Optimal Flood Insurance in a Second-Best World: Fiscal Spillovers, Reclassification Risk and Moral Hazard*

Jonathan Gruber
MIT & NBER

Adam Solomon
NYU Stern

September 15, 2025

Abstract

The National Flood Insurance Program's move to actuarially fair premiums corrects historical subsidies and reduces moral hazard, but has two unintended consequences: a fiscal spillover onto federal disaster aid as insurance coverage declines in flood-prone areas, and household exposure to uninsurable reclassification risk from uncertain climate projections. We find that each 1% increase in insurance coverage saves $203 per house in ex post federal aid when a flood occurs. We then use variance across climate models to estimate a household willingness-to-pay to avoid climate risk uncertainty of $65 annually. In contrast, reductions in moral hazard from correctly priced climate risk, through both migration and mitigation, are small. We estimate an optimal subsidy to flood insurance of 57%, comparable to the level of subsidization before RR2.0.

Keywords: flood insurance, disaster aid, climate risk, moral hazard, reclassification risk

*We thank Eileen Sadati for excellent research assistance. We are grateful to Ben Collier & Philip Mulder for helpful comments.

1 Introduction

The increased volatility of weather patterns around the world due to climate change has highlighted the essential role of natural disaster insurance. In the U.S. alone, over the last 10 years (2015-2024), there has been approximately $1.4 trillion of property damage from natural disasters (NOAA, 2025). Damage from random weather events is exactly the type of major risk against which individuals should naturally want to insure. Indeed, insurance against natural disasters is big business: the total premiums paid for homeowners insurance in 2024 were estimated at $172 billion (NAIC, 2024).

Yet, in the U.S., insurance against weather-induced damage is far from complete. Estimates suggest that 70% of the damage caused by floods is uninsured ex ante (Federal Reserve Bank of Philadelphia, 2016). But not necessarily ex post, as the U.S. government spends large amounts assisting individuals, businesses, and governments in the wake of disasters, primarily through Federal Emergency Management Agency (FEMA) assistance programs. The total ex post federal disaster relief obligations in the U.S. in fiscal year 2024 were $52 billion (U.S. Federal Emergency Management Agency, 2025), which is approximately 30% of the homeowners insurance premium spending in that year. This insurance gap is particularly acute for flooding: despite federal mandates for high-risk zones, we calculate that nationwide flood insurance coverage is 5%, exposing taxpayers to over $1 billion of ex post aid costs per year; even weighting by flood risk, the coverage rate is only 18%, and the highest risk areas have coverage rates below 25% (Amornsiripanitch et al. (2024)).

The market for disaster insurance is inhibited by multiple, cross-cutting, market failures. The first is the "Samaritan's Dilemma" induced by the availability of ex post coverage for disaster-related damage. As a result of these ex post promises, individuals may incompletely cover themselves ex ante with insurance. Efforts to resolve this incomplete coverage have failed, resulting in the low coverage rates noted above.

Historically, a second market failure—moral hazard—was a direct consequence of policy aimed at mitigating the first. Flood insurance has been provided by the public National Flood Insurance Program (NFIP), which for decades set premiums that were more than 50% lower than expected costs and which did not map to the true underlying flood risk, with the largest subsidies for the most flood-prone areas (Congressional Budget Office (2017)). This implicit subsidy, while encouraging insurance take-up, masked the true cost of risk and thus promoted inefficient over-development and under-mitigation in flood-prone areas (Fabian (2024), Ge et al. (2022)).

The NFIP's 2021 "Risk Rating 2.0" (RR2.0) reform, the program's first major overhaul in 50 years, undoes this historical subsidy by moving to actuarially-sound premiums. The new methodology integrates geospatial data and advanced catastrophe models to estimate property-level risk. This policy has been praised[1] for potentially offsetting the incentives to overbuild – and to insuffi-

[1] "By communicating flood risk more clearly, the new methodology should help policyholders make more informed decisions on the purchase of adequate insurance and on mitigation actions to protect against flooding." - Association of State Floodplain Managers, ASFPM (2021)

ciently remediate – in flood-prone areas.

But this reform is a classic example of the theory of the second best: fixing one market failure may not improve welfare in the presence of others. This paper argues that RR2.0 creates two significant unintended consequences. First, higher flood insurance premiums lead to lower coverage, which increases ex post recovery spending. Second, more tightly tying premiums to flood risk in a world of enormous uncertainty about the geographic distribution of coming flood risks imposes a welfare cost on risk-averse households through reclassification risk.

These offsetting market failures have not previously been explored or evaluated relative to the costs of subsidies. In this paper, we provide the first comprehensive measurement of the costs and benefits of subsidizing flood insurance.

We begin with a model for the optimal insurance subsidy in the presence of moral hazard, spillovers to ex post assistance, and dynamically evolving climate risk. Households face two forms of risk: the static uncertainty about whether a flood will occur or not, and the dynamic uncertainty as to whether their property will become higher or lower risk in the future. While typical models of optimal insurance focus on the former, we hold the generosity of insurance fixed and analyze the latter. Households choose whether to buy insurance, the amount of costly mitigation to undertake, and whether to locate in a higher or lower risk location. Anticipating this, the planner chooses a level of premium subsidy to trade off the gains from reclassification-risk insurance and the reduced spillover to ex post assistance against the induced moral hazard from distorted coverage, mitigation, and location choices.

The benefits from a premium subsidy are twofold: coverage increases thereby reducing future ex post assistance, and a portion of dynamic premium risk is insured. Our focus on the Samaritan's dilemma inverts the typical existing analyses: instead of quantifying how much ex post assistance crowds out insurance, we measure how subsidizing insurance reduces the fiscal externality from ex post assistance. In terms of reclassification risk, we are careful to distinguish and measure genuinely diversifiable risk from aggregate uncertainty across climate states that is uninsurable in the cross-section.

The cost of a premium subsidy is moral hazard: household actions that are not chosen optimally from a social point of view. Households make their insurance, mitigation, and location decisions trading off only the private costs and benefits. They do not consider costs to the planner, in particular, increased subsidy outlays or ex post assistance. Put another way, under a substantial subsidy, households are immunized from a portion of their true risk, and thus make decisions that increase that risk beyond the social optimum. It is this moral hazard that RR2.0 was intended to correct.

This model allows us to measure the optimal flood insurance structure as a function of five key reduced-form statistics: the fiscal spillover from NFIP coverage onto FEMA aid; the elasticity of demand for NFIP coverage with respect to its price; the certainty equivalent of the utility loss from higher reclassification risk; the elasticity of locational decisions with respect to insurance prices; and the elasticity of flood remediation with respect to insurance prices. We then turn to estimating these key empirical quantities.

First, we measure the fiscal spillover from the decrease in NFIP coverage onto FEMA aid. To do so, we assemble a county–year dataset combining flood insurance coverage, flood damages (insured and uninsured), and ex post government flood spending. Our empirical strategy faces two challenges: counties with severe floods may have both higher insurance uptake and higher ex post spending, and other unobserved factors may jointly influence coverage and recovery spending. We address the first by flexibly controlling for overall flood damage, comparing otherwise similar floods that strike counties with different insurance coverage levels. To address the second, we draw on the insight of Gallagher (2014): insurance demand temporarily increases after a flood in one's own or in a neighboring county. We instrument for flood insurance coverage with whether there was a flood in the past five years in a neighboring county. We show in an event study framework that there are no pre-trends in this relationship and run a host of other robustness checks as well.

Applying these methodologies, we estimate a sizable offset of ex post government spending due to more ex ante flood insurance. We identify these effects from quasi-random timing between floods; we do not compare flooded to non-flooded counties. We estimate that each additional 1% of houses that are insured lowers ex post government spending by $175-$203, should a flood occur.

Second, we estimate the elasticity of demand for NFIP insurance. To do so, we turn to a new source of identification: the shift to RR2.0. Specifically, through a FOIA request we obtained, for each property, its current premium and the full-risk premium it is transitioning to. Hence, our identification is within-property, as households face increasing prices along this transition path, without knowing in advance the final full-risk price. We estimate an elasticity of demand of -0.32, which is in the center of existing estimates.

Third, we measure the increased reclassification risk from risk-based pricing. As we document, there is substantial uncertainty about not only the mean path of climate change but also in the variance of that path across places. The more flood insurance prices are a function of local climate risk, the more households are exposed to premium volatility. This parallels the discussion of reclassification risk in health insurance markets (Handel et al. (2015), Ghili et al. (2023)): while more actuarially fair pricing sends correct price signals, it adds dynamic risk. Subsidies socialize some of this dynamic risk.

To quantify the reclassification risk households now face under RR2.0, we combine the universe of NFIP policies with state-of-the-art climate projections (Gori et al. (2022)). We estimate an empirical model of NFIP premiums, regressing RR2.0 premiums on a granular set of property characteristics and current flood risk factors. Holding property characteristics constant, we then use this model to project what premiums would be in 2070 under the eight distinct climate models (all calibrated to the same total level of warming) analyzed by Gori et al. (2022). In other words, we proxy for cross-sectional uncertainty, holding aggregate risk fixed, about local climate risk with the variability across climate models.

We find that there is substantial uncertainty in future projections. The average household faces a standard deviation of $3,317 in 2070 across the eight models. The premium risk households face has an aggregate component (all premiums are expected to rise) and an idiosyncratic one

(some properties' expected risk rises far more than others'). Community rating through a subsidy, i.e., undoing risk-based pricing, effectively insures the latter. We compute households' WTP for insurance against idiosyncratic reclassification risk. We find that the average household WTP against idiosyncratic reclassification risk is approximately $65. This represents roughly 3.6% of the average current premium and 1% of the expected 2070 premium. Finally, we estimate the impact of premium changes on moral hazard. To do this, we estimate the behavioral response of households to premium changes along two margins: long-run relocation decisions and short-run property-level mitigation efforts. These estimates allow us to quantify the primary intended benefit of RR2.0—reducing inefficient over-development and under-mitigation in flood-prone areas.

To quantify the effect of RR2.0 on migration out of high-risk areas, we use quarterly U.S. Postal Service data on occupied residential addresses at the census-tract level. We compare tracts where the median property saw a premium increase under RR2.0 to those in which the median property's insurance premium decreased. Our event study demonstrates that these tracts were on similar pre-trends prior to RR2.0, but that subsequently 0.86% of the population moved from the higher risk locations to the lower risk locations. For reference, Boustan et al. (2020) find that a typical natural disaster generates migration of about 1.5%. As such, finding almost half of this effect simply from an insurance price change is a relatively large effect. We use this estimated elasticity to quantify the welfare gains from the more efficient locational decisions induced by RR2.0.

To measure the effect of changed mitigation decisions, we use the universe of "Elevation Certificates" (EC) filed in Florida since 2017. An EC must be submitted prior to building to certify compliance with flood-zone-related building practices, and elevating a property is the primary form of risk-reduction against floods. Since these requirements differ inside and outside the high-risk flood zone, we compare those that faced a higher versus a relatively lower price increase inside the high-risk flood zone, and separately outside the high-risk flood zone. Properties within the high-risk zone are already subject to elevation requirements that were unchanged by the removal of premium subsidies. Accordingly, we find no differential effect on elevation by risk inside the high-risk flood zone. In contrast, outside of this zone, there are no requirements to elevate, potentially allowing individual decisions to be more sensitive to pricing incentives. Indeed, we find that properties in moderate-risk areas that experienced larger premium increases saw a 10% increase in the probability of being elevated by more than three feet compared to low-risk properties.

Having estimated these five key parameters, we then return to our model of optimal flood insurance subsidies. We find that, given the much higher FEMA spillover and reclassification risk reduction than the mitigation and migration moral hazard effects, the optimal subsidy to flood insurance is quite high. Our central estimate, 57%, is similar to the pre-RR2.0 level of subsidization. Interestingly, we find that this estimate is quite robust to demand elasticities (since both sides are affected linearly), that redistributive impacts are not significant, and that allowing for heterogeneity in our estimated costs and benefits only strengthens our conclusion, and that even applying an extreme set of assumptions on key parameters we obtain a lower-bound subsidy rate of 30%.

Our paper proceeds as follows. Section 2 provides background on flood insurance and a review

of previous literature on this topic. Section 3 sets out a model of optimal dynamic flood insurance with location and mitigation choices. Section 4 estimates the spillover to ex post FEMA spending from decreased NFIP coverage, while Section 5 estimates the demand elasticity for flood insurance. Section 6 then explains and explores the magnitude of welfare loss from reclassification risk with actuarially fair premiums. Sections 7 & 8 explore the two margins of moral hazard, location and mitigation. Section 9 puts all of these estimates together to calculate the optimal subsidy. Section 10 concludes.

2 Background

2.1 Flood Insurance in the U.S.

The vast majority of flood insurance in the U.S. is provided by the National Flood Insurance Program (NFIP). Launched in 1968 and administered by FEMA, the NFIP has roughly 4.7 million policies in force that provide coverage of $1.28 trillion (Federal Emergency Management Agency (2023)). This accounts for fewer than 1 in 20 households, although insurance penetration reaches up to 40% in certain coastal locations.

The NFIP issues a single type of policy: the Standard Flood Insurance Policy (SFIP). The SFIP pays for actual damages caused by floods up to residential coverage limits of $250,000 for the structure and $100,000 for personal contents. Most policies are written and claims administered by a "Write-Your-Own" private insurer, but contracts are set and the risk is ultimately borne by the NFIP.

Historically, there has been no non-negligible private flood insurance market, likely because of the subsidized rates offered by the NFIP. However, in recent years, as the NFIP has changed its pricing procedure to better capture property-specific risk, and also because of the relatively low coverage limits of the NFIP, the private flood insurance market has grown. As of 2023, approximately 350,000 policies were written by private insurers (National Association of Insurance Commissioners (2024)). This small, but growing, market is essentially irrelevant to the period we study (2010 to 2021). In a world with private coverage that mimics public coverage, our policy prescription would extend to market-wide subsidies; but if private coverage is more generous than public coverage—and in particular exceeds the limits reimbursed by FEMA—optimal private subsidies may be smaller.

2.2 Ex-Post Flood Recovery Spending

Federal ex post assistance for flooded households arrives mainly through three channels. First, there is FEMA's Individuals and Households Program (IHP), which provides grants for emergency repairs, temporary housing and limited personal-property losses. These grants typically cover only a fraction[2] of typical rebuilding costs. Second are SBA disaster home loans, which let owner-

[2] In our sample, the (weighted) average IHP payout was $20k against total damage of approximately $144k.

occupants borrow up to $500,000 at subsidized rates to restore their primary residence and up to $100,000 for damaged contents; these loans must be repaid and cannot be taken against damages covered by flood insurance. Third, after the acute response phase, communities can apply to FEMA's Hazard Mitigation Grant Program (HMGP) to finance buyouts, elevations or retrofits for individual homes.

Because federal law bars duplication of benefits, each of these programs is explicitly offset by any insurance proceeds. In practical terms, a homeowner with an NFIP policy may see their IHP grant reduced to zero and will borrow less—or nothing—from SBA, while widespread coverage in a disaster county lowers total IHP outlays and therefore the HMGP formula allocation. The average NFIP claim payment between 2016 and 2021 was about $66,000, roughly an order of magnitude larger than the typical IHP grant, so even modest increases in pre-event coverage can translate into sizable federal savings ex post (Congressional Research Service (2023)).

