I(x)$ denotes the claim payout. The policyholder's terminal wealth is therefore $-x + y(x)$ with insurance and $-x$ without insurance.

The insurer's problem is to minimize the expected payoff to the policyholder, subject to the following constraints: (1) the participation constraint, where the policyholder should prefer to participate in the insurance contract as opposed to staying uninsured[13] and (2) the incentive compatibility constraint, where the policyholder should prefer to take the action stipulated by the contract instead of deviating to the other. In this model, we assume that the insurer is risk-neutral.[14] We focus on the contract that induces effort from policyholders. A contract that does not induce effort would trivially imply complete insurance, but the data do not support this: the observed contracts do not take the form of full insurance. The insurer's problem can therefore be written as follows:

$$\max_{y(\cdot)} -E[y(x)|a = 1] \tag{4}$$

subject to:

$$\underbrace{E[u(-x + y(x), 1)|a = 1]}_{\text{Utility with Insurance}} \geq \underbrace{E[u(-x, 1)|a = 1]}_{\text{Utility with Outside Option}} \tag{P}$$

[13]We are thereby implicitly assuming that the best outside option is not having insurance, as opposed to having another insurance product. This simplification allows us to focus on contract design and risk-sharing. Moreover, empirical evidence suggests that local insurance markets, particularly in high-risk regions, tend to be highly concentrated. The literature also documents substantial search frictions and switching costs in insurance markets, which limit households' ability to costlessly switch across insurers. In addition, regulatory constraints in property insurance markets raise barriers to entry. Taken together, these factors reduce the extent to which alternative insurers constitute a frictionless outside option.

[14]The rationale is that insurers are sufficiently diversified that they are effectively risk-neutral when designing individual contracts. This is consistent with the lack of coinsurance, which Arrow (1965) predicts as an outcome when the insurer is risk-averse and therefore shares risk with the policyholder.

and

$$\underbrace{E[u(-x+y(x),1)|a=1]}_{\text{Utility with Effort}} \geq \underbrace{E[u(-x+y(x),0)|a=0]}_{\text{Utility without Effort}} \tag{IC}$$

We assume that both constraints above bind, as the insurer would not have any incentive to leave money on the table by allowing for slack in either of the constraints.

The first order condition FOC w.r.t. payoff y is:

$$\lambda + \mu\left(e^{\rho\phi} - \frac{f_0(x)}{f_1(x)}\right) = \frac{1}{\rho}e^{-\rho(x-y(x))} \tag{FOC}$$

where λ and μ denote shadow costs of P and IC, respectively.

The optimal payoff is thus:

$$y^*(x) = x + \frac{1}{\rho}\log\left(\rho\lambda + \rho\mu\left(e^{\rho\phi} - \frac{f_0(x)}{f_1(x)}\right)\right) \tag{5}$$

To check the optimality of this solution, we examine the second order condition 6. For the solution above to be optimal, the following must hold:

$$-\rho^2\left(\lambda + \mu\left(e^{\rho\phi} - \frac{f_0(x)}{f_1(x)}\right)\right)e^{-\rho(x-y(x))} < 0 \tag{SOC}$$

which is true if and only if:

$$\lambda + \mu\left(e^{\rho\phi} - \frac{f_0(x)}{f_1(x)}\right) > 0 \tag{6}$$

It can there be seen that if there exists $y^*(x)$ that is real for every x in the support of $f_1(x)$, the $y^*(x)$ indeed (locally) maximizes the insurer's objective function.

5 Structural Estimation

5.1 Identification

For the identification to be feasible, we make additional assumptions. First, we assume that the two constraints, P and IC, bind with equality. The binding P identifies the risk aversion ρ. The risk aversion should be such that the policyholder is indifferent between being insured and being uninsured:

$$-\frac{1}{\rho}\int_0^{\bar{x}} e^{-\rho(-x+y(x)-\phi)} f_1(x)dx = -\frac{1}{\rho}\int_0^{\bar{x}} e^{-\rho(-x-\phi)} f_1(x)dx \tag{7}$$

Given that the policyholder will exert effort in either case, the above can be further simplified:

$$\frac{1}{\rho}\int_0^{\bar{x}} e^{\rho(x-y(x))} f_1(x)dx = \frac{1}{\rho}\int_0^{\bar{x}} e^{\rho x} f_1(x)dx \tag{8}$$

The binding IC, combined with binding P and the FOC, identifies the shadow cost:

$$\lambda = \frac{1}{\rho\int_0^{\bar{x}} e^{\rho x} f_1(x)dx} \tag{9}$$

The intuition for the result above is that participation is harder to induce when the policyholder is happier without insurance.

Second, we assume that zero damage almost perfectly signals effort by the policyholder. In other words, we assume that $\frac{f_0(0)}{f_1(0)} \approx 0$. This assumption is consistent with the ex-ante contracting framework of Holmström (1979). Note that the likelihood ratio, $\frac{f_0(0)}{f_1(0)}$, is defined over the *unconditional* distribution of damage, integrating over disaster realizations. Zero damage is therefore, the outcome most informative about ex-ante effort relative to interior damage levels, where both effort and luck jointly determine outcomes.[15] This is also consistent with the fact that the minimum net cost $(x - y(x))$ under insurance to the

[15]Similar assumptions are common in structural contracting models — see, e.g., Gayle and Miller (2015).

policyholder is typically well-defined in practice: insurance premium for zero damage. This assumption, combined with the FOC, provides the following moment condition:

$$-\frac{1}{\rho\mu}e^{\rho y(0)} + \frac{\lambda}{\mu} + e^{\rho\phi} = \frac{f_0(0)}{f_1(0)} \approx 0 \tag{10}$$

Let $p = -y(0)$ denote the premium. Then, the above condition can be rewritten as:

$$\frac{1}{\rho}e^{-\rho p} = \lambda + \mu e^{\rho\phi} \tag{11}$$

The FOC, in conjunction with the property of probability distribution that integrates to 1 (i.e., $\int_0^{\bar{x}} f_1(x)dx = 1$ and $\int_0^{\bar{x}} f_0(x)dx = 1$), provides the final moment condition:

$$\frac{1}{\rho}\int_0^{\bar{x}} e^{-\rho(x-y(x))} f_1(x)dx = \lambda + \mu(e^{\rho\phi} - 1) \tag{12}$$

The two conditions above, Equations (11) and (12) jointly identify the shadow cost of incentive compatibility μ and the cost of effort ϕ.

For the estimation, we use the observed contract $y(x)$ and the distribution of damage given the policyholder's effort $f_1(x)$. We define the following moments as functions of risk aversion ρ:

$$\alpha(\rho) = \frac{1}{\rho}\int_0^{\bar{x}} e^{-\rho(x-y(x))} f_1(x)dx \tag{13}$$

$$\beta(\rho) = -\frac{1}{\rho}\log\left(\int_0^{\bar{x}} e^{\rho(x-y(x))} f_1(x)dx\right) \tag{14}$$

$$\gamma(\rho) = -\frac{1}{\rho}\log\left(\int_0^{\bar{x}} e^{\rho x} f_1(x)dx\right) \tag{15}$$

$$\delta(\rho) = \frac{1}{\rho}e^{-\rho p} \tag{16}$$

Then, we take the following steps for estimating the parameters. The first step is to numerically estimate $\hat{\rho}$ from 8:

$$\beta(\hat{\rho}) = \gamma(\hat{\rho}) \tag{17}$$

Given that the payoff under insurance is less risky relative to that without insurance, we conjecture that there exists a unique $\hat{\rho}$ that satisfies the condition above. The second step is to compute the three moments, $\alpha(\hat{\rho})$, $\gamma(\hat{\rho})$, $\delta(\hat{\rho})$ for the estimated risk aversion $\rho = \hat{\rho}$. The final step is to compute the parameters λ, μ, and ϕ from the remaining moment conditions, Equations (9) - (12). $\hat{\lambda}$ follows immediately from Equation (9):

$$\hat{\lambda} = \frac{1}{\hat{\rho}} e^{\hat{\rho}\gamma(\hat{\rho})} \tag{18}$$

$\hat{\mu}$ and $\hat{\phi}$ follow jointly from Equations (11) and (12):

$$\hat{\mu} = \delta(\hat{\rho}) - \alpha(\hat{\rho}) \tag{19}$$

$$\hat{\phi} = \frac{1}{\rho} \log \left(\frac{\delta(\hat{\rho}) - \frac{1}{\hat{\rho}} e^{\hat{\rho}\gamma(\hat{\rho})}}{\delta(\hat{\rho}) - \alpha(\hat{\rho})} \right) \tag{20}$$

5.2 Cost of Moral Hazard

In the absence of moral hazard, the insurer would offer a "complete" insurance: $y(x) = -p^{FB} + x$. The optimal contract can trivially discourage negligence when it is observable. The first-best premium p^{FB} will be determined by the participation constraint:

$$\frac{1}{\rho} e^{\rho p^{FB}} = \frac{1}{\rho} \int_0^{\bar{x}} e^{\rho x} f_1(x) dx \tag{21}$$

This gives:

$$\hat{p}^{FB} = -\gamma(\hat{\rho}) \tag{22}$$

Note that the estimation of risk aversion does not depend on model assumptions other than the CARA utility of policyholders and the feasibility of inducing effort.