To the extent that owners view ex post aid as dependable, it can crowd out ex ante insurance purchase, a point discussed in the literature review section below. Yet the same ex post approach insures against reclassification risk: owners outside mapped flood zones who never purchase NFIP coverage can still receive federal grants if climate change or new flood maps suddenly expose them to increased flood risk. Our analysis asks how large the fiscal trade-off is—how much these household-facing programs shrink when more properties buy insurance up front—and whether that offset meaningfully defrays the cost of any premium subsidies.

2.3 Related Literature

A large body of work in economics seeks to explain the "disaster insurance puzzle"—why rational agents fail to insure against low-probability, high-consequence events (Kunreuther and Pauly (2004)). For flood insurance, the leading explanation is that risk perceptions are myopic and driven by salience. Gallagher (2014) shows that direct or nearby experience of a flood dramatically increases insurance take-up, an effect that decays over time. This finding that households learn from disasters is reinforced by subsequent work (Kousky (2010)). Other research explores how adverse selection and adaptation incentives interact (Wagner (2022)) and how new risk-based pricing regimes affect demand when risk information is not clearly provided (Mulder and Kousky (2023)). Furthermore, studies investigating the primary regulatory tool—the mandatory purchase requirement for federally-backed mortgages—find its effectiveness is limited by inconsistent compliance and lax enforcement by mortgage servicers (U.S. Government Accountability Office (2017, 2022)).

A second strand examines whether generous disaster assistance depresses subsequent insurance demand. Raschky et al. (2013) exploit cross-country variation in relief schemes to show that assured post-event aid crowds out private flood cover. In the United States, Kousky et al. (2018) match NFIP records to FEMA grant data and show that an individual-assistance grant reduces the next year's average insurance purchase on the intensive margin. Deryugina and Kirwan (2018) document a parallel "Samaritan's dilemma" in crop insurance, finding that farmers expecting large

federal bailouts both insure less and invest less in costly inputs. Philippi and Schiller (2024) demonstrate the tradeoff between ex ante subsidy and ex post disaster insurance theoretically, and exhibit empirical evidence from German winegrowers consistent with this tradeoff.

The link between ex ante insurance and ex post aid has been studied almost exclusively in one direction. The literature has documented how the promise of government relief can suppress insurance demand. Yet, to our knowledge, no empirical study has estimated the reverse effect: the direct, mechanical reduction in federal ex post spending that results from higher ex ante insurance coverage. This paper is the first to provide a credible estimate of this trade-off.

We also focus on reclassification risk in flood markets. Annual contracts are well known to expose policyholders to reclassification risk (Cochrane (1995); Hendel and Lizzeri (2003)). While the magnitude of this risk has been quantified primarily in health-insurance markets Handel et al. (2015); Ghili et al. (2023), our contribution is the first to leverage state-of-the-art climate models to estimate reclassification risk in the context of natural catastrophes, and to assess whether socializing that idiosyncratic volatility through targeted NFIP premium subsidies can generate substantial welfare gains.

Finally, we provide direct evidence for moral hazard in terms of mitigation and location choices due to price changes in the context of climate-risk-exposed home insurance. Existing evidence shows that lower flood insurance subsidies reduced new development (Fabian (2024)), and the availability of insurance stimulates population movements (Peralta and Scott (2024), Browne et al. (2019)). However, some studies (Hudson et al. (2017), Kousky (2019) find minimal or no effect of insurance prices on risk reduction. On the location margin, various papers find that natural disasters cause substantial population movements ((Boustan et al., 2020), (Aron-Dine, 2025), (Deryugina et al., 2018)). Relatedly, Henkel et al. (2024) show that whether or not a disaster received a FEMA-enabling presidential declaration affects recovery and subsequent out-migration. To our knowledge, this is the first paper to identify the effects of changes in disaster-insurance prices - as opposed to eligibility or map changes - on both household location (migration) and structural mitigation (elevation).

3 A Model of Optimal Insurance Subsidies

In this section, we study a model of optimal insurance subsidies in the presence of moral hazard, dynamic reclassification risk, and the crowd out of ex post assistance. From this model, we derive sufficient statistics for the optimal insurance subsidy that we subsequently estimate.

The Household's Problem

Let $P_i(\omega)$ be the baseline, actuarially fair premium. The planner chooses a subsidy level s. This determines a post-subsidy premium that the household actually pays $\tilde{P}_i(\omega) = (1-s)\,P_i(\omega)$.

Given the subsidy, household i makes three choices: insurance $q_i \in [0,1]$, location $x_i \in [0,1]$ and mitigation $m_i \in [0, w)$. For expositional clarity, we model all three as continuous. In Section VII, when we calibrate the formula for optimal flood insurance to the data, we use the realistic, discrete analogues. Mitigation m costs $\kappa_i(m_i)$. Choosing to locate in x delivers dollarized amenity value $\Gamma_i(x_i)$.

The possible flood damage L_i depends on the household's choice of location and mitigation, the future climate scenario ω and the idiosyncratic flood realization (or not) ξ. Of the uninsured losses, we assume a portion $f > 0$ is covered by the government (i.e. FEMA) ex post. Finally, the government budget is balanced by some (climate state-specific) tax $T_i(\omega)$. Hence, household consumption, after climate state ω and idiosyncratic flood shock ξ realize, is given by

$$c_i(\omega, \xi) = y_i \; + \; \Gamma_i(x_i) \; - \; C_i(m_i) \; - \; q_i \tilde{P}_i(\omega; s, x_i, m_i) \; - \; (1-f)(1-q_i)\, L_i(\omega, \xi; x_i, m_i) \; - \; T_i(\omega).$$

The household chooses (x_i, m_i, q_i) to solve

$$V_i(s) = \max_{(x_i, m_i, q_i)} \mathbb{E}_\omega \mathbb{E}_\xi [u(c_i(\omega, \xi; \cdot))],$$

with u strictly increasing and concave, and V_i is the indirect utility.

The household faces two forms of uncertainty. First, the 'standard' risk from a flood, encoded by ξ. This is insured by q. Second, and novel to this setup, the household faces risk over their future premium due to an uncertain aggregate climate state ω. Some of this reclassification risk is implicitly insured via a subsidy s; at the limit, where $s = 1$ there is no future premium risk. However, because ω represents a global climate state, only a portion of the reclassification risk can be insured.

Premium Decomposition We can decompose each individual's premium price $P_i(\omega)$ into an individual-specific component, a scenario-specific component, and an idiosyncratic component:

$$P_i(\omega) = \bar{P}(\omega) + a_i + \varepsilon_{i\omega}$$

where $\bar{P}(\omega)$ is the average premium in scenario ω, $a_i = \mathbb{E}_\omega[P_i(\omega) - \bar{P}(\omega)]$ is the type-specific fixed effect, and $\varepsilon_{i\omega}$ is the residual with $\mathbb{E}_\omega[\varepsilon_{i\omega}] = 0$.

We want to quantify the value of a premium subsidy s in insuring the *idiosyncratic* component only. The scenario-specific risk $\bar{P}(\omega)$ affects the entire world at once, and so cannot be diversified. Similarly, 'insuring' predictable differences in individual-specific risk a_i is simply redistribution, which we do not focus on here.

The Planner's Problem

The planner chooses s to maximize

$$\max_{s} V_i(s) \quad \text{subject to} \quad \sum_i T(\omega) \;=\; \underbrace{\sum_i s\, q_i\, P_i(\omega)}_{\text{premium subsidy on insured}} \;+\; \underbrace{\sum_i f\,(1-q_i)\, P_i(\omega)}_{\text{expected FEMA on uninsured}} \qquad \forall \omega.$$

Note, we have assumed that a budget deficit is funded by a lump-sum tax such that individuals do not internalize the impact their actions have on the aggregate budget deficit. That is, moral hazard occurs through individuals' choices of insurance, mitigation and location. This assumes that the flood insurance and disaster aid program are funded by the same people they benefit; we do not analyze the distributional consequences of the average taxpayer possibly differing from the average disaster aid recipient.

General Formula for the Optimal Subsidy

To state the general condition for the optimal subsidy level s, we define the semi-elasticity of insurance choice q with respect to a change in the post-subsidy price $\mathcal{E}_i^q \equiv \frac{\partial q_i^*}{\partial \ln \mathbb{E}_\omega[\tilde{P}_i(\omega)]}$, and identically for $\mathcal{E}_i^x$ and $\mathcal{E}_i^m$. The following proposition presents a general formula with dynamic risk, both aggregate and idiosyncratic, state-specific adverse selection, and a decomposed fiscal externality from moral hazard.

Proposition 1 (Optimal Subsidy: General Condition)**.** *The optimal premium subsidy s solves:*

$$
\begin{aligned}
&\underbrace{(1-s)\sum_i \Big(\mathsf{Cov}_\omega(u_i',\ q_i^*\varepsilon_{i\omega})\ -\ q_i^*\,\mathsf{Cov}_\omega(\bar{u}'(\omega),\ \varepsilon_{i\omega})\Big)}_{\textit{idiosyncratic insurance}\ -\ \textit{macro–idiosyncratic correction}} \\
&\quad+\ \underbrace{(1-s)\sum_i q_i^* a_i\Big(\mathbb{E}_\omega[u_i'] - \mathbb{E}_\omega[\bar{u}'(\omega)]\Big)}_{\textit{risk-class }(a_i)\textit{ redistribution}} \\
&\quad+\ \underbrace{(1-s)\,\mathbb{E}_\omega[\bar{P}(\omega)]\cdot\mathbb{E}_\omega\Big[\sum_i q_i^*(u_i'(\omega)-\bar{u}'(\omega))\Big]}_{\textit{average selection}} \\
&\quad+\ \underbrace{(1-s)\,\mathsf{Cov}_\omega\Big(\bar{P}(\omega),\ \sum_i q_i^*(u_i'(\omega)-\bar{u}'(\omega))\Big)}_{\omega\textit{-specific selection}} \\
&\qquad=\ -\mathbb{E}_\omega\Bigg[\mu(\omega)\Bigg(\underbrace{\sum_i\left((s-f)\,P_i(\omega)\right)\mathcal{E}_i^q}_{\textit{coverage fiscal externality}} \\
&\qquad\qquad+\ \underbrace{\sum_i\Big(s\,q_i^*+f(1-q_i^*)\Big)\,\tfrac{\partial P_i(\omega)}{\partial m}\,\mathcal{E}_i^m}_{\textit{mitigation fiscal externality}} \\
&\qquad\qquad+\ \underbrace{\sum_i\Big(s\,q_i^*+f(1-q_i^*)\Big)\,\tfrac{\partial P_i(\omega)}{\partial x}\,\mathcal{E}_i^x}_{\textit{location fiscal externality}}\Bigg)\Bigg].
\end{aligned}
$$

The first term on the left – idiosyncratic insurance – is the 'standard' value of insurance term; it is the continuous analogue of the differences in marginal utility in the good versus bad state, as we show in Example 1 below. The remaining terms arise generically in this dynamic environment, but for reasons we explain below, we rule them out by assumption. The right hand side is an almost standard accounting of the fiscal externalities due to changes in household actions. However, when the government sets policy in the present, they average over climate states ω in the future. We make an additional assumption to additionally rule out any covariance between $\mu(\omega)$ (the shadow cost of public funds) and future climate states $P_i(\omega)$.

Assumptions

A1 $\mathsf{Cov}_\omega(\mu(\omega), L_i(\omega)) = 0$; dynamic-flood risk is small relative to the economy.

A2 $\mathsf{Cov}_\omega(\bar{u}'(\omega), \varepsilon_{i\omega}) = 0$; individual idiosyncratic risk is orthogonal to aggregate climate risk.

A3 $\mathbb{E}_\omega\Big[\sum_i q_i^*(u_i'(\omega) - \bar{u}'(\omega))\Big] = 0$; no selection on average.

A4 $\mathsf{Cov}_\omega(\bar{P}(\omega), \mathsf{Cov}_i(q_i^*, u_i'(\omega))) = 0$; no additional selection in bad aggregate climate states.

A5 $\mathsf{Cov}_i(a_i, \mathbb{E}_\omega[u_i'(\omega)]) = 0$; rules out redistribution on predictable individual risk.

Assumption A1 assumes that the shadow cost of public funds is orthogonal to the aggregate flood-risk scenario. To the extent flood risk is a 'small' component of the economy, this is justified. Estimating this covariance would require that we take a view on how differentially distortionary the marginal dollar of taxation is in 2070 under different climate scenarios.

Assumption A2 assumes that the individual idiosyncratic deviations $\varepsilon_{i\omega}$ do not systematically covary with the aggregate climate state, and in particular the cost of raising a dollar of (lump-sum) tax in climate state ω, which is $\bar{u}'(\omega)$. By construction, $\varepsilon_{i\omega}$ is zero in expectation over climate states ω. Hence, our assumption is that individual deviation from their predictable risk and the average climate risk is not correlated with aggregate climate shock. Empirically, this is true: the correlation is approximately 0.03 to 0.09 depending on the utility function specification.[3]

Assumption A3 assumes that there is no selection into insurance on average. The premiums are risk-rated and personally priced, so there is no externality from selection, were it to exist. Selection could be generated by individual constraints or (e.g., risk preferences that change their demand for insurance, over and above risk differences, which are already priced in. Critically, this assumption, while supported by the literature (e.g. Wagner (2022)) *would only strengthen* the case for subsidies if selection were adverse. A subsidy would incentivize insurance take-up amongst those with particularly high marginal utility, and therefore be even more desirable for the planner than if take-up was uncorrelated with marginal utility.

Assumption A4 rules out state-specific selection. The covariance term $\text{Cov}_i(q_i, u_i'(\omega))$ measures the degree of selection: do those who choose insurance have different risk levels and therefore marginal utility? We assume that this selection pattern does not covary with the level of aggregate risk in the future climate scenario $\bar{P}(\omega)$.

Assumption A5 rules out redistribution based on predictable flood risk. Since we wish to study the insurance value of a premium subsidy, we exclude transfers based on a known risk type. We explore the redistributive implications of flood insurance subsidies quantitatively in Section 9.

Under these assumptions, and further restricting insurance to be a binary choice, and assuming that mitigation and location choices only change (smoothly) among the insured (Assumptions A6-A7)[4], we arrive at the simplified optimality condition that we will estimate.

Proposition 2 (Optimal Subsidy: Simplified)**.** *Under assumptions A1-A7, the optimal premium subsidy s solves:*

$$\underbrace{(1-s)\sum_i \text{Cov}_\omega(u_i',\, \varepsilon_{i\omega})}_{\textit{idiosyncratic reclassification ins.}} + \underbrace{\mathbb{E}_\omega\left[\mu(\omega)\sum_i f\,\bar{P}_i\,\mathcal{E}_i^q\right]}_{\textit{FEMA savings}} =$$

$$-\,\mathbb{E}_\omega\left[\mu(\omega)\left(\underbrace{\sum_i s\,\bar{P}_i\,\mathcal{E}_i^q}_{\textit{coverage FE}} + \underbrace{\sum_i s\,\mathbb{E}_\omega\left[\tfrac{\partial P_i}{\partial m}\right]\,\mathcal{E}_i^m}_{\textit{mitigation FE}} + \underbrace{\sum_i s\,\mathbb{E}_\omega\left[\tfrac{\partial P_i}{\partial x}\right]\,\mathcal{E}_i^x}_{\textit{location FE}}\right)\right].$$

[3] See Appendix A.1 for more details.

[4] See the proof of Proposition 2 for formal definitions.

The right-hand side measures all the costs from an expanded subsidy: the fiscal externalities from more insurance, less mitigation, and riskier location choices. The fiscal externalities are proportional to s because the subsidy is exactly the quantity the individual doesn't consider when making their choices that the planner still has to pay . For example, when an individual chooses their level of mitigation, they trade off the private costs and benefits; they do not account for the reduced subsidy the government will pay if mitigation reduces their premium, nor the expected reduced FEMA costs because of a risk reduction. The logic is similar for location choice to the extent that risk changes with location. When choosing to insure or not, the household trades off the benefits of risk protection against the premium net of the subsidy, but not the cost of the subsidy itself, nor the ex post FEMA assistance should they go uninsured.

The left-hand side measures all the benefits of an expanded subsidy. The second term is the FEMA-cost saving from increased insurance. Although it is a fiscal saving, not an risk-smoothing gain, it is the primary benefit of expanded subsidies. The first term is the value of insurance against idiosyncratic (reclassification) risk. It is the continuous analogue of the difference in marginal utilities seen in the typical binary-outcome model as Example 1 shows.

Example 1. *Consider a two-state example with $\omega \in \{B, G\}$ (bad and good climate scenarios, probabilities $1-\pi$ and π). Idiosyncratic risk is high in B and low in G: $\varepsilon^B = \pi\Delta\varepsilon, \quad \varepsilon^G = -(1-\pi)\Delta\varepsilon$, so $\mathbb{E}_\omega[\varepsilon_\omega] = 0$ and $\varepsilon^B - \varepsilon^G = \Delta\varepsilon > 0$. Fix aggregate and person-specific risk. The idiosyncratic insurance on the left-hand side of the condition for optimal s^* is proportional to the marginal utility difference:*

$$(1-s)\mathsf{Cov}_\omega(u', q^*\varepsilon_\omega) = (1-s)q^*\left(\frac{u'^B - u'^G}{\Delta\varepsilon}\right)\mathit{Var}_\omega(\varepsilon_\omega).$$

In sum, we have derived a formula for the optimal subsidy with dynamic reclassification risk (both aggregate and idiosyncratic) and multiple fiscal spillovers that nests the standard insurance-moral hazard tradeoff. To solve the formula in Proposition 2, we need to estimate five numbers: the value of insurance against idiosyncratic reclassification risk for flood; the change in ex post FEMA spending as insurance prices, and therefore take-up, change; the elasticity of flood-insurance demand with respect to the price; the elasticity of mitigation with respect to an insurance subsidy change; and the elasticity of movement from risky to less risky places as insurance subsidies changes. In the next five sections, we estimate each of these quantities.

4 Spillovers to Ex Post Disaster Aid

Data on presidentially declared (flood) disasters (PDDs) comes from FEMA's Disaster Declaration Summaries. This dataset provides the official record for all federally declared disasters, detailing the incident type, date, and geographic scope. For our analysis, we restrict to the counties where assistance under the Individuals and Households Program (IHP) was authorized.[5]

[5]This should cover the vast majority of significant flooding–the cutoff for public assistance is $4.72 in damage per capita which is low compared to the mean damage of thousands of dollars in our sample.

To estimate disaster-specific damages, we use the National Oceanic and Atmospheric Administration (NOAA) Storm Events Database. We extract all events classified as floods and aggregate the total estimated property damage for each affected county. These county-level damage figures are then matched to a PDD by county and date, with a one-week buffer.

Data on FEMA's Individuals and Households Program (IHP), the National Flood Insurance Program (NFIP) policies and claims, and the Hazard Mitigation Grant Program (HMGP) come from the OpenFEMA data portal. Similarly, disaster loan information for the Small Business Administration (SBA) comes from data.sba.gov. We process these raw datasets to calculate each program's total spending by declared disaster and county.

Data on mortgage performance for Freddie Mac and Fannie Mae come from the GSEs public data releases. We use their quarterly loan-level performance data to assess costs incurred from defaulted mortgages. We measure fiscal costs in the year after each declared disaster. Freddie Mac directly provides an estimate of actual net costs due to foreclosure. For Fannie Mae, we calculate the total loss for each defaulted loan by summing the unpaid principal balance, delinquent interest, and various costs associated with foreclosure and property maintenance, and then subtracting any proceeds from the sale of the property or other credit enhancements.

We use demographic and housing data from the U.S. Census Bureau's American Community Survey (ACS). Additionally, the Census Bureau provides the geographic concordance tables used to map county and state FIPS codes to their respective names and to link ZIP codes to counties.

Our datasets cover different time periods. The most restrictive is the NFIP Policies-in-Force dataset, which begins in 2009. As such, we restrict our analysis to the period 2010 onward. However, for the purposes of certain backward-looking variables, such as time since a recent flood, we use data on disaster declarations and IHP spending going back to 2002.

4.1 Summary Statistics for Ex Post Aid

Table 1 presents the sample summary statistics, where each county-year is weighted by flood damage in that county/year. The average dollar of flood damage occurs in a county of 1 million people, and has county wide property damages of $4.7 billion. About one in five households have NFIP insurance; the damage done to a house by flooding averages $32,000. On average, across all houses in the county, NFIP payouts are $3440. Given that only 18% of houses are insured, if insured houses are representative, this suggests that NFIP offsets 60% of the costs of flooding for insured houses; this is likely a lower bound, however, as those houses that buy insurance are likely to be the most valuable.

Table 1 about here

We also find that there are sizable uninsured public expenditures associated with flooding. Disaster insurance per house is $733, while on average another $704 is provided in SBA loans, and $147 is provided in HMGP. Given that SBA loans have a government cost of 13c per dollar (CRS),

we estimate that the total ex post government spending averages $947 on a flood-weighted basis. Adding this to the $3440 in NFIP payouts, we find ex ante and ex post payments amount to 12% of total flood costs – meaning that about 88% of the costs of flood damage in America are uninsured either ex ante or ex post.

4.2 Methods and Results

We begin by estimating a two-way fixed effects model of the effect of flood insurance on ex post flood recovery spending:

$$\text{ExPost Spending}_{c,f} = \alpha_c + \gamma_t + \text{NFIP Coverage}_{c,f} + f\left(\text{Flood Damage}_{c,f}\right) + \epsilon \tag{1}$$

where $\text{ExPost Spending}_{c,f}$ is ex post spending per house in county c for flood f that occurs at time t, $\text{Coverage}_{c,t-2m}$ is NFIP coverage measured two months before the start of the flood event, $\text{Flood Damage}_{c,f}$ measures flood damage per house in county c for flood event f, α_c are county fixed effects, and γ_t are calendar-year fixed effects. We control for flood damage by binning[6] and estimating a distinct quadratic function within each bin. All variables are measured per house in the county.

As noted in the introduction, such a model faces two challenges. The first is that NFIP coverage and ex post spending are both related to flood risk. Moreover, mitigation activities that affect the damage from a flood are related to NFIP coverage and past floods. Hence, we flexibly control for flood damage—insured and uninsured—so that we are really comparing two otherwise similar floods that hit counties at times in which NFIP coverage differs.

Table 2 shows our OLS estimates for the impact of flood insurance on ex post spending. We estimate that insuring one more house per 100 would lead to a $175 reduction in ex post spending, which is highly significant, but is quite small relative to the total costs imposed by flooding. We estimate that 8% of houses make a FEMA claim. Assuming that these are the uninsured houses, and assuming that they also benefit from other spending (e.g., SBA loans), then this amounts to roughly $2,200 per house. If the average damaged house gets this assistance, this is only one-sixteenth of total damages.

The second empirical challenge is that it is possible that our OLS estimate is biased by factors that jointly determine insurance coverage and ex post spending. Places with more motivated owners might both buy insurance and more aggressively seek aid (biasing our estimate downward), while counties that can reliably secure recovery funds—perhaps for political reasons—may insure less (biasing the estimate upward). County fixed effects absorb time-invariant differences but not shifts (e.g., changes in political representation that reduce insurance demand by promising more aid). To address this, we implement an instrumental variables strategy, following Gallagher (2014). In particular, we exploit the short-term rise in flood insurance coverage in a county following floods in neighboring counties.

[6] Bins cutoffs at 1, 10, 100, 1,000, 10,000, 100,000, and 1,000,000 dollars of flood damage per house in the county.

We begin by replicating the impact of flood timing in neighboring counties on demand for flood insurance, in Figure 1. This graph shows the results from regressing NFIP coverage per house in county c on the time since a flood in the county c's closest neighbor, controlling for time since a flood in county c. We find that flood insurance demand picks up and remains elevated for the five years after a neighboring flood, rising by roughly 3 percentage points off a (weighted) mean of 18 percentage points, or approximately 17%. These estimates are very similar to the effects found by Gallagher (2014)), whose Figure 8 estimates a 3.1 percentage point increase in neighboring counties.[7]

Figure 1 about here

Given this relationship, it suggests a natural instrument for flood insurance demand: whether there was a flood in a neighboring county in the last five years. The specific timing of past floods in neighboring counties should be independent of insurance demand in own county, other than through its operation as a taste shifter for insurance demand.

Figure 2 shows the reduced-form relationship between time since neighboring flood and own ex post recovery spending. This figure mirrors Figure 1: there is no pre-trend, and in the roughly five years after a neighbor's flood, there is reduced spending on ex post recovery.

Figure 2 about here

The second column of Table 2 shows the results of our regression using neighbor-flood timing as an instrument. Doing so, we obtain a comparable estimate of $-203, which is statistically significant at the 1% level. As we decompose in the Appendix Tables A2 - A6, the majority of this – $150 – comes from FEMA's IHP. To benchmark this, FEMA payouts are capped at $85,000, but the average uninsured household receives approximately $14,500 in assistance relative to flood damage of $135,000. Hence, if people who take up NFIP coverage would otherwise cost FEMA $14,500, this rationalizes a FEMA saving of $150 per household as coverage increases by 1%.

And so if an extra percent of insurance coverage reduced FEMA payouts from this maximum to zero for these marginals, that would imply a $850 reduction in FEMA spending per total household. Similarly the maximal saving from SBA loans would be $650 per house[8] Accounting for these primary cost drivers caps the spillover effect of a percentage increase in NFIP coverage at about $\$1,500$ per house. Our effect is 13.5% of this upper bound.

Table 2 about here

Given these results, the only identification concern would be that there is an omitted variable which is jointly correlated with neighbor flooding, own insurance demand, and (in the opposite way)

[7] Note, Gallagher (2014) uses data from 1980–2007, earlier than our sample, and most analyses are done at the NFIP community, not county, level. His definition of neighbor is also based on media markets, while ours is based on geographic proximity of the counties' centroids. Gallagher also uses information on floods in own and neighboring counties; our results are robust to using both as well.

[8] The maximum loan size is $500,000, which costs the SBA (in terms of interest rate subsidy and expected default) $13\% \times \$500,000 = \$65,000$ per marginal household, or at most $650 per total household for a 1% coverage change.

own ex post flood spending. While this seems unlikely, we can go further to rule out underlying trends that drive this relationship by including a linear trend in time since flood in neighboring county c. We include these linear time-since-flood controls in Appendix Table A1, and our findings are robust, with an OLS estimate of -$175 and an IV of -$221.

In Appendix Tables A2 - A6, we decompose the spending impacts into the different programs included in our total. The largest and most significant effects come from FEMA's IHP program, with smaller and less precisely estimated fiscal savings coming from the HMGP and SBA loan programs. Fiscal savings from reduced foreclosures in Fannie Mae and Freddie Mac are negligible. We further show in the Appendix Table A7 that our primary estimates are robust to a wide variety of different ways to control for damage.

5 Demand Estimation

We use a policy-level panel of National Flood Insurance Program (NFIP) policies obtained via a Freedom of Information Act (FOIA) request. Relative to the public OpenFEMA release, our data include the *full risk-based* premium under Risk Rating 2.0 (RR2.0) for each policy and year, in addition to the billed premium. We restrict to policies priced under RR2.0 and construct a panel by matching policies on stable property identifiers, including geographic coordinates (latitude and longitude), census tract, flood map attributes, construction year, and building replacement cost.

Let $P_{i,t}$ denote the billed premium for policy i in year t and $R_{i,t}$ the full risk–based premium under RR2.0 applicable for year t. By regulation, NFIP premium increases cannot exceed 18% per year. As such, the renewal premium offered for the next period is $P^{\text{offer}}_{i,t+1} = \min\Big\{(1+0.18)\,P_{i,t},\ R_{i,t+1}\Big\}$. Define the *relative price* (subsidy ratio) as $\tilde{P}_{i,t} \equiv \frac{P^{\text{offer}}_{i,t+1}}{R_{i,t+1}}$.

Let $y_{i,t+1} \in \{0,1\}$ indicate whether policy i renews into year $t+1$. We estimate the following linear probability model with contract fixed effects:

$$y_{i,t+1} = \beta \ln \tilde{P}_{i,t} + \mu_i + \varepsilon_{i,t+1}, \tag{2}$$

where μ_i are policy (contract) fixed effects, absorbing all time-invariant heterogeneity. Standard errors are clustered at the policy level. Our identification comes from *within-policy* variation in the premium due to RR2.0.

The coefficient β is a *semi-elasticity*; the change in renewal probability with respect to a percentage change in renewal price, holding the full-risk premium constant:

$$\left.\frac{\partial \Pr(y_{i,t+1}=1)}{\partial \ln P^{\text{offer}}_{i,t+1}}\right|_{R_{i,t+1}} = \beta.$$

We additionally estimate two versions of (2) with additional heterogeneity terms. One with an interaction between the semi-elasticity and subsidy level, one with an interaction between the

semi-elasticity and an indicator for a house being in a SFHA. The results are in Table 3.

Table 3 about here

We estimate a semi-elasticity of demand of -0.32. This is roughly in the middle of the range estimated by the literature. Less elastic estimates include -0.19 to -0.28 (Lee (2024)) and -0.25 (Wagner (2022)); more elastic estimates are as high as -0.75 (Landry and Jahan-Parvar (2011)). When allowing for heterogeneity, we find an elasticity that increases (in absolute magnitude) as the subsidy gets smaller. That is, demand is more elastic for the first dollar of subsidy than the last. This result suggests that at least some small subsidy is likely to be optimal as it delivers particularly large FEMA savings. Moreover, we find lower elasticities (approximately -0.25) in SFHA relative to non-SFHA. This is likely due to the mandate that exists, but is imperfectly enforced, for households in the SFHA to buy and hold flood insurance for their mortgage to be guaranteed by Freddie Mac or Fannie Mae.

Additionally, in Appendix A.4 to complement the within-house estimates in this section, we analyze differences in demand responses across counties. Specifically, we compare counties differentially exposed to RR2.0 (in the same way as Section 8) and find, in an event study framework, an implied elasticity of -0.2, comparable to Lee (2024).

6 Reclassification Risk

The NFIP's transition to Risk Rating 2.0 (RR2.0) ties premiums directly to risk, present and future. While climate models broadly agree on the aggregate trend of increased future flood risk, they often produce divergent predictions for specific localities, even under the same aggregate warming scenario (Gori et al. (2022)). This creates two distinct forms of uncertainty for households. The first is aggregate uncertainty over the average increase in future risk. The second, which is the focus of this section, is idiosyncratic reclassification risk: the uncertainty a household faces about its premium relative to the average, driven by model disagreement over future local flood patterns.

While aggregate risk is largely unavoidable, this idiosyncratic risk is, in principle, diversifiable across the nation. Historically, the NFIP's subsidized, non-actuarial premiums implicitly socialized this risk. By moving to property-level actuarial rates, RR2.0 undoes this implicit insurance. Our goal is to quantify the magnitude of this newly exposed risk and estimate households' willingness to pay to insure against it.