The cost of moral hazard would be the difference between the risk premium that can be charged in the complete insurance and that charged in the partial insurance we observe in the data. The risk premium for the complete insurance is given as:

$$RP^{FB} = p^{FB} - E[x|a = 1] \tag{23}$$

That for the insurance we observe is given as:

$$RP = p - E[I(x)|a = 1] \tag{24}$$

Therefore, the cost of moral hazard becomes:

$$\Delta V = RP^{FB} - RP = (p^{FB} - E[x|a = 1]) - (p - E[I(x)|a = 1]) \tag{25}$$

5.3 Estimation

Having established the model and identification, we now turn to estimation. This section outlines how we use the data to recover the model's key parameters and illustrates how the estimation approach operationalizes the mechanisms described above.

The first step is to estimate the distribution of loss rate, the damage as a portion of recovery value ($\frac{x}{H} \in [0, 1]$). Considering the pattern we document in Section 3 that the distribution features a very low mean and a thick tail, we fit a zero-inflated (mass at zero) beta-family distribution. Among the beta-family distribution, we choose the Kumaraswamy distribution, which is very close to the beta distribution but is more numerically stable and better behaved around the edges.[16] The bounded nature of the

[16]The behavior at the edges is non-trivial, because many contracts feature very low deductibles ($\frac{D}{H} \approx 0$) and very high coverages ($\frac{C}{H} \approx 1$).

distribution precludes outcomes that perfectly reveal negligence, which may lead to trivial optimal contracts that achieve near first-best efficiency.[17] The probability distribution function (PDF) of the zero-inflated Kumaraswamy distribution is given as:

$$f(x;\pi_0,a,b) = \pi_0\delta\left(\frac{x}{H}\right) + (1-\pi_0)abx^{a-1}\left(1-\left(\frac{x}{H}\right)^a\right)^{b-1} \tag{26}$$

where $\delta(t)$ denotes a Dirac delta function that captures the mass of π_0 at 0, and a and b are non-negative shape parameters. As the distribution is modeled for the loss rate, the damage x is scaled by the recovery value H.

We make use of three moments of loss rate provided by CoreLogic. For each property, we observe: mean denoted by m, 98th percentile denoted by q_{98}, and 99th percentile denoted by q_{99}. With these three moments (m, q_{98}, q_{99}), the parameters (π_0, a, b) are just identified.

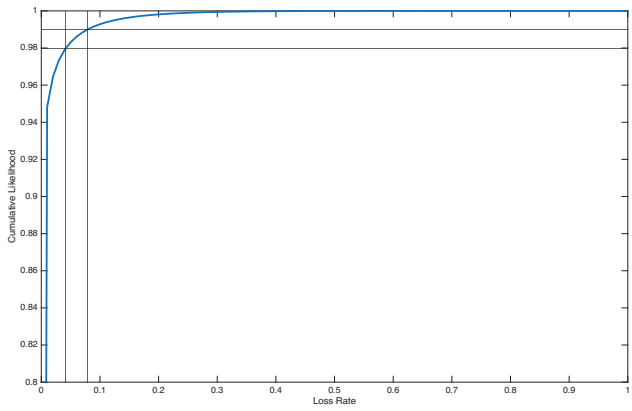


Figure 2: CDF of Estimated Loss Rate Distribution This figure plots the fitted loss rate CDF distribution for a property with $(m, q_{98}, q_{99}) = (0.003, 0.042, 0.079)$.

Figure 2 plots an example of the cumulative distribution function (CDF) of the estimated loss rate distribution for a contract. Note that it closely replicates the quantile moments. The two key aspects of this distribution are: (i) a heavy point mass at 0 (π_0) and (ii) a thick tail that diminishes much slower than the normal or the exponential distributions. These jointly explain the small mean and the large quantiles. As we assume that the contract

[17] Mirrlees (1974)

induces effort by the policyholder, this equilibrium damage distribution $f(x; \pi_0, a, b)$ will serve as $\hat{f}_1(x)$.

The second step is to decompose the premium into the expected claim payout and the risk premium. Specifically, we compute the negative expected net payoff of a contract of premium p, deductible D, and coverage C to the policyholder. This is equivalent to the risk premium that the risk-neutral insurer collects from the risk-averse policyholder:

$$-E[y(x; p, D, C)] = p - E[\min(\max(x - D, 0), C)] \tag{27}$$

The third step is to estimate the risk aversion (ρ), using the participation condition in Equation (P). In the numerical implementation, we compare the certainty equivalent (CE) with and without insurance, as opposed to comparing the expected utility. Using the CE offers both numerical stability for high level of risk aversion, as well as smooth convergence to risk neutrality as the risk aversion approaches zero..[18] Based on the CE representation, we numerically solve for the risk aversion parameter ρ that satisfies:

$$-\frac{1}{\rho}\log(E[e^{\rho(x-y(x))}|a=1]) = -\frac{1}{\rho}\log(E[e^{\rho x}|a=1]) \tag{28}$$

[18]While Bertomeu et al. (2025a) show that the CARA utility does converge to risk neutrality as ρ approaches to 0, the numerical implementation suffers from dividing by 0.

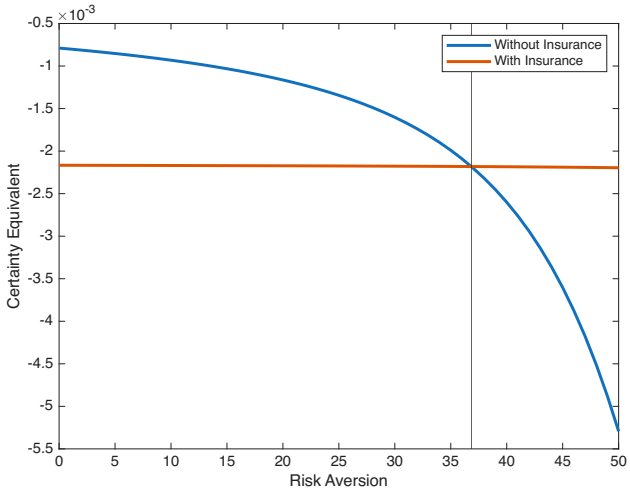


Figure 3: CE with and without Insurance This figure plots certainty equivalent, with and without insurance, as a function of risk aversion for a hypothetical contract with $(p, D, C) = (2000, 2000, 94000)$.

For an illustration, we plot the CE with and without insurance with respect to the risk aversion parameter ρ in Figure 3 based on the same contract used for Figure 2. Given that the payoff under insurance features lower mean and higher risk, the two CEs cross only once, offering identification of a unique risk aversion of $\hat{\rho}$. The estimated risk aversion for this hypothetical policyholder is about 37, where her CE with insurance equals her CE without an insurance.

The fourth step is to estimate the cost of moral hazard, which involves the premium and the expected damage under the complete insurance. Following from the analysis in Section 5.2, the premium of this first-best contract, p^{FB}, should equal to the negative CE with and without insurance. We expect to find $\hat{p}^{FB}$ higher than the observed premium. The expected damage under complete insurance is simply the untruncated expected damage: $E[x|a = 1]$, which would be higher than the expected payout under the second-best contract. The risk premium that the insurer could have earned in the first-best case is the premium less the expected payout, and we expect it to be higher than the risk premium in the second-best case we observe. The difference between the two is the cost of moral hazard:

$$\text{Cost of Moral Hazard} = RP^{FB} - RP$$

which represents the risk premium that the insurer could have earned had there been no moral hazard.

The fifth step is to estimate the remaining parameters, including the policyholder's cost of exerting effort, ϕ. To that end, we first compute compute the three moments, $\alpha(\hat{\rho})$, $\gamma(\hat{\rho})$, $\delta(\hat{\rho})$ for the estimated risk aversion $\rho = \hat{\rho}$. Then, we input the computed moments into Equations (18) to (20) and obtain estimates $(\hat{\lambda}, \hat{\mu}, \hat{\phi})$.

The sixth and final step is to infer the counterfactual damage distribution $f_0(x)$ from the FOC. By rearranging the FOC, we obtain the following expression for the $f_0(x)$:

$$f_0(x) = \left(e^{\rho\phi} - \frac{1}{\mu} \left(\frac{1}{\rho} e^{-\rho(x - y(x))} - \lambda \right) \right) f_1(x) \tag{29}$$

For consistency with the assumption in Section 5.1 that zero damage fully reveals effort (Equation (10)), we rebalance the recovered distribution so that it does not have a mass at zero. This adjustment ensures that the likelihood ratio $\frac{f_0(0)}{f_1(0)} = 0$ and has no effect on the relative informativeness among positive damages.