6.1 Data on Future Flood Risk

To understand the uncertainty in future flood risk, we utilize synthetic flood data as developed by Gori et al. (2022). That paper generates a set of weather event 'seeds' that, under present climate conditions, generate a number of eventual climatic events that is well calibrated to historical flood frequency. Gori et al. (2022) then takes those same seed events, and runs them through eight

different models from the Coupled Model Intercomparison Project (CMIP) of the climate as it is projected to be in 2070. To ensure a consistent basis for comparison, each climate model is applied to the same aggregate climate scenario: SSP5-8.[9] This produces, for each model and for each location, a variable number of flood events. We record, in these synthetic future climate events, the risk of varying levels of precipitation, storm surge and simultaneous precipitation and surge risk.

Since NFIP policy locations do not perfectly align with the climate data grid, we assign hazard values to each policy based on Inverse Distance Weighting (IDW) using the closest three points on the climate data grid. For each grid point, we use three key hazard metrics: 100-year return period rainfall level (in millimeters); 100-year return period storm surge water level (in millimeters); the annual probability of a joint-100-year (rain and surge) flood event. Panels A and B of Table 4 summarize the pre-RR2.0 and RR2.0 premiums, and the current hazard metrics.

Mapping Premiums to Risk Factors. We have data on future hazard factors, but not directly on future premiums. To project future premiums, we generate a mapping from hazard factors to premiums using current data, and then use that mapping to predict future premiums given future hazards. We estimate, for property i's RR2.0 full-risk premium,

$$\text{Full Risk Premium}_i = \beta(\text{Rain, Surge, Joint})_i + \gamma(\text{Building Controls})_i + \epsilon_i$$

where we control for a rich vector of granular property characteristics, including building replacement cost, elevation, foundation type, number of floors, construction date, and occupancy type. The estimates for β are in Table A9. An extra millimetre of rainfall increases premiums by $\$3.93$; an extra millimeter of storm surge increases premiums by \$1.07; an increase in the probability of a joint storm and rainfall event by 1% increases expected premiums by \$367.63.

6.2 Simulating Future Premiums

No discounting. All amounts are in constant 2023 dollars. That is, our future premiums are calculated using 2070 levels of risk but with 2023 risk-to-premium factors. We do not apply any discounting to the 2070 premiums because premiums scale with building value and, as such, increase with house prices. Long-run real house-price growth averages about 1–2% per year (Jordà et al., 2019), close to the 2.0% real social discount rate recommended for regulatory analysis (omb, 2023). These approximately offset. Regardless, as we show in Appendix A.2, our policy conclusions are not sensitive to these discounting assumptions. Intuitively, this is because reclassification risk insurance value is highest for the first dollar of subsidy and then declines approximately quadratically after that. In the range of subsidies that we are typically finding as optimal - 50 to 60% or so - the marginal impact from a bit more insurance against reclassification risk has decayed to be very

[9] SSP5-8.5 is at the higher end of the ensemble of climate scenarios considered by the Intergovernmental Panel on Climate Change, assumes high economic growth and energy demand and low mitigation, and 8.5 W/m^2 radiative forcing by 2100. The best estimate of the warming implied by SS5-8.5 is about 4.4 degrees Celsius relative to 1850-1900(IPCC (2021).

small, regardless of additional discounting.

Gori et al. (2022) simulate 2070 rain, surge, and joint-hazard measures under eight general circulation models (GCMs). For each policy and each GCM, we replace historical hazard values with the projected 2070 values for that GCM, using inverse distance weighting (IDW) from the three nearest grid points; all other property characteristics (replacement cost, elevation, etc.) are held fixed. We then use our estimated pricing regression (in 2023 dollars) to predict 2070 premiums.

Our results are summarized in Table 4. Panel A measures premiums prior to RR2.0, which are on average $\$913$. Panel B shows the effect of RR2.0 on premiums (they rise to $\$1,772$ on average) and the risk factors estimated for each property using current hazard levels. Our regression above, the results of which are in Table A9, is estimated on the premiums and risk factors in Panel B. Panel C then uses simulated future climate hazards from Gori et al. (2022) and the model trained on Panel B to project future premiums across eight climate scenarios. On average across the scenarios, rainfall hazard increases by a factor of 4.69, storm-flooding risk by 45% and joint hazard by more than a multiple of 17. This leads to an average premium of $\$6,382$ in 2070 (denominated in 2023 dollars).

Table 4 about here

However, future premium uncertainty is substantial even under common aggregate warming assumptions. The average household faces considerable uncertainty around this future premium, with a standard deviation of $\$3,317$ across the eight climate outcomes (panel C). Future premium uncertainty combines (i) idiosyncratic, diversifiable component $\varepsilon_{i\omega}$ and (ii) aggregate scenario risk. We decompose these in panel D. The idiosyncratic component—diversifiable within a scenario—is large: removing scenario and household fixed effects still leaves a standard deviation of $\$1{,}554$. This is the uncertainty we can insure through a premium subsidy. The uncertainty across climate scenarios within a household is even more acute but not insurable in the cross-section through premium subsidies.

To further illustrate the degree of uncertainty in the idiosyncratic, diversifiable-across-households, risk, Figure 3 plots the percentage premium change by 2070 for a random policy sample under two of the least-correlated models (MPI on the horizontal axis; IPSL on the vertical). The dashed line represents pure aggregate risk (uniform scaling of premiums). Deviations from this line reflect diversifiable idiosyncratic reclassification risk. The degree to which some households are above the diagonal and others below represents the diversifiability of premium risk; in one scenario house A is high risk and house B is low, in the other scenario, the inverse. Visually, there is substantial diversifiable risk. We next turn to quantifying the benefits of insuring that diversifiable risk.

Figure 3 about here

6.3 Reclassification Risk Results

We measure the reduction in insurance against reclassification risk induced by RR2.0. Our estimates imply that RR2.0 removed a level subsidy of approximately $s = 0.5$.[10] Let $\widetilde{P}_i(s)$ denote the *future* premium distribution for household i under policy s, and let $\mathsf{CE}_i(\cdot)$ denote the certainty equivalent. RR2.0 corresponds to $\mathsf{CE}_i\big(\widetilde{P}_i(0)\big)$; undoing RR2.0 corresponds to $\mathsf{CE}_i\big(\widetilde{P}_i(0.5)\big)$. The raw value of undoing RR2.0 for household i is

$$\mathsf{CE}_i\big(\widetilde{P}_i(0)\big) - \mathsf{CE}_i\big(\widetilde{P}_i(0.5)\big)\,.$$

To isolate the pure *insurance* value (net of mechanical transfers), we define the change in the mean premium

$$c_i \equiv \mathbb{E}_\omega\big[\widetilde{P}_i(0) - \widetilde{P}_i(0.5)\big]\,,$$

and compute

$$\underbrace{\mathsf{CE}_i\big(\widetilde{P}_i(0) - c_i\big) - \mathsf{CE}_i\big(\widetilde{P}_i(0.5)\big)}_{\text{pure insurance value of undoing RR2.0}}\,.$$

We assume CARA utility, $u(c) = -\exp(-\gamma c)$, and report results for multiple γ.[11]

Table 5 about here

Table 5 indicates that RR2.0 exposed households to substantial reclassification risk. Under our central risk-aversion parameter $\gamma = 5\times10^{-4}$, the level shift reduces insurance value by \$64.5, about 7% of the subsidized pre-RR2.0 premium, 3.6% of the RR2.0 premium and 1% of the projected future premium.

7 Moral Hazard: Mitigation

Flood risk mitigation encompasses a range of strategies aimed at reducing potential damage. However, the most effective and commonly incentivized measure is elevating a structure above expected flood levels. This involves raising the lowest floor, typically via pilings, piers, or extended foundations. While elevation can substantially lower expected damages (e.g., by 30–70% depending on height and flood depth (U.S. Army Corps of Engineers (2021))), it is often prohibitively expensive (estimates range from \$30,000 to \$90,000 for a typical house (Congressional Budget Office (2017))). This means that when flood insurance is subsidized, households might not choose to voluntarily elevate their house. On the other hand, elevation might be incentivized if premiums are adjusted to fully reflect risk, and if premium discounts are offered commensurate with the reduction in damage.

[10] Per Table 4, the mean pre-RR2.0 premium is \$913, the mean post-RR2.0 premium is \$1,772.

[11] We use $\gamma = 5 \times 10^{-5}$ as a lower bound as estimated by Cohen and Einav (2007), and $\gamma = 5 \times 10^{-4}$ as our primary estimate. Note, this still errs conservatively: Sydnor (2010) and Barseghyan et al. (2011) estimate $\gamma = 5 \times 10^{-3}$, an order of magnitude larger.

In this section, we quantify the responsiveness of elevation decisions to insurance price changes. To tackle this difficult empirical challenge, we turn to data from the state of Florida, which requires data on mitigation efforts. While not representative of the US, Florida is more representative of the most flood-prone areas in the country: the average full-risk premium is $\$2,996$ relative to the national average of $\$1,941$.

7.1 Data

Since 2017, all elevation certificates (ECs) in Florida have been required to be submitted to the Florida Division of Emergency Management, providing a comprehensive dataset for analyzing mitigation trends. These certificates adhere to a standardized FEMA format, capturing critical details such as the structure's location (in particular, its flood zone designation), base flood elevation, actual lowest floor elevation, foundation type (e.g., slab, crawlspace, or elevated), presence of enclosures or vents below the elevated floor and so on. We collect all ECs filed Florida Division of Emergency Management (2025) and extract the relevant textual data. In addition, we use our RR2.0 pricing data to calculate, by flood zone, the RR2.0-induced price changes.

7.2 Policy Change

Risk Rating 2.0 affected incentives to elevate new buildings distinctly in high and low flood risk areas. Special Flood Hazard Areas (SFHAs) are high-risk flood zones defined by FEMA as areas with a 1% annual chance of flooding. Flood insurance is mandatory for properties with federally-backed mortgages inside the SFHA. New buildings in SFHAs are subject to elevation mandates to meet or exceed the Base Flood Elevation (BFE) under NFIP minimum standards. This requirement was unchanged by RR2.0. As such, we do not expect substantial responsiveness of elevation to flood insurance price changes inside the SFHA since elevation is already mandated, although imperfectly enforced,[12] by building codes.

In contrast, outside of the SFHA there is no such elevation mandate. Areas outside of SFHA include moderate risk (1-in-100 to 1-in-500 year flood risk), and low risk (less than 1-in-500 year flood risk) properties. Because there is no mandate, elevation outside of the SFHAs is entirely voluntary and therefore likely more responsive to insurance price incentives. For these reasons, identifying the responsiveness of elevation to insurance price is best done by comparing properties both inside or both outside the SFHA; comparisons across are less informative.

7.3 Analysis

We test the responsiveness of elevation decisions to insurance price changes, separately inside and outside of the SFHA.

[12] An audit by Mathis and Nicholson (2006) finds 89% compliance with the SFHA elevation requirements.

Inside the SFHA, we compare very high risk properties to the high risk. All properties in SFHAs are, by definition, exposed to a greater than 1% annual chance of flood. The very high risk properties, located near the coast, are additionally exposed to high velocity wave flooding; the high risk are not. The average full-risk premium for the very high risk is $4,478; for the high risk $2,288. The price changes under RR2.0 were correspondingly sharper: the mean premium differences pre- and post-RR2.0 were $2,142 (very high risk) versus $1,202 (high risk). If the very high risk have a greater elevation response to RR2.0 than the high risk, this would indicate that elevation is responsive to price incentives even conditional on the stringent elevation requirements already in place in SFHAs.

Outside the SFHA, we compare moderate risk to low risk properties. Moderate risk properties saw their premiums increase under RR2.0 by $1,075, reaching an average full-risk premium of $1,724. Low risk properties' premiums increased by a smaller $702 to an average full-risk premium of $1,312. We will infer the responsiveness of elevation decisions to insurance prices, when there is no elevation mandate in the building codes, by comparing the changes in moderate-risk versus low-risk properties.

We estimate the probability of a house i in NFIP community j and year t elevating by 3 feet or more relative to the base flood elevation (inside the SFHA) or local ground height (outside the SFHA). We run the analyses separately inside and outside the SFHA. Our treatment variables are being in the higher-risk flood zone (in the SFHA: very high risk; outside of the SFHA: moderate risk). Specifically, we estimate:

$$\text{Elevate}_{it} = \beta_0 + \sum_{t \neq 2022} \beta_{1,t} \text{ Higher Risk Category}_i + \gamma_t + \delta_j + \epsilon_{it} \tag{3}$$

We include fixed effects for year (γ_t) and for NFIP community (δ_j) to control for community-specific mitigation incentives. The coefficients of interest are $\beta_{1,t}$, the effects at different time periods of being in the very high (moderate) versus high (low) risk areas inside (outside) the SFHAs. The coefficient estimates are displayed in Figure 4.

Figure 4 about here

We find that elevation responds to the removal of RR2.0 subsidies outside of SFHAs, but not inside. The top panel of Figure 4 shows that, despite facing a price change almost twice as high, very high risk properties did not change their elevation decisions relative to high risk properties. Unfortunately, there are fairly strong pre-trends for this area, but if anything they further confirm the fact that mitigation inside SFHAs didn't respond positively to the policy change. As discussed, this is likely due to existing and continuing SFHA-specific elevation requirements and building codes.

In contrast, outside of the SFHA, where elevation is voluntary, there is a 10% increase in elevating by more than 3 feet in the moderate risk areas relative to the low-risk areas. The difference in mean price increases was approximately $372, against a mean full-risk price of $1,518, indicating that an approximate price change of 25% yielded a 10% elevation increase. This implies an

elasticity of approximately 0.4, calculated as:

$$\text{Elasticity} = \frac{\%\,\Delta \text{ in Elevation}}{\%\,\Delta \text{ in Price}} = \frac{10\%}{25\%} = 0.4$$

This elasticity, in combination with the returns to 3 feet of mitigation – a 40% reduction in expected risk – are the key inputs we use to calibrate the fiscal externality from under-mitigation due to subsidies in Section 9.

8 Moral Hazard: Location

Subsidized flood insurance has long been criticized for dulling price signals and incentivizing excessive construction in high-risk areas (Ben-Shahar and Logue (2016)). This is in addition to evidence that when flood insurance becomes available for the first time in an area, this spurs development (Peralta and Scott (2024)). By removing subsidies and exposing households to full-risk prices, RR2.0 aims to reduce overdevelopment. In this section, we study population movements from locations that experienced dramatic price increases under RR2.0 to those that experienced lesser increases or even decreases.

8.1 Data

To measure population changes, we use quarterly data on the number of occupied residential addresses at the census-tract level, provided by the U.S. Postal Service (USPS). This serves as our primary outcome variable. We link this with NFIP data on pre- and post-RR2.0 premiums, also aggregated to the tract level.

We compare census tracts whose median property experienced a price increase under the NFIP to those that experienced a price decrease. Specifically, we define the indicator variable $\text{MedianIncrease}_{ct}$ to be equal to 1 if the median premium in tract c at time t increased under RR2.0 and 0 otherwise. These treated and control tracts have very similar prices before RR2.0, reflecting the subsidies of the previous system; the median premium in treated census tracts was \$929 while in non-treated tracts it was \$769. But these tracts received very different price shocks from RR2.0, with premiums rising by \$1063, to \$1992, in treated tracts, but only by \$152, to \$921, in non-treated tracts.