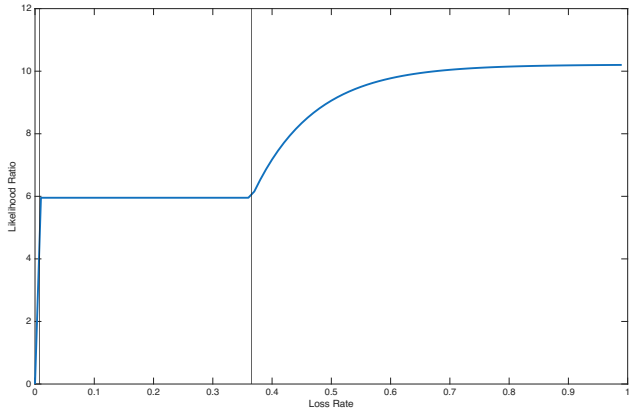


Figure 4: Likelihood Ratio This figure plots likelihood ratio, $\frac{f_0(x/H)}{f_1(x/H)}$, of a hypothetical contract with $(p, D, C) = (2000, 2000, 94000)$ and $(m, q_{98}, q_{99}) = (0.003, 0.042, 0.079)$.

Figure 4 plots the likelihood ratio, $\frac{f_0(x)}{f_1(x)}$, implied by the insurance contract we analyzed for Figure 2 and Figure 3. The two vertical lines represent the deductible and the sum of coverage scaled by recovery value. The likelihood ratio increases sharply in the deductible

range (left of the left vertical line), remains stable in the intermediate range where the insurer covers the damage, and increases substantially in the coverage range (right of the right vertical line) where the insurer limits the coverage.

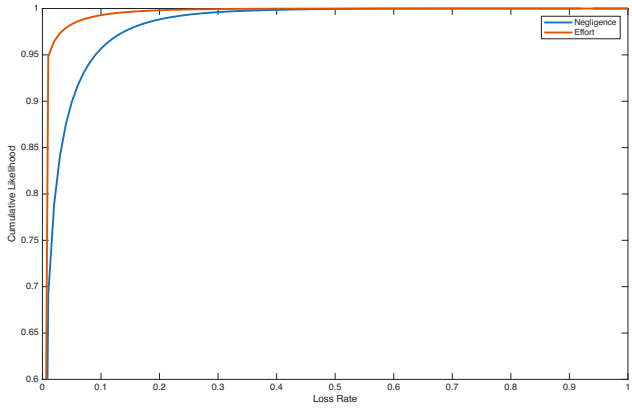


Figure 5: Cumulative Likelihood under Negligence vs. under Effort This figure plots cumulative likelihood under negligence (F_0) vs. effort (F_1) of a hypothetical contract with $(p, D, C) = (2000, 2000, 94000)$ and $(m, q_{98}, q_{99}) = (0.003, 0.042, 0.079)$.

Figure 5 plots the same contract's CDF of loss rate x under negligence (F_0) and that under effort (F_1). It can be immediately seen that the damage under negligence first-order stochastically dominates (FOSD) that under effort. Both the expected damage and the likelihood of extremely high damage is higher under negligence. For this contract, the expected damage increases substantially, which is 5 times higher than that under effort. That the expected increase in damage exceeds the monetary cost of effort is consistent with our assumption, Equation (2), that the policyholder prefers to exert effort in the absence of insurance.

These estimates are useful for counterfactual analysis, which we examine in detail in Section 7. In the counterfactual world where insurers are forced by regulation to provide full insurance at the premium that the policyholder would be willing to pay, the insurer would choose to withdraw from the market. By comparing the expected claim payout, which is the expected damage under negligence ($E[x|a = 0]$), with the premium that policyholder would be willing to pay for a full insurance ($\hat{p}^{FB}$), the structural model allows

us to assess whether taking away the insurer's ability to provide incentives will result in market failure.

6 Empirical Analysis

We apply the steps we outline in the previous section, contract by contract. In this section, we report the estimation results and validate our estimates based on economic hypotheses.

6.1 Estimation Results

Table 2 reports summary statistics of the risk aversion, risk premium, cost of moral hazard, and increase in exposure due to moral hazard estimated at the *contract level*. This is possible because we exploit proprietary data on property-level moments (mean and quantiles) of loss distribution used by insurers, rather than inferring from realized losses that tend to be sparse and infrequent for home insurance. The latter would likely require a pooled estimation from observations across properties and time to infer the loss distribution. However, since the underlying loss information is already estimated by insurers at the contract level, we are able to estimate the policyholders' risk aversion, risk premium component of the premium, the cost of moral hazard, and the increase in exposure due to moral hazard at the contract level.

	Mean	St.Dev.	P10	Median	P90
Risk Aversion	79	64	15	61	172
Risk Premium ($)	1,135	690	463	970	2,022
Risk Premium/Total Premium (%)	73	19	47	77	95
Moral Hazard ($)	7	37	.27	1.6	9.5
Moral Hazard/Risk Premium (%)	.7	3.6	.031	.16	.92
Increase in Exposure ($)	71	82	20	37	179
Increase in Exposure/Expected Payoff (%)	40	64	3.5	12	120

Table 2: Estimation Results This table reports summary statistics of the estimated variables.

6.1.1 Risk Aversion

Our estimated risk aversion ranges mostly between 15 and 170, with a mean of 79. In contrast to the existing estimates based on a pooled estimation, we provide risk aversion estimates at the *contract level*, which allows us to study its cross-sections. To our knowledge, this is the first paper offering estimates of policyholder-level risk aversion from a large sample and characterizing its rich heterogeneity spanning multiple cross-sections.

First, we show that high risk aversion is associated with low property risk, which adds validity to our estimates. Figure 6 shows the bilateral relationship, indicating that risk-averse policyholders are drawn to low risk (safer) properties.

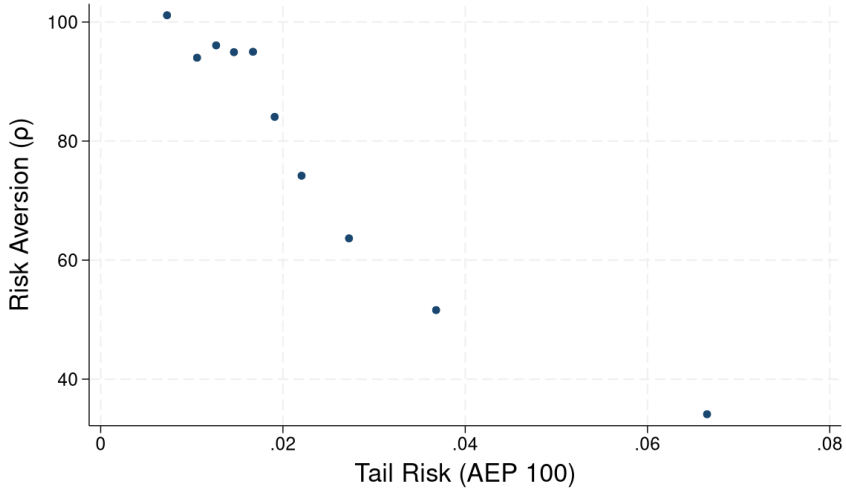
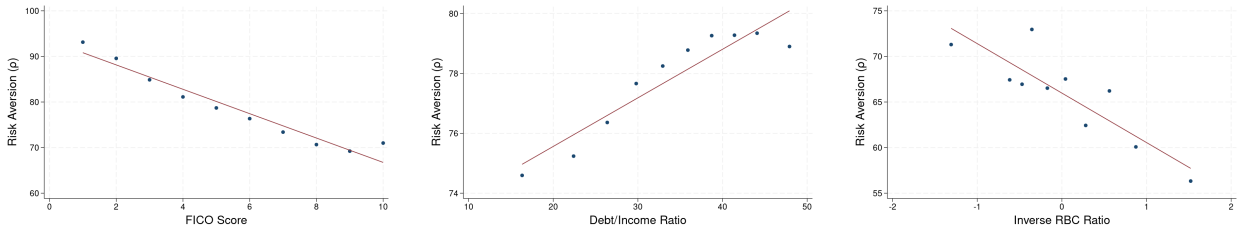


Figure 6: Risk Aversion Estimates vs. Damage Risk This figure shows a binned scatter plot of ρ estimates vs. the 99-percentile damage (AEP 100).

Next, we examine the relationship between risk aversion and borrower characteristics, which are equilibrium outcomes. Note that the goal of this exercise is to test relationships that are ex-ante unclear, rather than to establish causality. The relationship between estimated risk aversion and FICO can be both positive and negative due to countervailing forces. On one hand, lower-FICO policyholders face tighter financial constraints, making them more sensitive to consequences of loss and are therefore more risk averse.[19] On the

[19] Campello et al. (2010)

other hand, higher-FICO policyholders likely have higher value at risk and therefore more risk averse.[20] Figure 7a shows that, overall, risk aversion declines in FICO, supporting the financial constraint channel. However, for the highest FICO bins, we find that risk aversion slightly bounces back, which is consistent with the at-risk channel.



(a) FICO Score **(b)** DTI Ratio **(c)** Insurer Fin. Constraint

Figure 7: Risk Aversion, Policyholder Characteristics, and Insurer Financial Constraint

The relationship between risk aversion and the debt-to-income (DTI) ratio is also theoretically ambiguous. A higher DTI can be associated with tighter financial constraints and therefore higher marginal cost of financial distress, leading to a higher risk aversion. Therefore, risk aversion can increase in DTI. At the same time, a higher DTI indicates a higher leverage, reflecting greater willingness to take risk. In this case, risk aversion will decrease in DTI. We find evidence supporting the former as shown in Figure 7b.