8.2 Empirical Strategy

We compare occupied residential addresses in census tract ct measured in quarter q, $\text{OccupiedResidential}_{ct,q}$, across tracts that on average experienced a price increase versus a decrease under RR2.0. Specifically, we estimate

$$\log(\text{OccupiedResidential}_{ct,q}) = \alpha_{ct} + \gamma_q + \sum_{q' \neq \text{Q1 2022}} \beta_{q'} \Big(\text{MedianIncrease}_{ct} \times \mathbf{1}\{q = q'\} \Big) + \varepsilon_{c,t,q}, \quad (4)$$

where α_{ct} are census–tract-by-time fixed effects, γ_q are quarter fixed effects. The coefficients of interest are $\beta_{q'}$, which capture the relative difference in occupied residences in the more- versus less-exposed census tracts, normalized to zero in Q1 2022. Figure 5 plots the estimated $\beta_{q'}$.

Figure 5 about here

Figure 5 shows a net movement of about 0.75% of the population from census tracts treated with a large price increase under RR2.0 to those with smaller (or negative) price changes. This movement begins in 2023 (the full rollout of RR2.0 concluded at the end of 2022) and rises gradually, appearing to stabilize by mid-2025.

Translated into dollars, RR2.0 induces 0.75% of the population to move from locations with an average risk of $\$1{,}992$ to locations with an average risk of $\$921$. We use this moral hazard response as an input to our optimal policy analysis in Section 9.

These effects are precisely estimated but relatively small. For reference, estimates of outmigration impacts of actual disasters range from about 1.5% for severe disasters (Boustan et al., 2020), to roughly 7% following Hurricane Maria (Aron-Dine, 2025), to 12% after Hurricane Katrina (Deryugina et al., 2018).

9 Optimal Policy

In this section, we calibrate a discrete approximation of the formula in Proposition 2. We use this calibration to evaluate the efficiency of a marginal subsidy reduction starting either at the pre-RR2.0 prices or at the full-risk post-RR2.0 prices.

Throughout, we assume that household choices are discrete. Specifically, we assume that insurance fully covers all losses and households either purchase or go uninsured; the mitigation margin we consider is whether to elevate a new house or not; the location choice is binarized to the household choosing between a high-risk or low-risk location. We keep the choice of the semi-elasticity of demand $\mathcal{E}^q$ and the average flood insurance premium (i.e. expected risk) amongst the marginals $\bar{P}$ general. We derive optimal subsidy rates for many pairs of values of ($\mathcal{E}$, $\bar{P}$).

9.1 FEMA Savings Calibration

The FEMA savings term is equal to $\sum_i f P_i \mathcal{E}_i^q$. We need to estimate the spillover from a lack of insurance to increased FEMA spending f, and the elasticity of insurance demand with respect to price $\mathcal{E}_i^q$. Since, additionally, we assume a homogeneous elasticity and FEMA spillover across people, we need to estimate $\Delta\text{FEMA}(s) = f\mathcal{E}^q\bar{P}$.

To calibrate the FEMA spillover, we use our estimate from Section III that measures the change in FEMA spending with respect to coverage increases *conditional on a flood*. The term in the equation is unconditional. Hence, we need to estimate the probability of a flood. We take the

expected damage per house *given a flood* to be $L_f = \$31{,}921$ (see Table 1). Under actuarial fairness, an average premium $\bar{P}$ implies a flood probability $p_f = \bar{P}/L_f$ each year.

Therefore, the change in *unconditional* expected FEMA spending per house for a 1 percentage point (0.01) increase in coverage is $\beta_{\text{uncond}} = \beta_{\text{cond}} \cdot p_f = (-\$203) \cdot \frac{P}{L_f}$. For a small ad valorem subsidy s that induces a (log) price change $\Delta \ln \tilde{P} \simeq -s$, coverage changes by $\Delta q \simeq \mathcal{E}^q \Delta \ln \tilde{P}$. Hence, the FEMA saving per house-year from a small subsidy s is

$$\Delta\text{FEMA}(s) = (-\beta_{\text{cond}}) \frac{\bar{P}}{L_f} \cdot 100 \cdot \mathcal{E}^q \cdot \Delta \ln \tilde{P}.$$

For example, plugging in $\Delta \ln \tilde{P} = 0.01$, $L_f = 31{,}921$, $P = \$1{,}500$, and $|\mathcal{E}^q| = \mathbf{0.32}$, yields $\Delta\text{FEMA}(0.01) \approx$ **$3.05** per house-year in response to a 1% price change.

9.2 Insurance Against Reclassification Risk Calibration

We directly calculate the WTP for a 1% reduction in log prices similarly to Section 6. That is, starting at a given s by $WTP(s) = CE(\tilde{P}(s + \Delta s)) - CE(\tilde{P}(s))$ where Δs is the subsidy change induced by a 1% log price change.

A 1% price reduction starting from $s = 0$ provides insurance against reclassification risk worth $\$1.04$ per insured household. At $s = 0.5$ this falls to $\$0.25$, and at $s = 1$ to zero. The difference is due to the concavity of utility. A dollar of insurance against reclassification risk is more valuable when the household is otherwise uninsured ($s = 0$) than when the household is already substantially insured ($s = 0.5$), and is worth nothing when the household is fully insured $s = 1$. As a back of the envelope check, for normal risks and CARA utility, the WTP for insurance is quadratic in s, consistent with the $\approx 4\times$ drop between $s = 0$ and $s = 0.5$.

9.3 Coverage Fiscal Externality Calibration

The coverage term is equal to $\sum_i sP_i\mathcal{E}_i^q$. On average, since we assume $\mathcal{E}_i^q$ does not vary across people, and scaled by the price change $\Delta \ln \tilde{P}$, this becomes:

$$\Delta\text{Coverage FE} = s \times \mathcal{E}^q \times \bar{P} \times \Delta \ln \tilde{P}.$$

At our central estimates of $\Delta \ln \tilde{P} = 0.01$, $\bar{P} = \$1{,}500$, and $\mathcal{E}^q = 0.32$, this coverage externality is zero at $s = 0$, about $2.40 per house-year at $s = 1/2$, and about $4.80 at $s = 1$.

This term represents the fiscal externality in the form of higher subsidy outlays as households change their insurance choices. Specifically, when households change their insurance choices they do so trading off private costs – the *post-subsidy* premium – against private benefits. They do not consider the public fiscal cost of selecting insurance, i.e. the subsidy itself.

For this reason, the marginal fiscal externality scales with s. This logic mirrors the Harberger's triangle for taxes, where the deadweight loss from a tax increases quadratically with the tax rate.

The first dollar of tax has no deadweight loss (as households internalize the entire price), but this grows rapidly due to greater distortions in behavior. Similarly, here, the fiscal externality *for the marginal household* from subsidies scales with s, such that the total fiscal externality is quadratic in s. This is the main reason we find an optimal subsidy level that is strictly between 0 and 1: the FEMA benefits of a subsidy scale with s while the coverage costs of a subsidy scale with s^2.

9.4 Location Fiscal Externality Calibration

For the location term, we focus on the intensive-margin response among the insured. We therefore set $q = 1$ and $f = 0$, since uninsured households neither pay nor respond to insurance premiums. Under this restriction, the location component on the right-hand side reduces to $s \cdot \mathbb{E}[\partial P_i/\partial x] \cdot \mathcal{E}^x$.

When an insured person moves, the internalize the new post-subsidy insurance price but not the subsidy itself. The subsidy they do not consider becomes the fiscal externality from their location choice. Identically to the coverage term, the fiscal externality from the marginal mover is 0 at $s = 0$ and then increases in s.

We calibrate $\mathbb{E}[\partial P_i/\partial x]$ from the observed change in expected annual loss among movers in our RR2.0 analysis: movers go from tracts with average expected loss \$1,992 to \$921, so $\Delta P \approx -\$1{,}071$. The location semi-elasticity $\mathcal{E}^x$ comes directly from the event-study: high-premium tracts experience a net out-migration of -0.75% in response to a 98% premium increase (RR2.0 increased high-risk location premiums from $\$1,063$ to $\$1,992$). Hence, $\mathcal{E}^x \approx (-0.0075)/(+0.874) \approx -0.0086$. This gives a location fiscal externality of

$$\Delta\text{Location FE} = s \times 0.0086 \times \Delta P \times \Delta \ln \tilde{P}.$$

For a 1% log price cut, the location fiscal externality is zero at $s = 0$, $\$0.046$ at $s = 1/2$, and $s = \$0.09$ at $s = 1$.

9.5 Mitigation Fiscal Externality Calibration

For the mitigation term we again focus on the intensive-margin response among the insured and set $q = 1$ and $f = 0$, since uninsured households neither pay nor respond to insurance premiums. Under this restriction, the right-hand-side contribution reduces to $s \cdot \mathbb{E}[\partial P_i/\partial m] \cdot \mathcal{E}^m$, so a small policy change that induces a price change $\Delta \ln \tilde{P}$ contributes $s \cdot \mathbb{E}[\partial P_i/\partial m] \cdot \mathcal{E}^m \cdot \Delta \ln \tilde{P}$ per house-year. Identically to the location and coverage fiscal externalities, the marginal mitigation fiscal externality is zero at $s = 0$ and then scales with s.

We define mitigation as elevating the structure by 3 feet. Our engineering (U.S. Army Corps of Engineers (2021)) estimate is that elevating the structure by 3 feet causes a reduction in expected loss (i.e. the full-risk premium) of 40%. We use our event-study (outside of the SFHA) to calibrate the mitigation semi-elasticity. We take the difference in price increases from the moderate versus the low risk relative to the initial moderate risk price: $(\$1075 - \$702)/\$649 = 57\%$. That is, relative to the price change in the low risk areas, the moderate risk areas had a 57% increase in prices,

and (per the event study) this caused an increase in mitigation of 10% (relative than the low risk areas). Hence, $\mathcal{E}^m \approx 0.10/0.57 = 0.18$

In summary, 1% subsidy yields a mitigation fiscal externality of $s \times 0.4 \times 0.18 \times \bar{P} \times 0.01. = 0.00072 \times s \times \bar{P}$ At $\bar{P} = 1500$, this is equal to $\$0.54$ per affected house at $s = 0.5$. But the mitigation response to insurance price changes was only observed outside of the SFHA (which makes up 48% of total insured properties). Hence, in terms of per total houses per year, the mitigation fiscal externality is $0.48 \times 0.00072 \times s \times \bar{P}$, or only about $\$0.25$ per house-year at $s = 1/2$.

9.6 Optimal Subsidy

The optimal subsidy is the s^* that solves the equation in Proposition 2. That condition is of the form

$$\underbrace{WTP(s)}_{\text{Reclass risk: quadratic in } s} + \underbrace{\Delta FEMA(s)}_{\text{constant in } s} = \underbrace{\Delta\text{Coverage FE}}_{\text{linear in } s} + \underbrace{\Delta\text{Location FE}}_{\text{linear in } s} + \underbrace{\Delta\text{Mitigation FE}}_{\text{linear in } s} \quad (5)$$

The right-hand side (costs) are zero at $s = 0$ and increase linearly thereafter. The left-hand side (benefits) are highest at $s = 0$ and then either remain constant (FEMA) or decline quadratically (WTP for insurance against reclassification risk). Although not guaranteed theoretically, for all reasonable parameter choices the optimal s is well-defined and interior to $[0, 1]$.

We find the optimal s^* for different combinations of $\bar{P}$ and $\mathcal{E}^q$. The results are illustrated in 6.

Figure 6 about here

At our central estimates – $\mathcal{E}^q \approx$ **0.32** and $\bar{P} \approx \$1,500$ – the optimal subsidy is approximately **57%**.

Even at demand elasticities substantially higher or lower, and at average premiums among the insured starkly different from our central estimate, the optimal subsidy remains in the range 55% - 60% range. The relative insensitivity of s^* is because the primary benefit (FEMA savings) and the primary cost (ΔCoverage FE) are both linear in $\bar{P}$ and in $\bar{\mathcal{E}^q}$. This is because both terms are similar functions of the number of marginals (hence $\mathcal{E}^q$) and the their riskiness (hence $\bar{P}$).

An implication of this is that our use of demand elasticities and premium data from the set of people treated by RR2.0 is relatively innocuous. Even if the demand elasticity or riskiness of the insured pool – i.e. their annual probability of flood or expected damage – were to change, this analysis would produce a similar policy prescription.

9.6.1 Redistribution.

Our formula for the optimal subsidy considers only efficiency. We did not, in particular, worry about the difference in wealth or income, and therefore in marginal utility, between those who are in higher risk flood areas and those in lower risk areas. Because we are considering proportional subsidies, there is an implicit redistribution from the latter to the former.

To understand in which direction redistributive concerns would push, we proceed as follows. Since we do not have household income data at the NFIP policy level, we assign to each policyholder the median income of their census tract, which is the most granular geographic identifier in the data. We simulate the net gains and losses from a 1% subsidy funded by a lump sum tax on all the insured. The results are in Figure 7.

Figure 7 about here

Figure 7 bins everyone by their income, and then within each bin calculates the average net gain or loss from the simulated subsidy. The height of the bars represents the average gain/loss; the color represents the number of people within each bin.

We see that the redistributive implications of the subsidy we consider are minor. The bulk of the affected individuals have incomes from $\$50,000$ to $\$100,000$. The subsidy effectively transfers a dollar from those earning $\$50,000 - \$70,000$ to those earning $\$80,000 - \$100,000$. Additionally, there is a transfer among the relatively few policy holders at very high or very low income: those earning more than $\$170,000$ or less than $\$50,000$ subsidize those earning $\$150,000$ to $\$170,000$.

Overall, there is a relatively weak positive association between income and gains under the subsidy. This could be reversed by funding any subsidy out of a progressive tax (e.g. income tax), but we do not pursue this issue further here.

9.6.2 SFHA Regulations

An SFHA designation has two main effects: it mandates mitigation, reducing that margin to zero, and it mandates insurance take-up (although not perfectly enforced) so the demand semi-elasticity is 0.25 instead of 0.32 (see Table 3). The analogue of Figure 6 with the SFHA restrictions is displayed in Figure 8.

Figure 8 about here

Because the SFHA regulations mandate mitigation, a premium subsidy has a lower cost. Hence, at any combination of $\mathcal{E}^q$ and $\bar{P}$, the optimal subsidy is higher than in the baseline. In particular, at the SFHA-estimated demand elasticity of 0.25, and the baseline premium estimate of $\bar{P} = \$1,500$, the optimal subsidy is 62.8%, an increase relative to the baseline of about 5%.

9.6.3 Heterogeneity in Demand and FEMA Savings

As a robustness check, we relax homogeneity along two margins. First, we allow the FEMA spillover to vary linearly with the subsidy. We estimate, shown in Appendix A.5, that instead of a flat $\$-203$ effect, an increasing FEMA saving with respect to s: starting at $\$-116$ at $s = 0$ rising (in magnitude) with a slope of $\$-112$. This attenuates the average FEMA-savings benefit over the relevant range—despite a mild upward slope in s—so the marginal-benefit curve shifts down and the optimal subsidy falls.

Second, we allow the demand semi-elasticity to vary linearly with the subsidy. Per Table 3, the elasticity at $s = 0$ is ≈ 0.61 falling linearly to -0.27 as the subsidy increases to 100%. This makes the costs more convex in s, offset only by a linear increase on the benefit side. The former dominates and the optimal subsidy increases to be above our baseline. We conclude that our baseline estimate of approximately 62% is, if anything, conservative.

9.6.4 'Worst-Case' Sensitivity

Given the limitations of our data, one might be concerned that we have overstated the benefits and understated the costs. For example, the location and mitigation elasticities we have measured are relatively short run (only three or four years since RR2.0) and might increase over time. We check the sensitivity of our policy prescriptions to assumptions on costs and benefits.

We assume that the migration response is 1.5% instead of 0.86%. The former is the estimate from Boustan et al. (2020) of population responses to natural disasters, surely an upper bound for responses to a subsidy change. We assume the mitigation response is twice as large as in our baseline (instead of about 25 cents per house year at $s = 1/2$, we assume 50 cents). Correspondingly, we take conservative estimates for the benefits: we discount the reclassification risk insurance benefits by 2% a year for 35 years, and we assume the FEMA spillover is only $\$116$ per percent of coverage, the minimum we found in our heterogeneity analysis. We recompute the optimal subsidy with these assumptions. The results are in Figure 9.

Figure 9 about here

Under these very high estimates for the costs of a subsidy, and very low estimates for the benefits, the optimal subsidy is between 30 and 35%. Although half as big as our baseline estimates, this is still far from zero. Our qualitative conclusion is that a non-trivial premium subsidy is optimal under a wide variety of assumptions and estimates.

10 Conclusion

The ongoing increase in "global weirding"[13] has raised attention to the risks posed to households from natural disasters. One of the most salient and most discussed is flood damage. In 2024 alone, over 100,000 houses suffered flood damage for a total cost of more than $10 billion (Reinsurance News (2024)).

In theory, such damage is offset by mandated flood insurance offered by the National Flood Insurance Program. But in practice, NFIP coverage is very incomplete, with only 5% of households nationwide holding policies. Coverage is incomplete even in designated high risk Special Flood Hazard Areas averaging only 23% (Amornsiripanitch et al. (2024)).

[13] See Friedman (2010)

One response to this lack of coverage was extensive subsidies offered for the purchase of the NFIP – both in terms of a reduced average premium, and mitigated variation by risk level of the area. But this has long been criticized as promoting development – and retarding mitigation – in our most flood-prone areas. As a result, in 2021 the government introduced a new actuarially fair pricing system for flood insurance, Risk Rating 2.0, which properly reflected the cost of flood risk around the country.

But what has largely been ignored in the praise for this new policy is the fact that flood insurance operates in a second-best world. In particular, the existence of an incomplete mandate for insurance, combined with ex post government spending to bail out flooded homes, means that removing subsidies leads to a rise in ex post disaster spending. We estimate that this is a very large offsetting source of increased government spending, with each 1% fall in flood insurance coverage leading to $203 per household in additional (flood-risk weighted) recovery spending.

Moreover, the large idiosyncratic risk faced by households due to uninsurable locational variation in the impacts of future climate changes means that there is a welfare loss arising from actuarially fair pricing. We estimate that this reclassification risk is also sizeable, amounting to $64 per insured household for standard risk preference assumptions.

To assess the importance of these offsetting costs to moving to RR2.0, we also estimate the benefits in terms of both population movement and mitigation. In fact, we find that these benefits are much smaller than the costs, an effect an order of magnitude lower than the direct benefits from the FEMA spillover and subsidy expenditures from changed coverage.

We develop a model of optimal flood insurance which allows us to put these estimates together to calibrate optimal flood subsidies. We find that the optimal subsidy is 57%, and that it is quite robust to parameter variation. Importantly, any income redistribution done through these subsidies is minor, and does not offset our conclusion.

As with any exercise of this nature, more work is needed to assess the robustness of our findings. In particular, our estimates of the benefits of RR2.0 are necessarily short run. Over a much longer time period, population movements and mitigation may respond more to price signals and lower the optimal subsidy rate. But it seems very unlikely that the optimal subsidy rate will approach zero.

Moreover, we critically take FEMA policy as given for this exercise. In fact, if FEMA cuts back on its ex post assistance to flood victims, that would substantially affect our conclusions. Indeed, President Trump recently announced his intention to “wind down” FEMA. [14] While that would dramatically impact our conclusions, it remains to be seen whether such a radical reform can survive the political process.

In any case, the main point from our analysis holds: in the second-best world of flood insurance, careful thought is needed about the various forces that determine optimal insurance pricing policy.

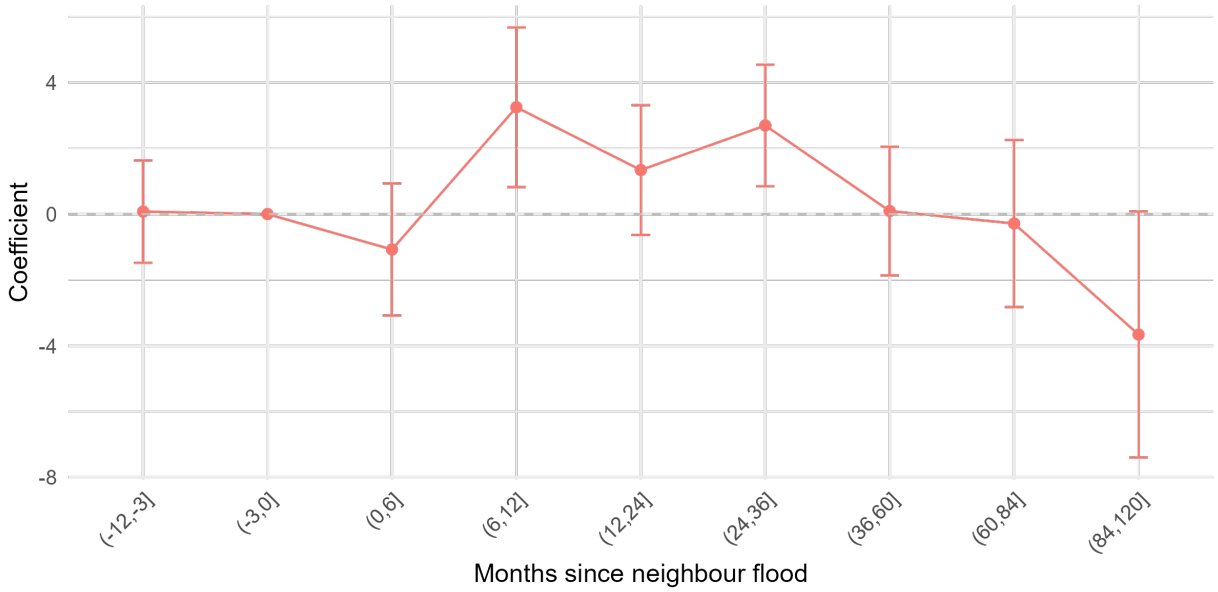
[14] See: Reuters (2025)

Table 1. Damage-Weighted Summary Statistics (2010–present)

Variable	Weighted Mean	Weighted Median	90th Percentile
Panel A: Aggregate (levels)			
Property Damage (*millions USD*)	4,700	3,200	10,000
Total Population	980,771	345,995	4,525,519
Total Housing Units	380,678	142,115	1,714,340
Panel B: Per house			
NFIP Policies per House	0.18	0.16	0.49
NFIP Payouts per House ($)	3,440.41	1,798.16	8,255.15
Total Non-NFIP Spending per House ($)	955.30	700.14	2,566.32
FEMA Assistance Claims per House	0.08	0.06	0.20
Property Damage per House ($)	31,920.97	27,454.25	81,466.06
Disaster Assistance per House ($)	733.16	431.34	1,822.86
Fannie Mae Losses per House ($)	3.11	1.17	6.14
Freddie Mac Losses per House ($)	1.71	0.29	2.53
HGMP Expenditures per House ($)	147.42	88.94	282.36
Approved SBA Loans per House ($)	703.57	301.30	1,404.00

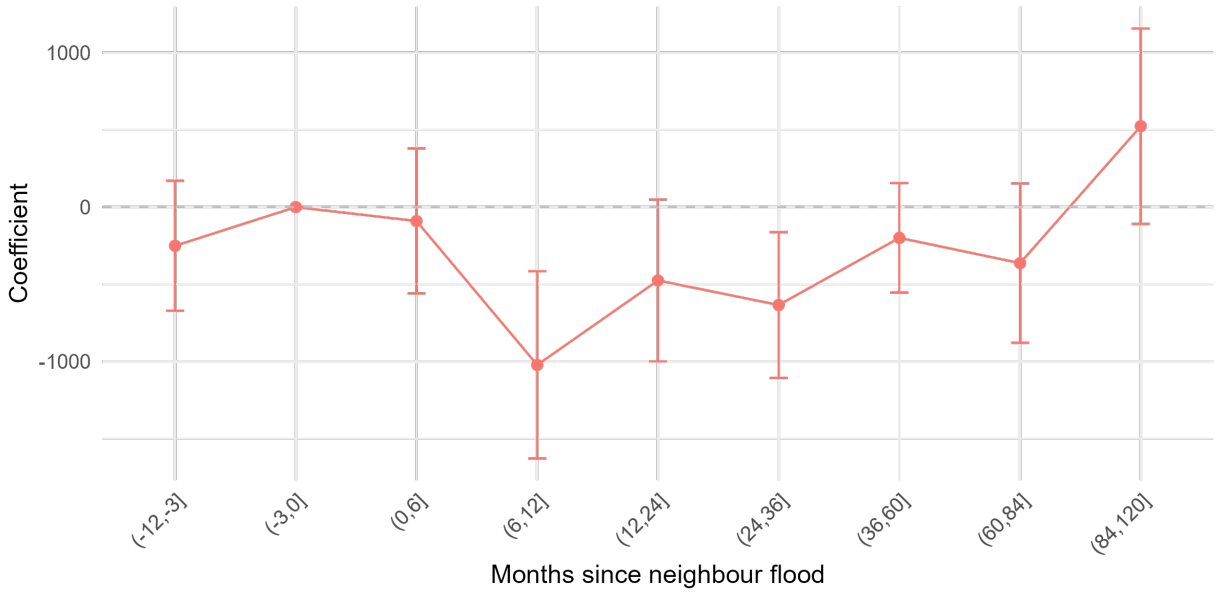
Notes: Sample statistics are from counties with presidentially declared disasters from 2010 to 2024 where assistance under the Individuals and Households Program (IHP) was authorized. All statistics are *damage-weighted* using county-level flood property damage from the NOAA Storm Events database. Panel A reports aggregate (level) outcomes for each disaster-county; Panel B reports per house in the county averages. Dollar amounts are inflated to 2023 dollars.

Figure 1. Response of NFIP coverage to neighboring-county floods



Notes: Event-study estimates from a regression of NFIP coverage per house (measured two months before flood onset) on indicators for years since the most recent flood in the closest neighboring county, controlling for years since own-county flood and flexible quadratic functions of flood damage within bins (cutoffs at $1, $10, $100, $1,000, $10,000, $100,000, and $1,000,000 per house). Sample includes U.S. counties with presidentially declared flood disasters from 2010 to 2024 where Individuals and Households Program (IHP) assistance was authorized; observations weighted by county-level flood property damage (NOAA Storm Events Database) and inflated to 2023 dollars. Coefficients normalized to zero in the year before the neighboring flood; 95% confidence intervals based on standard errors clustered at the county level.

Figure 2. Response of ex-post spending to neighboring-county floods



Notes: Event-study estimates from a regression of total ex-post disaster spending per house on indicators for years since the most recent flood in the closest neighboring county, controlling for years since own-county flood and flexible quadratic functions of flood damage within bins (cutoffs at $1, $10, $100, $1,000, $10,000, $100,000, and $1,000,000 per house). Sample includes U.S. counties with presidentially declared flood disasters from 2010 to 2024 where Individuals and Households Program (IHP) assistance was authorized; observations weighted by county-level flood property damage (NOAA Storm Events Database) and inflated to 2023 dollars. Coefficients normalized to zero in the year before the neighboring flood; 95% confidence intervals based on standard errors clustered at the county level.

Table 2. Main Results: Total Fiscal Cost

	OLS	**IV (Neigh. County Flood $\leq$ 5 Years)**
NFIP Policies per 100 Houses	-171.510^{***}	-203.188^{**}
	(34.672)	(73.416)
Num. Obs.	815	815
Within R^2	0.915	0.912

Notes: Dependent variable is total ex-post disaster spending per house (sum of FEMA IHP assistance, net fiscal cost of SBA loans, HMGP expenditures, and GSE foreclosure losses; all inflated to 2023 dollars). Each column reports coefficients from a two-way fixed effects regression of ex-post spending on NFIP coverage per 100 houses (measured two months pre-flood), with quadratic controls within bins of flood damage per house. Sample includes U.S. counties with presidentially declared flood disasters from 2010 to 2024 where Individuals and Households Program (IHP) assistance was authorized; observations weighted by county-level flood property damage. The IV specification instruments NFIP coverage using an indicator for a flood in the closest neighboring county within the past 5 years. Standard errors clustered at the county level in parentheses. Standard errors in parentheses. Significance: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$.

Table 3. Renewal vs. Price Relative to Full-Risk: Semi-Elasticity and Heterogeneity

	(1) Semi-Elasticity Only	(2) + Relative Price Interaction	(3) + SFHA Interaction
Log(Price/Full-Risk)	-0.322^{***}	-0.270^{***}	-0.396^{***}
	(0.002)	(0.002)	(0.003)
Log(Price/Full-Risk) $\times$ Price/Full-Risk		-0.338^{***}	
		(0.007)	
Log(Price/Full-Risk) $\times$ SFHA			0.146^{***}
			(0.004)
Policy FE	Yes	Yes	Yes
N	4,905,305	4,905,305	4,905,305
Adj. R^2	0.53	0.53	

Notes: Each column reports estimates from a linear probability model of NFIP policy renewal (an indicator for renewal into year $t+1$) on the log of the relative renewal price $\tilde{P}_{i,t} \equiv P^{\text{offer}}_{i,t+1}/R_{i,t+1}$, where $P^{\text{offer}}_{i,t+1}$ is the offered renewal premium (capped at 18% annual increase or full-risk level) and $R_{i,t+1}$ is the full risk-based premium under Risk Rating 2.0 (RR2.0). Columns (2) and (3) add interactions for heterogeneity by subsidy level (Price/Full-Risk) and Special Flood Hazard Area (SFHA) indicator, respectively. Standard errors clustered at the policy level in parentheses. Significance: $p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$.

Table 4. Summary Statistics of Premiums and Hazard Metrics

Metric	Mean	SD	P10	P90
Panel A: Pre RR2.0 Premiums				
Pre RR2.0 Premium	913.00	1604.0	341.00	1763.00
Panel B: RR2.0 Premiums & Current Hazard Levels (IDW)				
RR2.0 Premium	1772.00	4107.0	386.00	3772.00
100-year Rainfall Depth (mm)	171.00	54.0	77.00	235.00
100-year Storm Tide Water Level (mm)	2523.00	865.0	1287.00	3647.00
Joint Probability (%) (x 100)	0.32	0.1	0.17	0.46
Panel C: Future Premiums and Hazard Levels				
Future Premium (Scenario Avg)	6382.00	3317.0	3668.00	8665.00
100-year Rainfall Depth (mm)	802.00	292.0	309.00	1204.00
100-year Storm Tide Water Level (mm)	3668.00	1335.0	1824.00	5601.00
Joint Probability (%) (x 100)	5.59	2.0	2.60	8.46
Panel D: Decomposition Across Climate Scenarios				
Household Fixed Effect	0.00	3317.0	-2715.00	2283.00
Scenario Fixed Effect (vs current mean)	4611.00	3489.0	1974.00	9634.00
Idiosyncratic Risk	0.00	1554.0	-1342.00	1666.00

Notes: Notes: Summary statistics at the property level for NFIP policies priced under Risk Rating 2.0 (RR2.0). Panels A–B report historical premiums (pre- and post-RR2.0 implementation) and current (IDW-interpolated) hazard metrics from Gori et al. (2021) synthetic flood data under present climate. Panel C projects 2070 premiums (scenario average) and hazards under SSP5-8.5 across eight CMIP6 general circulation models, holding property characteristics fixed and mapping hazards to premiums via OLS regression on rainfall depth (mm), storm tide level (mm), joint exceedance probability (%), granular building controls. Panel D decomposes cross-scenario variability in future premiums into household fixed effects (γ_i), scenario fixed effects (vs. current mean), and idiosyncratic residuals ($\varepsilon_{i\omega}$). All values in constant 2023 dollars.