We also examine the relationship between policyholder risk aversion and insurer characteristics. Motivated by papers emphasizing the effect of financial constraint on insurer behavior, we focus on insurers' risk-based capital (RBC) ratios. We invert the RBC ratio to proxy it as the financial constraint. Again, the relationship between policyholder's risk aversion and insurer's financial constraint is ex-ante unclear. We find that more financially constrained insurers tend to provide insurance to less risk averse policyholders. Building upon the finding on risk aversion increasing in risk (Figure 6), this would be consistent with mechanism where more constrained insurers taking more risk, selling insurance to properties with higher risk, and policyholders who live in risky properties tend to have low risk aversion.

[20] Paravisini et al. (2017)

We estimate the relationships discussed above using the following regression specification for policyholder-level variables, FICO score and DTI ratio:

$$Y_{c,s,i} = \alpha + \beta^{HH} HH_{c,s,i} + \beta' X_c + \gamma_s + \delta_i + \varepsilon_{c,s,i} \tag{30}$$

where $HH_{c,s,i}$ denotes a household characteristic, either standardized FICO score or DTI ratio of contract c in state s, offered by insurer i. The outcome variable $Y_{c,s,i}$ denotes the estimated parameter of interest— risk aversion parameter *rho* in this subsection— and the same specification is applied when Y represents other estimated parameters in the next subsections. The dependent variable X_c includes contract-level controls including AEP 100, property recovery value. We include state-fixed effects (γ_s) and insurer- fixed effects (δ_i), and cluster standard errors at the zip-code level.

To examine the cross-section of insurer characteristics, we aggregate the contract-level data to the insurer level by taking the premium-weighted average of each variable. We then estimate the following regression:

$$Y_{c,s,i} = \alpha + \beta^{Insurer} Insurer\ FC_i + \beta' X_i + \varepsilon_i \tag{31}$$

where $Insurer\ FC_i$ denotes insurer financial constraint, proxied by the inverse RBC ratio. X_i denotes insurer-level controls, including size, leverage, and equity-to-premium ratio. Because this regression is conducted at the insurer level, we do not include state or insurer fixed effects.

Table 3 reports the results, showing that the coefficients are significant after controlling for the fixed effects. Taken together, we find that households with low credit scores and high DTI ratios tend to show high risk aversion. Insurers facing tighter financial constraint tend to offer insurance to those with low risk aversion, who tend to own properties that are more exposed to damage risk.

	(1) Risk Aversion	(2) Risk Aversion	(3) Risk Aversion
FICO Score	-2.511***		
	(-31.87)		
DTI Ratio		0.111**	
		(2.33)	
Inverse RBC Ratio			-5.425**
			(-2.34)
Controls	Y	Y	Y
State FE	Y	Y	N
Insurer FE	Y	Y	N
N	6,121,694	3,439,829	139

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Risk Aversion, Policyholder Characteristics, and Insurer Financial Constraint

These results highlight the importance of accounting for endogenous matching between households and properties.

6.1.2 Risk Premium

The observed annual premium is decomposed into the expected claim payout and the risk premium:

$$\underbrace{p}_{\text{Observed Premium}} = \underbrace{E[I(x)|a=1]}_{\text{Expected Claim Payout}} + \underbrace{RP}_{\text{Risk Premium}} \tag{32}$$

Table 2 shows that, on average, the risk premium constitutes about 73% of the annual premium. We first confirm that higher risk aversion is associated with a higher risk premium. This relationship is expected because more risk-averse agents would offer greater compensation to the insurer for bearing the same level of risk. Likewise, the estimated risk premium should increase with the underlying level of risk. Figure 8 documents both patterns, serving as a consistency check for our estimates. Moreover, in light of previous work, this risk premium is not unusually high; it follows directly from our estimated risk aversion, which lies toward the lower end of estimates in the literature

(e.g., Cohen and Einav, 2007; Einav et al., 2010; Sydnor, 2010; Handel, 2013; Hendren, 2020).

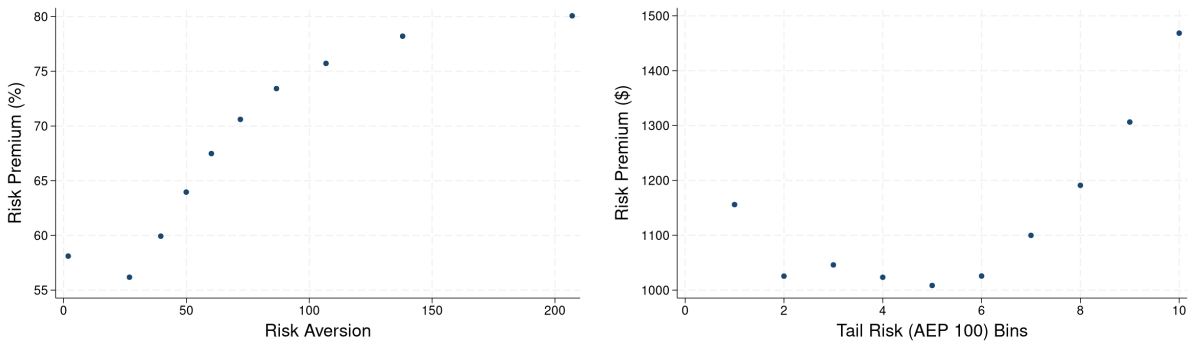
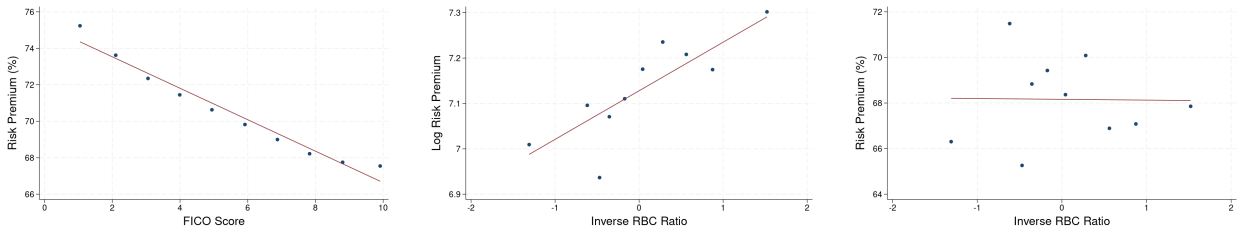


Figure 8: Risk Premium, Risk Aversion, and Damage Risk

While the risk premium represents policyholders' willingness to pay for transferring the risk to the insurer, from the insurer's perspective, it represents the expected profit that insurer can make from the respective contract. We find that this risk premium, either in dollar value or as a share of premium, decreases in FICO score. The relationship between scaled premium and FICO score is as presented in Figure 9a. Taken together with Figure 7a, a mechanism behind the premium increasing in credit risk is low-FICO policyholders having higher risk aversion, plausibly due to the household's financial constraint. These findings are consistent with low-FICO policyholders acquiring insurance to avoid having to borrow at exceedingly high rates (i.e., risk-taking capacity channel).



(a) FICO Score **(b)** Insurer Fin. Constraint **(c)** Insurer Fin. Constraint

Figure 9: Risk Premium, FICO Score, and Insurer Financial Constraint

Similar to the previous subsection, we investigate the role of insurer's financial constraint on risk premium. Consistent with the association between constrained insurers

and high risk (from Figure 7c and Figure 6), we find that more financially constrained insurers charge higher risk premium, as presented in Figure 9b. However, this does not necessarily imply that financially constrained insurers are extracting high profitability by selling insurance policies to households with risky properties. In fact, Figure 9c shows that risk premium as a share of total premium does not increase in the inverse RBC ratio.

Using the same specifications as in equations (30) and (31), we test the relationship between the risk premium, FICO, and the insurer's inverse RBC ratio. The results are reported in Table 4. Consistent with the description above, we find that the risk premium and the FICO score are negatively associated. In the cross-section of insurer financial constraint, we find that while more constrained insurers tend to collect more risk premium, the risk premium as a proportion of total premium is flat across insurers' financial constraint. This is consistent with the finding in the previous section that constrained insurers tend to insure more risky properties, and therefore, the expected claim payout proportion in (32) is large and the proportion of risk premium is relatively small.