Figure 3. Variability in Future Flood Risk Scenarios



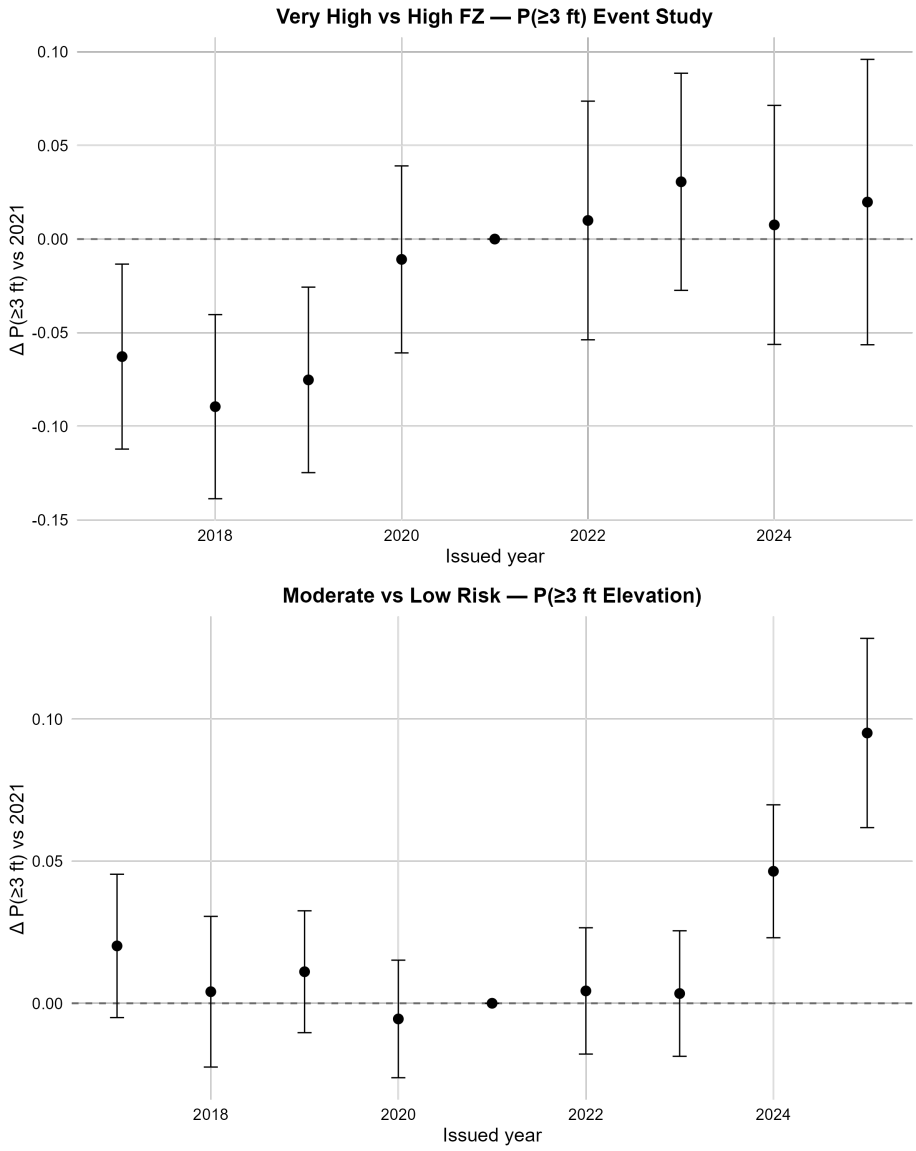
Notes: Scatter plot of projected 2070 percentage premium changes (relative to current RR2.0 levels) for a random sample of NFIP policies under two climate models (MPI on horizontal axis; IPSL on vertical), using synthetic flood data from Gori et al. (2021) under SSP5-8.5 warming scenario. Dashed 45-degree line represents pure aggregate risk (uniform scaling across households); deviations above/below reflect idiosyncratic reclassification risk, while cross-model flips (e.g., high in one, low in the other) illustrate diversifiability across households.

Table 5. Reclassification Risk Results

Coefficient of Absolute Risk Aversion	Reclassification Risk Insurance Value (Level)
5×10^{-5}	\$8.5
1×10^{-4}	\$18.7
2×10^{-4}	\$39.8
5×10^{-4}	\$64.5

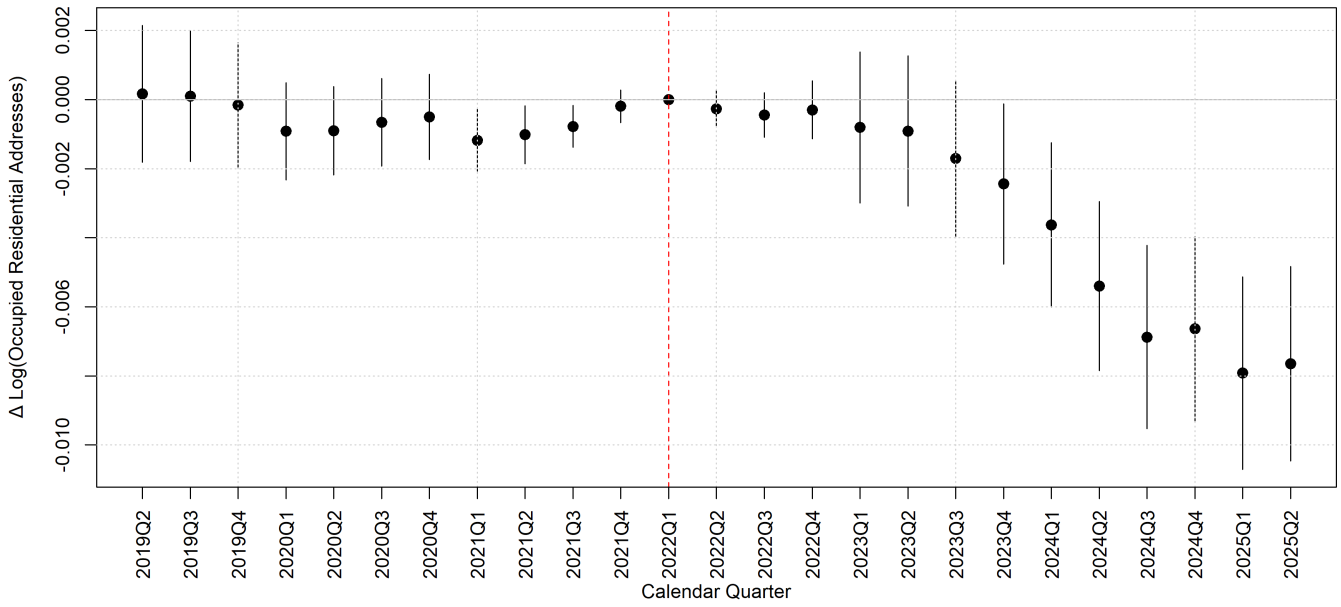
Notes: Each row reports the pure insurance value (in 2023 dollars) of reinstating a 50% average subsidy against idiosyncratic reclassification risk, net of mechanical transfers (c_i), under constant absolute risk aversion (CARA) utility $u(c) = -\exp(-\gamma c)$. γ values span literature estimates: 5×10^{-5} (lower bound, Cohen et al. 2007); 5×10^{-4} (central estimate, conservative relative to Sydnor 2010 and Barseghyan et al. 2011). At central $\gamma = 5 \times 10^{-4}$, value is \$64.5 (7% of pre-RR2.0 premium, 1% of projected future premium).

Figure 4. Mitigation Changes Under RR2.0



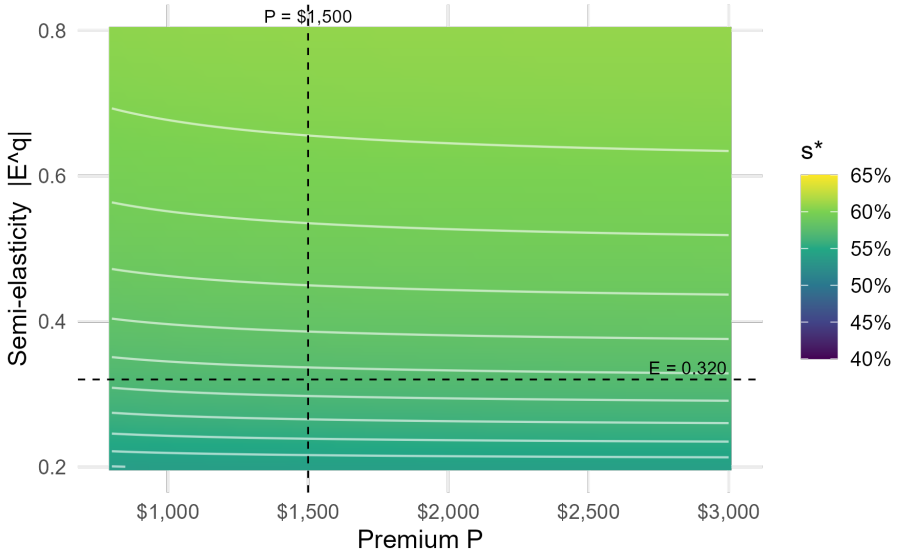
Notes: Event-study estimates from regressions of an indicator for elevating the lowest floor by 3+ feet (relative to base flood elevation inside SFHA or local ground outside) on indicators for years relative to RR2.0 implementation (2022 as reference), interacted with higher-risk flood zone category, separately inside (top panel) and outside (bottom panel) Special Flood Hazard Areas (SFHAs). Inside SFHA: higher-risk = very high (V zones with velocity waves) vs. high (A zones); outside: moderate (shaded X) vs. low (unshaded X). Sample: all Florida elevation certificates filed 2017–2024. The estimating equation is (3).

Figure 5. Location Changes Under RR2.0



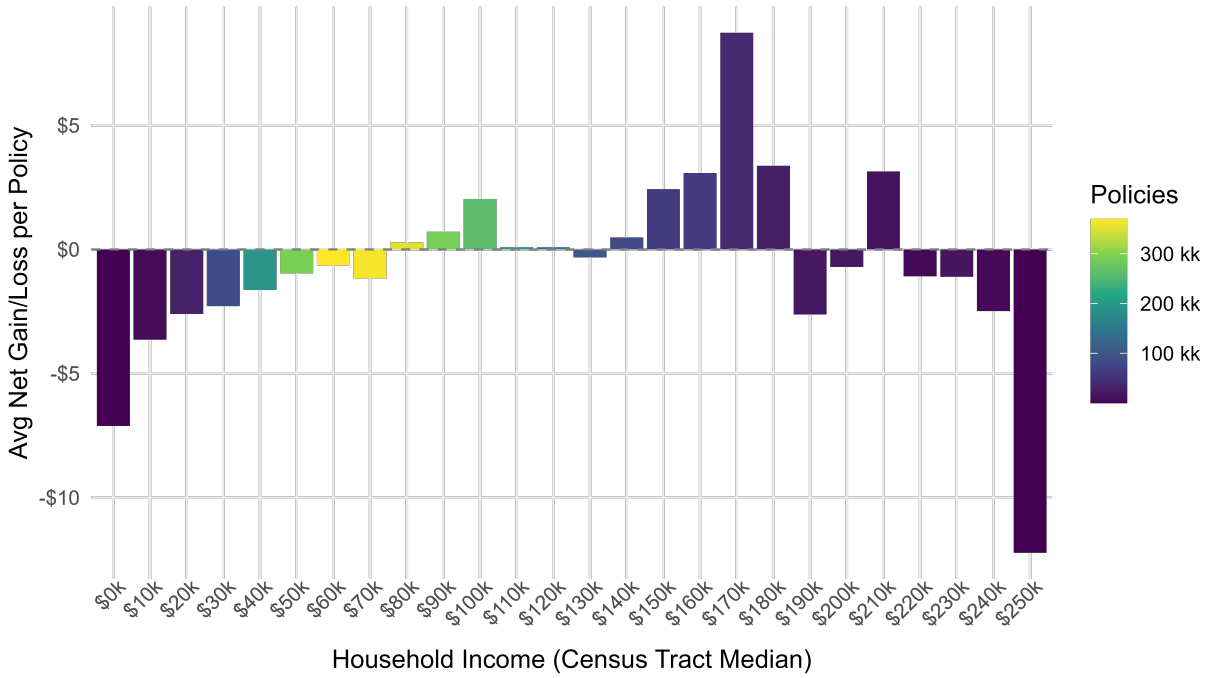
Notes: Event-study estimates from regressions of $\log(1 + \text{OccupiedResidential}_{c,t,q})$ (quarterly occupied residential addresses from USPS) on indicators for quarters relative to Q1 2022 (RR2.0 rollout completion), interacted with $\text{MedianIncrease}_{ct}$ (indicator for census tract c where median NFIP premium increased under RR2.0).

Figure 6. Optimal Subsidy



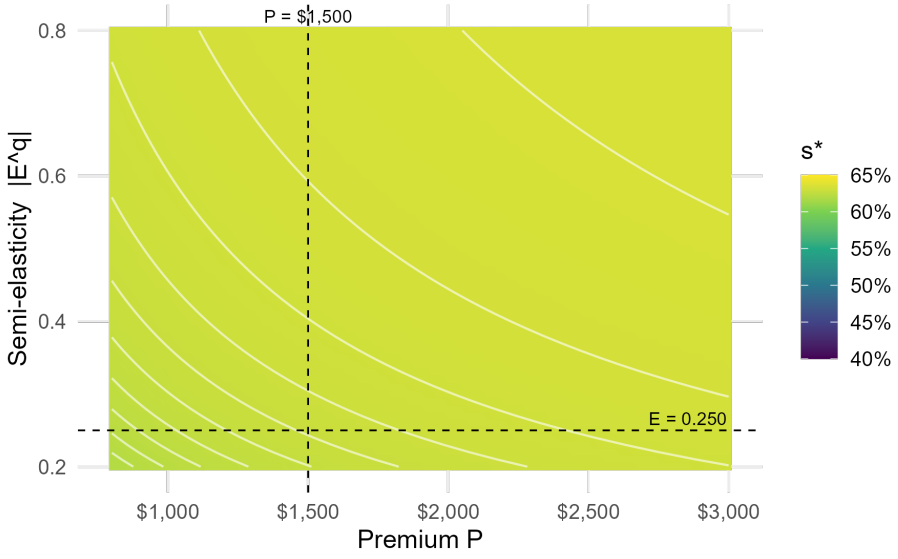
Notes: Heatmap of optimal subsidy rate s that solving the calibrated benefit-cost condition from Proposition 2 for discrete household choices (insure/uninsure; elevate/not; high-/low-risk location), across combinations of average full-risk premium among marginals $\bar{P}$ (horizontal axis, \$800 – \$3,000) and insurance demand semi-elasticity $\mathcal{E}^q$ (vertical axis, -0.2 to -8).

Figure 7. Redistribution.....