	(1)	(2)	(3)	(4)
	Log Risk Premium	Risk Premium (%)	Log Risk Premium	Risk Premium (%)
FICO Score	-0.0817***	-1.403***		
	(-98.92)	(-62.99)		
Inverse RBC Ratio			0.107***	-0.0344
			(3.15)	(-0.03)
Controls	Y	Y	Y	Y
State FE	Y	Y	N	N
Insurer FE	Y	Y	N	N
N	7,691,488	6,886,496	139	139

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Risk Premium, Policyholder Characteristics, and Insurer Financial Constraint

6.1.3 Cost of Moral Hazard

Recall that we compute the risk premium in a counterfactual first-best world to compute the moral hazard. Specifically, in the first-best world in the absence of information asymmetry, full insurance would be optimal; insurers would offer full insurance in exchange for

counterfactual premium, and the coverage limit and deductible would be unnecessary because there is no moral hazard. The counterfactual premium can be decomposed into the expected damage and the counterfactual (first-best) risk premium:

$$\underbrace{p^{FB}}_{\text{Counterfactual Premium}} = \underbrace{E[I(x)|a=1]}_{\text{Expected Claim Payout}} + \underbrace{RP^{FB}}_{\text{Counterfactual Risk Premium}} \tag{33}$$

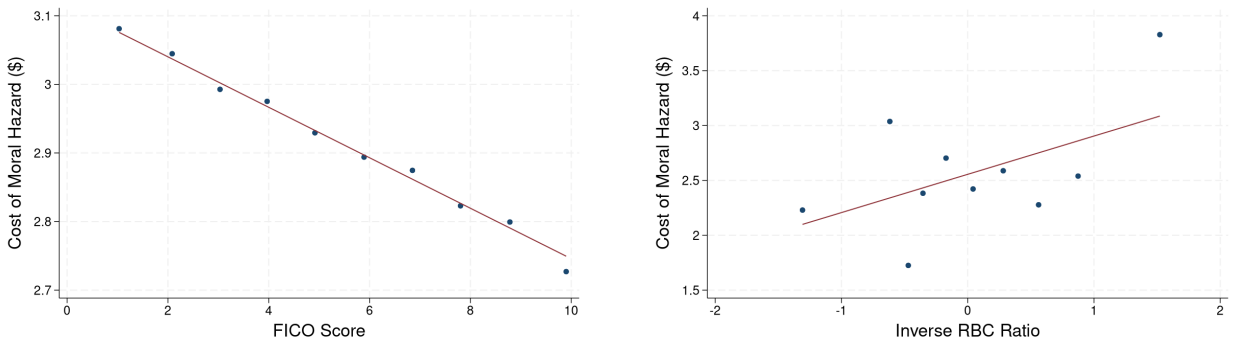
The cost of moral hazard estimate is:

$$\underbrace{\Delta V}_{\text{Cost of Moral Hazard}} = \underbrace{RP^{FB}}_{\text{Counterfactual Risk Premium}} - \underbrace{RP}_{\text{Observed Risk Premium}}$$

from the equations (32) and (33).

We find that the cost of moral hazard is low. Table 2 shows that the average cost of moral hazard is only \$7, or 0.7% of the estimated risk premium. This indicates that the contracts are well designed to discourage moral hazard; at the same time, this suggests that the deductible and the coverage, which reduces moral hazard, *increases* policyholder's risk exposure. We quantify this increase in risk exposure in the next subsection.

In the cross-section of FICO score, we find that the cost of moral hazard is higher for low-FICO policyholders, as presented in Figure 10a. This is consistent with the economic intuition that there would be more uncertainty regarding the unobservable actions by low-FICO policyholders than by high-FICO ones. In the cross-section of insurer's financial constraint, we find that more constrained insurers tend to face higher cost of moral hazard as shown in Figure 10b. Combined with previous results, constrained insurers are willing to take on more risk (i.e., insure riskier properties), facing higher costs of moral hazard as a result. Table 5 shows consistent results from regressions using the same specifications as (30) and (31).



(a) FICO Score **(b)** Insurer Fin. Constraint

Figure 10: Cost of Moral Hazard, FICO Score, and Insurer Financial Constraint

	(1) Cost of Moral Hazard	(2) Cost of Moral Hazard
FICO Score	-0.0400***	
	(-10.36)	
Inverse RBC Ratio		0.349**
		(2.49)
Controls	Y	Y
State FE	Y	N
Insurer FE	Y	N
N	5,442,243	139

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Cost of Moral Hazard, FICO Score, and Insurer Financial Constraint

6.1.4 Increase in Exposure due to Moral Hazard

The increase in exposure due to the coverage and deductible, which are in place to discourage moral hazard, is computed as:

$$\underbrace{\Delta I}_{\text{Increase in Exposure}} = \underbrace{E[x|a=1]}_{\text{Expected Damage}} - \underbrace{E[I(x)|a=1]}_{\text{Expected Claim Payout}}$$

from the equations (32) and (33). We scale this by the expected claim payout. Table 2 shows that the average increase in exposure relative to the expected payoff is 40%. That is, although the cost of moral hazard is low, the coverage limit and deductible that are designed to curb the moral hazard, exposes policyholders to 29% of expected payoff.[21]

We first examine the relationship between the increase in exposure and tail risk. Controlling for risk aversion and the average expected loss, we expect to find a positive relationship because increase in exposure is effectively the sum of the left tail cut off by the deductible and the right tail curtailed by the coverage. Therefore, the increase in exposure naturally increases with the tail risk that is not covered by the insurance contract. We control for risk aversion to isolate the negative correlation between risk aversion and tail risk. Because tail risk, not the average risk, is relevant for the increase in exposure, we control for the AAL. Figure 11 confirms this relationship. That is, households exposed to higher damage tail risk tend to retain *more* risk even after getting insurance.

In the cross-section of FICO, we find that low-FICO policyholders tend to be left with *more* exposure, as presented in Figure 12a. On the other hand, while increase in exposure is weakly negatively related with insurer's financial constraint (Figure 12b), this relationship is not significant in regression. Table 6 reports the regression results using the same specifications (30) and (31).

[21] That the insurer's claim payout increases by 40% under full insurance means that the partial insurance we observe in equilibrium covers only 71% of expected losses.

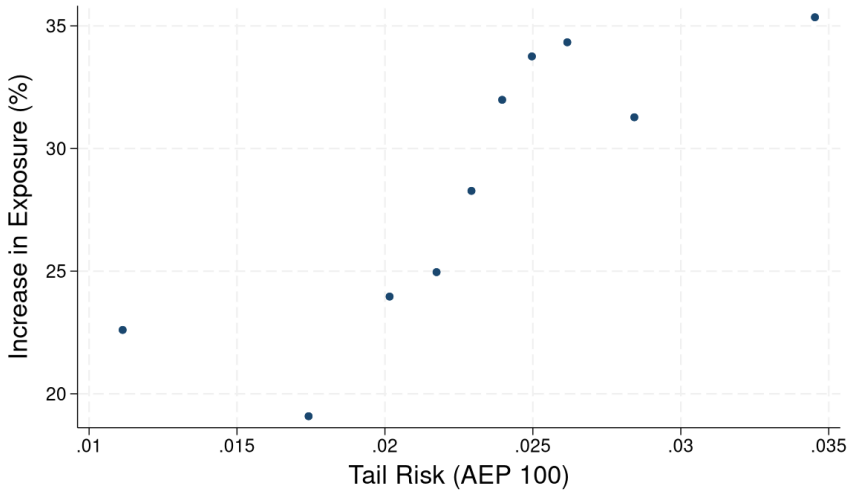
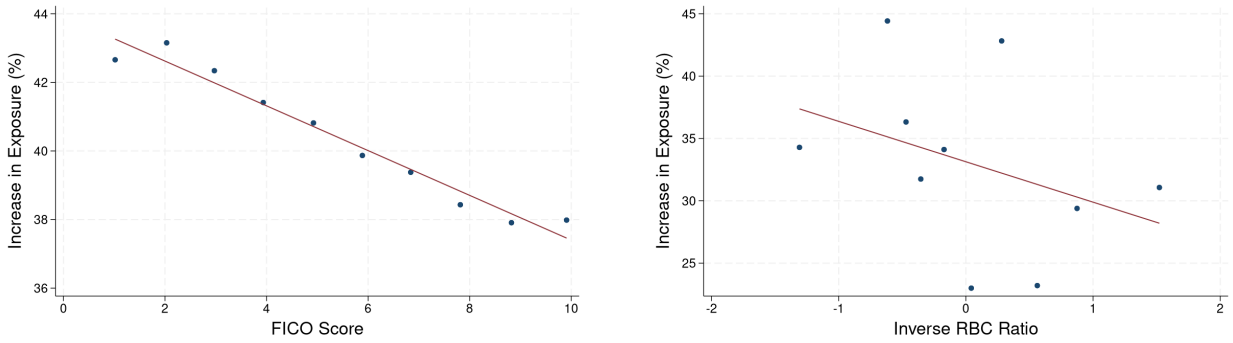


Figure 11: Increase in Exposure vs. Tail Risk This figure plots relationship between increase in exposure divided by expected payoff and tail risk (AEP 100), after controlling for risk aversion and expected loss (AAL).



(a) FICO Score **(b)** Insurer Fin. Constraint

Figure 12: Increase in Exposure due to Moral Hazard, FICO Score, and Insurer Financial Constraint

	(1) Increase in Exposure (%)	(2) Increase in Exposure (%)
FICO Score	0.181**	
	(2.01)	
Inverse RBC Ratio		-3.240
		(-1.11)
Controls	Y	Y
State FE	Y	N
Insurer FE	Y	N
N	6,955,719	139

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Increase in Exposure due to Moral Hazard, FICO Score, and Insurer Financial Constraint

6.1.5 Spatial Distribution

We map the spatial distribution of the estimated parameters in Figure A.2. To illustrate, we focus on Florida, a region with high exposure to natural disasters, as shown in Figure A.1.

Panel (a) shows that estimated risk aversion is relatively low, consistent with the idea that less risk-averse households are more willing to own property in higher-risk areas. Despite this, the elevated underlying risk leads to higher risk premiums (panel b). The average cost of moral hazard is moderate, reflecting contract designs that effectively limit opportunistic behavior (panel c). However, these same features leave households with substantial residual risk, as shown by the increase in uninsured exposure in panel (d). Comparing these results with the tail-risk patterns in Figure A.1 highlights that policyholders in regions with greater disaster exposure tend to retain the most residual risk.

Taken together, our estimates suggest that, although the cost of moral hazard is low, the contract design substantially increases the policyholders' risk exposure. From the cross-sectional analyses, we find that policyholders with low credit scores and high DTI ratio tend to be more risk averse and pay higher risk premium to insure property damage.

6.2 Validation

In this section, we test the validity of the cost of moral hazard estimates exploiting variation in the loan-to-value and variation in regulations across states.

6.2.1 Exploiting variation in "skin in the game"

To evaluate the economic validity of our cost of moral hazard estimates, we examine their cross-sectional relationship with the loan-to-value (LTV) ratio at origination. This test is not intended to establish causality, but rather to assess whether the estimates behave in a manner consistent with economic intuition.

In the cross-section, we hypothesize that moral hazard decreases as the property becomes more "owned." We proxy ownership by the LTV ratio. A lower LTV indicates that the homeowner has greater equity in the property and thus more "skin in the game." With more equity at stake, homeowners have stronger incentives to maintain the property, take preventive actions, and avoid filing excessive claims. Accordingly, we expect moral hazard to be less severe for properties with lower loan-to-value ratios.

Importantly, our estimate of moral hazard is derived exclusively from insurance contract characteristics—premium, deductible, and coverage—along with the underlying risk distribution, $f_1(x)$. It does not incorporate any borrower or financing information such as credit risk, loan-to-value, or other mortgage attributes. Consequently, examining how the estimated moral hazard varies with these external borrower characteristics provides an out-of-sample validation of our measure.

To test the hypothesis, we confine our sample to the loans that originated within one year and examine the relationship. Figure 13 shows that the moral hazard indeed increases in LTV. More formally, we regress the moral hazard estimate on LTV ratio, controlling for FICO, risk, and insurer fixed effects:

$$Cost\ of\ Moral\ Hazard_{c,i} = \beta^{LTV} LTV_c + \beta' X_c + \delta_i + \varepsilon_{c,i} \tag{34}$$

where the outcome variable is cost of moral hazard estimate of contract c underwritten by insurer i and the LTV_c denotes the policyholder's loan-to-value ratio. We control for the policyholder's FICO score, property's damage risk, recovery value, and the insurer fixed effects. Standard errors are clustered at the zip-code level.

As a result, we find significant and positive coefficients on the loan-to-value ratio.

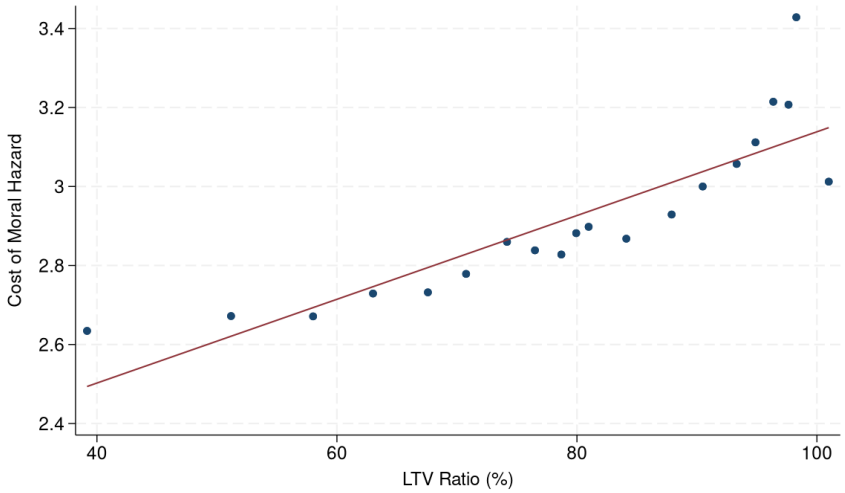


Figure 13: Cost of Moral Hazard vs. Loan-to-Value Ratio

	(1) Moral Hazard	(2) Moral Hazard	(3) Moral Hazard
LTV Ratio	0.221***	0.119***	0.125***
	(15.53)	(20.95)	(22.44)
Tail Risk (AEP100)		0.205***	0.0639***
		(9.74)	(3.36)
FICO Score		-0.0736***	-0.0631***
		(-12.59)	(-11.36)
Log Recovery Value		-1.051***	-1.027***
		(-39.32)	(-39.47)
Insurer FE	N	N	Y
N	1,261,431	1,261,431	1,261,424

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Cost of Moral Hazard vs. LTV Ratio

6.2.2 Exploiting variation in regulatory environment

Another approach to validating the cost of moral hazard estimate is to use cross-state differences in the insurance regulatory environment. Because insurance is regulated at the state level in the US, policies written on properties in a given state must comply with that state's requirements. These requirements include not only formal statutes, but also supervisory guidance and binding market practices that shape underwriting standards. The regulatory environment has several dimensions relevant for moral hazard.

The most prominent dimension is the requirement that homeowners provide documentation or inspection reports on the condition of key structural features, such as roof age and condition, hurricane tie-downs, or wind-mitigation features, before a policy is issued. Florida imposes these through statute. Several other high-risk Gulf and Southeastern states, insurers commonly require documentation of roof or structural condition, and coverage or pricing often depends on verified structural features. Although these requirements are not statutory in these states except Florida, they reflect state guidance and common underwriting practice.

We exploit this cross-state variation in regulation as a strategy to validate the cost of moral hazard estimates; we expect less severe moral hazard in the states with the inspection-related regulations discussed above. We test this hypothesis by running the following regression:

$$Cost\ of\ Moral\ Hazard_{c,s,i} = \beta^{Regulation} Inspection\ Regulation_s + \beta' X_c + \delta_i + \varepsilon_{c,s,i} \quad (35)$$

where the outcome variable is cost of moral hazard estimate of contract c, sold in state s by insurer i and the $Inspection\ Regulation$ dummy takes a value of 1 if the state has inspection-related regulation and 0 otherwise. We control for the policyholder's FICO score, property's damage risk, recovery value, and the insurer fixed effects. Standard errors

are clustered at the zip-code level. We expect $\beta^{Regulation}$ to be negative, as the regulations would likely reduce the cost of moral hazard by reducing the information asymmetry.

Consistent with the hypothesis, Figure 14 shows that the $\beta^{Regulation}$ is negative and significant.

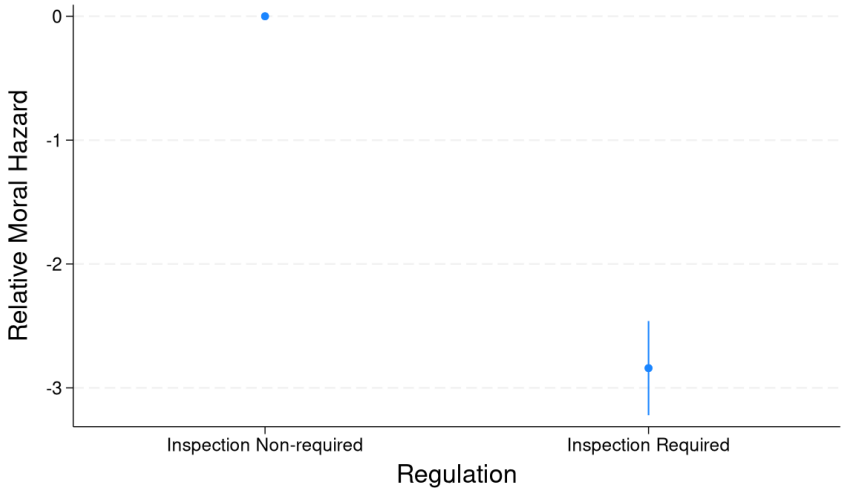


Figure 14: Cost of Moral Hazard vs. Inspection Requirement Regulation

The inspection requirement is the regulatory feature most directly connected to the moral hazard mechanism in our model. However, several additional regulations potentially affect the cost of moral hazard, either by influencing the incentives to file or inflate claims or by affecting the observability of risk. We therefore compile a broader set of regulatory provisions and combine them into a single index, where a higher value corresponds to a *lower* cost of moral hazard.

We consider five types of regulations in addition to the inspection requirement used above. First, some states restrict assignment of benefits (AOB), preventing contractors or third-party adjusters from assuming control of the claim and negotiating directly with the insurer. These restrictions limit the scope for repair-cost inflation and should therefore reduce moral hazard. Second, some states cap public adjuster fees, which reduces the financial incentive to expand the claimed scope of damage. Third, states differ in the strength of fraud investigation and enforcement (e.g., mandatory insurer Special

Investigation Units or active state fraud bureaus), which increases the expected cost of submitting exaggerated claims. For these three categories, we define indicator variables equal to 1 when the regulation is present.

The remaining two regulatory categories operate through information asymmetry. Some states restrict the use of prior claims history or credit information in pricing or renewal decisions. By limiting insurers' ability to condition premiums on observable risk indicators, these regulations weaken dynamic incentives to avoid marginal or inflated claims. For these two categories, we define indicator variables equal to –1 when the restriction is present.

	(1)	(2)	(3)
	Moral Hazard	Moral Hazard	Moral Hazard
MH Reducing Regulation	-1.228***	-1.062***	-1.026***
	(-106.49)	(-79.55)	(-12.99)
Tail Risk (AEP 100)		-1.065***	-1.015***
		(-55.25)	(-14.04)
FICO Score		-0.260***	-0.225***
		(-17.21)	(-7.34)
Log Recovery Value		-1.695***	-1.598***
		(-48.87)	(-10.17)
Constant	7.891***	29.91***	28.62***
	(457.53)	(66.25)	(13.98)
Insurer FE	N	N	Y
N	5,926,005	5,926,005	5,925,999

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Moral Hazard vs. Regulations Reducing Moral Hazard

Our moral-hazard regulation index is defined as the sum of these five indicator variables. Accordingly, higher values correspond to regulatory environments that make moral hazard more costly. By regressing the cost of moral hazard on the regulation index controlling for the same set of control variables and insurer fixed effects, we find that the coefficient on the regulation index is negative, as reported in Table 8. This confirms that

regulations that are associated with reducing moral hazard, indeed, reduces the cost of moral hazard.

While not intended to establish causality, the results in this section add validity to our estimates.

7 Counterfactual Analysis and Policy Implications

Our results so far have shown that policyholders, particularly those exposed to tail risk and those with low FICO scores, retain considerable risk even after being insured. This prompts a policy-relevant question: what if insurers are required to offer full coverage?

Our structural model is useful for answering this question. In this section, we focus on one application of it to consider a hypothetical regulation that mandates full coverage of insurance. In this case, insurers would not be able to incentivize effort from policyholders and be required to cover all the property damage under the policyholder's negligence. If the expected damage under negligence exceeds the premium that the insurer can earn from a complete insurance, the insurer will choose to leave the market, resulting in market failure. In short, the regulation that aims to extend insurance coverage may leave policyholders entirely exposed to property risk.

We first estimate the counterfactual damage distribution under negligence, $f_0(x)$, using the model. Specifically, we obtain the likelihood ratio $LR = \frac{f_0(x)}{f_1(x)}$ from Equation (FOC). Then, we recover $f_0(x)$ using the LR and the $f_1(x)$ that we had estimated earlier. This allows us to infer the distribution of property damage when policyholders are not incentivized to properly manage their properties.

The median increase in expected damage is $\$1,107$; the damage increases by the factor of 3.3 under negligence relative to effort. In the counterfactual world where insurers are forced by regulation to provide full insurance at the premium that the policyholder would be willing to pay, the insurer would choose to withdraw from the market. For

approximately 46% of the contracts, the expected claim payout, $E[x|a = 0]$ exceeds the premium that the policyholder would be willing to pay for a complete insurance, which is p^{FB}. For the other 54%, the insurer still finds it profitable to provide full insurance. Our result suggests that taking away the insurer's ability to provide incentives would leave a substantial portion of households uninsured, exposing them entirely to property risk.

8 Conclusion

Housing represents the largest component of household wealth in the United States, making property risk a key financial for households. Losses due to natural disasters often exceed the liquid savings of most households. Therefore, property insurance plays a critical role in smoothing these shocks. Yet the protection households ultimately receive from insurance depends crucially on contract design, specifically, the deductible and coverage limit that determine how much loss remains with the household. These features, designed to balance risk sharing and incentives, shape how disaster shocks propagate through household balance sheets and, by extension, the broader economy.

Drawing on rich contract-level data matched to property-level disaster risk, this paper presents one of the first comprehensive analyses of the structure of property insurance contracts and the economic forces underlying their design. We develop a structural model linking household risk aversion, property risk, and insurer pricing to observed contract terms.

Our estimation results show that insurance contracts successfully mitigate moral hazard—the estimated cost of moral hazard is only 0.7% of the risk premium—but this incentive alignment comes at the cost of substantial uninsured exposure, roughly 29% of total expected losses on average. The burden of this retained risk falls disproportionately on low-FICO households and those in high-risk regions, suggesting important distributional consequences. A counterfactual analysis further shows that mandating full insurance

would lead many policyholders to lose coverage altogether, highlighting the trade-off between incentive alignment and risk sharing.

References

Arrow, Kenneth J., "Uncertainty and the Welfare Economics Of Medical Care," *The American Economic Review*, 1965.

Aunon-Nerin, Daniel and Paul Ehling, "Why firms purchase property insurance,," *Journal of Financial Economics*, 2008.

Bertomeu, Jeremy, Ivan Marinovic, and Mario Milone, "Explaining the Compensation of Superstar CEOs: The Case of Elon Musk," *Stanford University Graduate School of Business Research Paper No. 4863514*, 2025.

_ , Kyle (Jaehoon) Jung, and Ivan Marinovic, "Moral Hazard and the Value of Information: A Structural Approach," *Working Paper*, 2025.

Bhutta, Neil and Benjamin J Keys, "Moral Hazard during the Housing Boom: Evidence from Private Mortgage Insurance," *The Review of Financial Studies*, 05 2021, *35* (2), 771–813.

Blonz, Joshua A., Mallick Hossain, and Joakim Weill, "Pricing Protection: Credit Scores, Disaster Risk, and Home Insurance Affordability," *Working Paper*, 2024.

Campello, Murillo, John R. Graham, and Campbell R. Harvey, "The real effects of financial constraints: Evidence from a financial crisis," *Journal of Financial Economics*, 2010, *97* (3), 470–487. The 2007-8 financial crisis: Lessons from corporate finance.

Cohen, Alma and Liran Einav, "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, June 2007, *97* (3), 745–788.

Cookson, J. Anthony, Emily Gallagher, and Philip Mulder, "Coverage Neglect in Homeowners Insurance," *Working Paper*, 2024.

Czajkowski, Jeffrey, Kevin M. Simmons, and James M. Done, "Demonstrating the Intensive Benefit to the Local Implementation of a Statewide Building Code," *Risk Management and Insurance Review*, 2017, *20* (3), 363–390.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen, "Estimating Welfare in Insurance Markets Using Variation in Prices*," *The Quarterly Journal of Economics*, 08 2010, *125* (3), 877–921.

_ , _ , Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen, "Selection on Moral Hazard in Health Insurance," *American Economic Review*, February 2013, *103* (1), 178–219.

Federal Insurance Office, "US Department of the Treasury Report: Homeowners Insurance Costs Rising, Availability Declining as Climate-Related Events Take Their Toll," *US Department of the Treasury Report*, 2025.

Federal Reserve Board, "Economic Well-Being of U.S. Households in 2024," *Research & Analysis Report*, 2025.

Frazzini, Andrea and Lasse Heje Pedersen, "Betting against beta," *Journal of Financial Economics*, 2014, *111* (1), 1–25.

Gayle, George-Levi and Robert A. Miller, "Has Moral Hazard Become a More Important Factor in Managerial Compensation?," *American Economic Review*, December 2009, *99* (5), 1740–69.

_ and _, "Identifying and Testing Models of Managerial Compensation," *The Review of Economic Studies*, 01 2015, *82* (3), 1074–1118.

_, Chen Li, and Robert A. Miller, "Was Sarbanes–Oxley Costly? Evidence from Optimal Contracting on CEO Compensation," *Journal of Accounting Research*, 2022, *60* (4), 1189–1234.

Ge, Shan, "How Do Financial Constraints Affect Product Pricing? Evidence from Weather and Life Insurance Premiums," *Journal of Finance*, 2021.

_, Stephanie Johnson, and Nitzan Tzur-Ilan, "Rising Insurance Premiums Increase Mortgage Delinquency and Drive Relocation to Safer Areas," *Working Paper*, 2024.

Handel, Benjamin R., "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 2013.

Hendren, Nathaniel, "Measuring Ex Ante Welfare in Insurance Markets," *The Review of Economic Studies*, 06 2020, *88* (3), 1193–1223.

Holmström, Bengt, "Moral Hazard and Observability," *The Bell Journal of Economics*, 1979, *10* (1), 74–91.

Insurance Information Institute, "Facts + Statistics: Homeowners and Renters Insurance," 2025.

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

Jung, Hyeyoon, "Real Consequences of Shocks to Intermediaries Supplying Corporate Hedging Instruments," *Review of Financial Studies*, 2025.

_, Robert F. Engle, Shan Ge, and Xuran Zeng, "Physical Climate Risk Factors and an Application to Measuring Insurers' Climate Risk Exposure," *FRB of New York Staff Report No. 1066*, 2025.

Jung, Kyle (Jaehoon), "Green Moral Hazard: Estimating the Financial and Non-financial Impacts of CEO Incentives," *New York University Stern School of Business Research Paper Series*, 2025.

Keys, Benjamin J. and Philip Mulder, "Property Insurance and Disaster Risk: New Evidence from Mortgage Escrow Data," *Working Paper*, 2024.

Koijen, Ralph S. J. and Motohiro Yogo, "The Cost of Financial Frictions for Life Insurers," *American Economic Review*, 2014.

_ **and** _ , "Shadow Insurance," *Econometrica*, 2016.

LexisNexis Risk Solutions, "U.S. Home Insurance Trends Report Highlights Rising Peril Severity as Catastrophic Claims Hit Seven-Year High," 2025.

Margiotta, Mary M. and Robert A. Miller, "Managerial Compensation and the Cost of Moral Hazard," *International Economic Review*, 2000, *41* (3), 669–719.

Mirrlees, James, "Notes on welfare economics, information and uncertainty," *Essays on economic behavior under uncertainty*, 1974, pp. 243–261.

Mossin, Jan, "Aspects of Rational Insurance Purchasing," *Journal of Political Economy*, 1968, *76* (4), 553–568.

Oh, Sangmin, Ishita Sen, and Ana-Maria Tenekedjieva, "Pricing of Climate Risk Insurance: Regulation and Cross-Subsidies," *Journal of Finance, Forthcoming*, 2025.

Paravisini, Daniel, Veronica Rappoport, and Enrichetta Ravina, "Risk Aversion and Wealth: Evidence from Person-to-Person Lending Portfolios," *Management Science*, 2017, *63* (2), 279–297.

Perez-Gonzalez, Francisco and Hayong Yun, "Risk Management and Firm Value: Evidence from Weather Derivatives," *Journal of Finance*, 2010.

Sastry, Parinitha, Ishita Sen, Ana-Maria Tenekedjieva, and Therese C. Scharlemann, "The Limits of Insurance Demand and the Growing Protection Gap," *Working Paper*, 2025.

Sydnor, Justin, "(Over)insuring Modest Risks," *American Economic Journal: Applied Economics*, October 2010, *2* (4), 177–99.

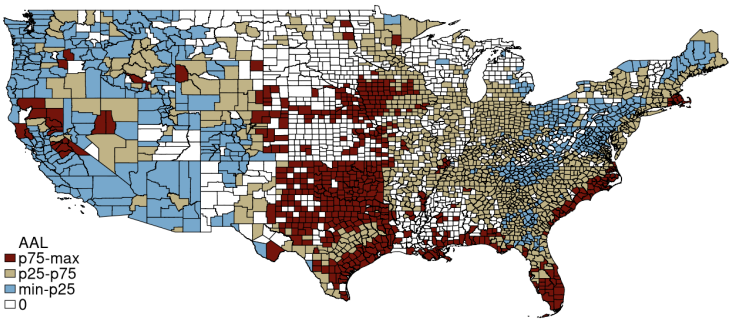
Vera-Hernandez, Marcos, "Structural estimation of a principal-agent model: Moral hazard in medical insurance," *The Rand journal of economics*, Winter 2003, *34* (4), 670–693.

Wagner, Katherine R. H. and Kendra Marcoux, "Fifty Years of U.S. Natural Disaster Insurance Policy," in Georges Dionne, ed., *Handbook of Insurance*, Vol. I, Springer Nature, 2024, pp. 55–79.

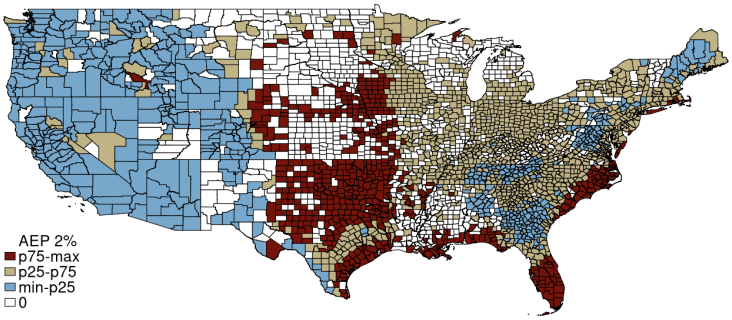
Woleben, Jason, "In Industry First, US P&C Insurers Exceed $1 Trillion in Direct Annual Premiums," *S&P Global Market Intelligence*, March 2025.

Appendix

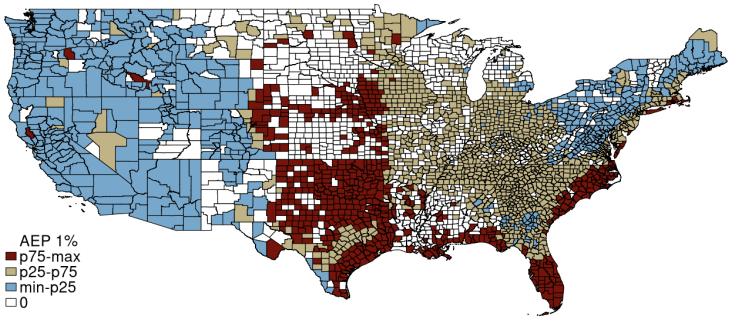
A Additional Figures



(a) Average Loss (AAL)

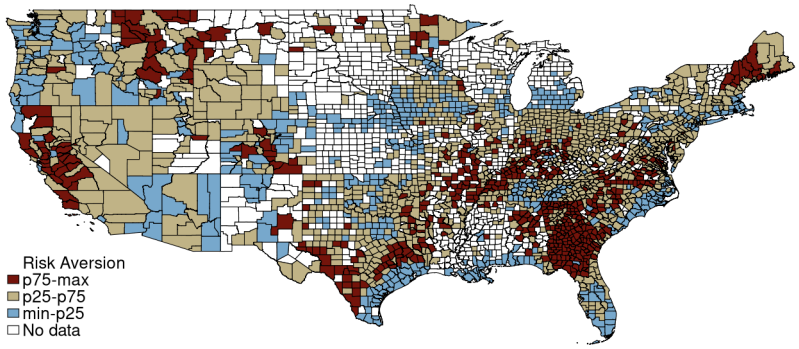


(b) 98 Percentile Loss (AEP 2%)

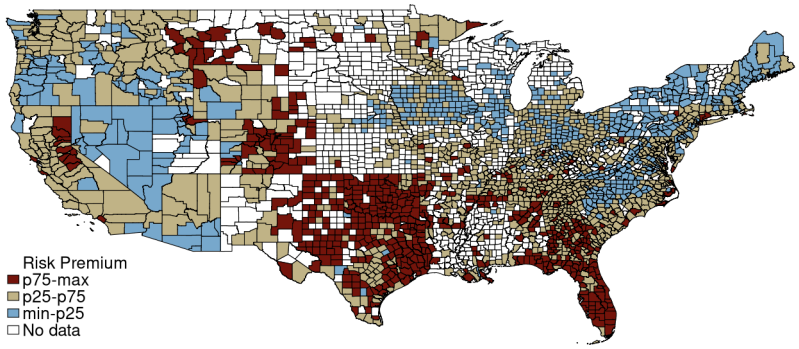


(c) 99 Percentile Loss (AEP 1%)

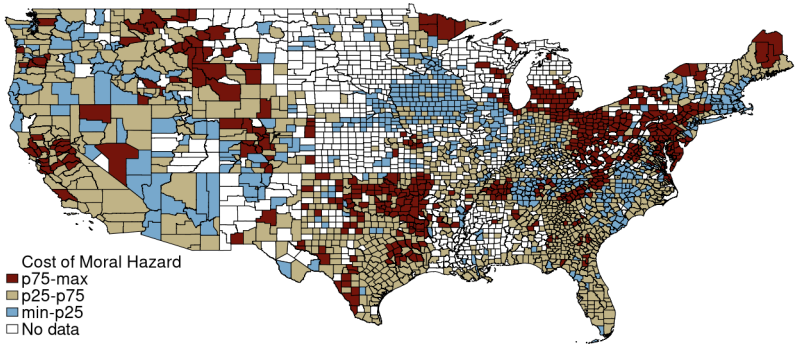
Figure A.1: Disaster Risk Map These figures show spatial distribution of the three disaster risk moments, mean, 98 percentile, and 99 percentile. The data source is CoreLogic.



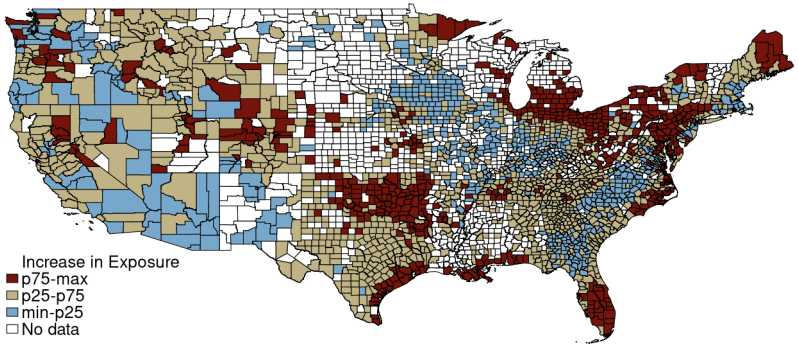
(a) Risk Aversion



(b) Risk Premium



(c) Cost of Moral Hazard



(d) Increase in Exposure due to Moral Hazard

Figure A.2: Spatial Distribution of Estimated Variables