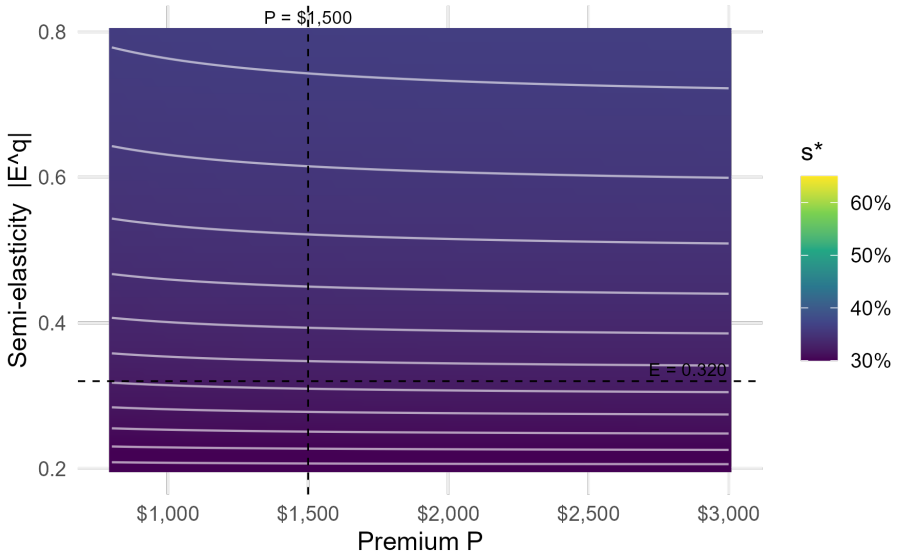
Notes: Bar chart of average net gain/loss (height, 2023 dollars) from a simulated 1% subsidy on NFIP premiums (funded by lump-sum tax on all insured households), binned by census tract median household income (x-axis bins: $0–$50k, $50k–$70k, etc., up to $170k+). Color intensity represents number of policies (sample size) per bin. Simulation assigns tract median income (ACS) to each NFIP policy; net gain/loss = subsidy received minus tax paid, averaged within bin.

Figure 8. Optimal Subsidy with SFHA Restrictions Imposed



Notes: Analogous to Figure 6, but incorporates SFHA regulations: mitigation elasticity set to zero (mandated) and demand semi-elasticity to -0.25.

Figure 9. Optimal Subsidy with Upper Bounds on Costs and Lower Bounds on Benefits



Notes: Analogous to Figure 6, but with 'worst-case' assumptions: location elasticity -0.015 (Boustan et al. 2020 upper bound); mitigation externality 2x baseline; reclassification WTP discounted 2%/year over 35 years; FEMA spillover fixed at -$116/pp (minimum from heterogeneity, Appendix A.5).

References

(2023). Circular no. a-4: Regulatory analysis. Circular A-4, Office of Management and Budget, Executive Office of the President, Washington, DC. Issued November 9, 2023.

Amornsiripanitch, N., Biswas, S., Orellana, J., and Zink, D. (2024). Flood underinsurance. Technical Report Working Paper 24-23, Federal Reserve Bank of Philadelphia.

Aron-Dine, S. (2025). Rebuild or relocate? recovery after natural disasters. Job market paper, Stanford University.

ASFPM (2021). New interactive maps provide visibility into flood insurance premium changes coming with fema's risk rating 2.0. *Association of State Floodplain Managers (ASFPM)*.

Barseghyan, V., Molinari, F., and Silverman, J. (2011). Are risk preferences stable across contexts? evidence from insurance choices. *American Economic Review*, 101(2):591–630.

Ben-Shahar, O. and Logue, K. D. (2016). The perverse effects of subsidized weather insurance. *Stanford Law Review*, 68:571–624.

Boustan, L. P., Kahn, M. E., Rhode, P. W., and Yanguas, M. L. (2020). The effect of natural disasters on economic activity in us counties: A century of data. *Journal of Urban Economics*, 118:103257.

Browne, M. J., Dehring, C. A., Eckles, D. L., and Lastrapes, W. D. (2019). Does national flood insurance program participation induce housing development? *Journal of Risk and Insurance*, 86(4):835–859.

Cochrane, J. H. (1995). Time-consistent health insurance. *Journal of Political Economy*, 103(3):445–473.

Cohen, A. and Einav, L. (2007). Estimating risk preferences from deductible choice. *American Economic Review*, 97(3):745–788.

Congressional Budget Office (2017). The national flood insurance program: Financial soundness and affordability. Technical report, Congress of the United States.

Congressional Research Service (2023). Introduction to the national flood insurance program (nfip). Technical Report R44593, Congress of the United States. Accessed: September 9, 2025.

Deryugina, T., Kawano, L., and Levitt, S. (2018). The economic impact of hurricane katrina on its victims: Evidence from individual tax returns. *American Economic Journal: Applied Economics*, 10(2):202–233.

Deryugina, T. and Kirwan, B. (2018). Does the samaritan's dilemma matter? evidence from u.s. agriculture. *Economic Inquiry*, 56(2):983–1006.

Fabian, J. (2024). The price of risk: Flood insurance premium reform and local development. Job Market Paper, Brown University.

Federal Emergency Management Agency (2023). The watermark: Official publication of the national flood insurance program. Technical report, U.S. Department of Homeland Security. FY2023 Q4 Edition. Accessed: September 9, 2025.

Federal Reserve Bank of Philadelphia (2016). The Uninsured Cost of Flooding: A State-by-State Analysis.

Florida Division of Emergency Management (2025). Florida elevation certificate database. https://www.floridadisaster.org/mitigation/floodplain-management/elevation-certificates/. Accessed: September 9, 2025.

Friedman, T. L. (2010). Global weirding. *The New York Times*. Accessed on September 14, 2025.

Gallagher, J. (2014). Learning about an infrequent event: Evidence from flood insurance take-up in the united states. *American Economic Journal: Applied Economics*, 6(3):206–33.

Ge, S., Lam, A., and Lewis, R. (2022). The effect of insurance premiums on the housing market and climate risk pricing. *Available at SSRN 4192699*. Last revised: 17 Mar 2025.

Ghili, S., Handel, B., Hendel, I., and Whinston, M. D. (2023). Optimal long-term health insurance contracts: Characterization, computation, and welfare effects. *The Review of Economic Studies*, 91(2):1085–1121.

Gori, A., Lin, N., Xi, D., and Emanuel, K. (2022). Tropical cyclone climatology change greatly exacerbates US extreme rainfall–surge hazard. *Nature Climate Change*, 12(2):171–178.

Handel, B., Hendel, I., and Whinston, M. D. (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica*, 83(4):1261–1313.

Hendel, I. and Lizzeri, A. (2003). The role of commitment in dynamic contracts: Evidence from life insurance. *The Quarterly Journal of Economics*, 118(1):299–327.

Henkel, M., Kwon, E., and Magontier, P. (2024). The Unintended Consequences of Post-Disaster Policies for Spatial Sorting. MIT Center for Real Estate Research Paper No. 22/08, Available at SSRN: https://ssrn.com/abstract=4237083 or http://dx.doi.org/10.2139/ssrn.4237083.

Hudson, P., Botzen, W. J. W., Kreibich, H., and Aerts, J. C. J. H. (2017). Moral hazard in natural disaster insurance markets: Empirical evidence from germany and the united states. *Land Economics*, 93(2):179–208.

IPCC (2021). Summary for policymakers. Technical report, Intergovernmental Panel on Climate Change, Geneva, Switzerland. In: Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on

Climate Change [Masson-Delmotte, V., P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu and B. Zhou (eds.)]. Cambridge University Press.

Jordà, Ò., Knoll, K., Kuvshinov, D., Schularick, M., and Taylor, A. M. (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics*, 134(3):1225–1298.

Kousky, C. (2010). Learning from extreme events: Risk perceptions after the flood. *Land Economics*, 86(3):395–422.

Kousky, C. (2019). The role of natural disaster insurance in recovery and risk reduction. *Annual Review of Resource Economics*, 11(Volume 11, 2019):399–418.

Kousky, C., Michel-Kerjan, E., and Raschky, P. A. (2018). Does federal disaster assistance crowd out flood insurance? *Journal of Environmental Economics and Management*, 87:71–90.

Kunreuther, H. and Pauly, M. (2004). A behavioral economics perspective on the disaster insurance puzzle. *The American Economic Review*, 94(2):400–404.

Landry, C. E. and Jahan-Parvar, M. R. (2011). Flood insurance coverage in the coastal zone. *Journal of Risk and Insurance*, 78(2):361–388.

Lee, J. (2024). Homeowners' responses to the new price of flood insurance. Job market paper / AEA-AFA materials.

Mathis, M. L. and Nicholson, S. (2006). An evaluation of compliance with the national flood insurance program: Part b are minimum building requirements being met? Technical report, American Institutes for Research, Washington, D.C. Prepared for the Federal Emergency Management Agency (FEMA).

Mulder, P. and Kousky, C. (2023). Risk rating without information provision. *AEA Papers and Proceedings*, 113:299–303.

NAIC (2024). 2024 Homeowners Insurance Premiums Data. Accessed via the NOAA news report cited.

National Association of Insurance Commissioners (2024). Private flood insurance data call report. Technical report, Center for Insurance Policy and Research (CIPR). Report on data for the 2023 calendar year. Accessed: September 9, 2025.

NOAA (2025). Billion-Dollar Weather and Climate Disasters: Overview. Data as of January 10, 2025.

Peralta, A. and Scott, J. B. (2024). Does the national flood insurance program drive migration to higher risk areas? *Journal of the Association of Environmental and Resource Economists*, 11(2):287–318.

Philippi, T. and Schiller, J. (2024). Abandoning disaster relief and stimulating insurance demand through premium subsidies. *Journal of Risk and Insurance*, 91(2):339–382.

Raschky, P. A., Schwarze, R., Schwindt, M., and Zahn, F. (2013). Uncertainty of government relief and the crowding out of private flood insurance. *Environmental and Resource Economics*, 54(2):179–200.

Reinsurance News (2024). Total nfip losses from 2024 flood events could reach $10.5bn: Neptune. *Reinsurance News*. Accessed on September 14, 2025.

Reuters (2025). Trump says FEMA to be wound down after hurricane season. *Reuters*.

Sydnor, J. (2010). (over)insuring modest risks. *American Economic Journal: Applied Economics*, 2(4):177–199.

U.S. Army Corps of Engineers (2021). National economic development (ned) procedures manual - national flood risk management. Technical Report IWR Report 2021-R-01, Institute for Water Resources, Alexandria, VA.

U.S. Federal Emergency Management Agency (2025). Disaster Relief Fund: Fiscal Year 2024 Annual Report. Technical report, U.S. Department of Homeland Security, Washington, D.C.

U.S. Government Accountability Office (2017). Flood insurance: Fema needs to improve efforts to address noncompliance with purchase requirements. Technical Report GAO-17-805T, U.S. Government Accountability Office. Accessed: September 9, 2025.

U.S. Government Accountability Office (2022). National flood insurance program: Options for improving the financial condition and addressing affordability. Technical Report GAO-22-105342, U.S. Government Accountability Office. Accessed: September 9, 2025.

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

A Empirical Appendix

Table A1. Main Results: Total Fiscal Cost

	OLS	**IV**
NFIP Policies per 100 Houses	-175.468^{***}	-221.144^{**}
	(33.363)	(75.604)
Damage Control	Binned	
First-stage F-stat.		152.17
Num. Obs.	783	783
Within R^2	0.916	0.911

Notes: Standard errors in parentheses. Significance: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$.

Table A2. Fiscal Cost Decomposition: FEMA Disaster Assistance

	OLS	**IV**
NFIP Policies per 100 Houses	-150.659^{***}	-169.732^{**}
	(29.187)	(57.250)
Damage Control	Binned	
First-stage F-statistic		183.11
Num. Obs.	815	815
Within R^2	0.907	0.906

Notes: Standard errors in parentheses. Significance: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$.

Table A3. Fiscal Cost Decomposition: SBA Fiscal Cost

	OLS	**IV**
NFIP Policies per 100 Houses	-18.537^{*}	-10.443
	(8.559)	(17.998)
Damage Control	Binned	
Num. Obs.	815	815
Within R^2	0.801	0.795

Notes: Standard errors in parentheses. Significance: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

Table A4. Fiscal Cost Decomposition: HGMP Payouts

	OLS	**IV**
NFIP Policies per 100 Houses	-2.917	-21.664
	(9.751)	(19.516)
Damage Control	Binned	
Num. Obs.	815	815
Within R^2	0.367	0.342

Notes: Standard errors in parentheses. Significance: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$.

Table A5. Fiscal Cost Decomposition: Freddie Mac

	OLS	**IV**
NFIP Policies per 100 Houses	0.611^{*}	0.603
	(0.288)	(0.835)
Damage Control	Binned	
First-stage F-statistic		173.83
Num. Obs.	810	810
Within R^2	0.259	0.259

Notes: Standard errors in parentheses. Significance: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$.

Table A6. Fiscal Cost Decomposition: Fannie Mae

	OLS	**IV**
NFIP Policies per 100 Houses	0.244	-1.212
	(0.161)	(0.995)
Damage Control	Binned	
First-stage F-statistic		167.41
Num. Obs.	809	809
Within R^2	0.700	0.597

Notes: Standard errors in parentheses. Significance: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$.

Table A7. Robustness to Different Damage Controls

Panel A: Baseline, Cont. (Cubic), Baseline + Cubic

	Baseline		**Cont. (Cubic)**		**Baseline + Cubic**	
	OLS	**IV**	**OLS**	**IV**	**OLS**	**IV**
NFIP Policies per 100 Houses	-171.5^{***}	-203.2^{**}	-221.4^{**}	-304.8^{**}	-166.6^{***}	-207.1^{**}
	(34.672)	(73.416)	(84.873)	(97.919)	(39.899)	(72.232)
Num. Obs.	815	815	815	815	815	815
Within R^2	0.915	0.912	0.831	0.813	0.916	0.913

Panel B: Cont. (Quintic), Log-Linear, Hybrid

	Cont. (Quintic)		**Log-Linear**		**Hybrid**	
	OLS	**IV**	**OLS**	**IV**	**OLS**	**IV**
NFIP Policies per 100 Houses	-189.3^{*}	-296.9^{**}	-320.6^{***}	-382.7^{***}	-159.2^{**}	-217.0^{**}
	(85.751)	(100.095)	(30.326)	(22.572)	(53.146)	(79.119)
Num. Obs.	815	815	815	815	815	815
Within R^2	0.845	0.820	0.783	0.762	0.892	0.884

Panel C: Baseline + Log, Wide Bins, Piece-wise Linear Steps

	Baseline + Log		**Wide Bins**		**Piece-wise Linear Steps**	
	OLS	**IV**	**OLS**	**IV**	**OLS**	**IV**
NFIP Policies per 100 Houses	-246.0^{***}	-345.3^{***}	-231.6^{**}	-307.0^{***}	-165.8^{***}	-211.8^{**}
	(30.786)	(29.042)	(86.586)	(78.068)	(47.868)	(77.730)
Num. Obs.	815	815	815	815	815	815
Within R^2	0.823	0.794	0.853	0.839	0.887	0.883

Notes: Standard errors in parentheses. Significance: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$.

A.1 Testing Assumption A2

We empirically test Assumption A2 – whether household–level idiosyncratic premium residuals tend to be high in bad aggregate climate states by computing. We compute, for each policy i, the correlation across scenarios ω between the state aggregate marginal utility $\bar{u}'_T(\omega)$ and the residual $\varepsilon_{i\omega}$ from the premium decomposition

$$P_i(\omega) = \bar{P}(\omega) + a_i + \varepsilon_{i\omega}, \qquad \bar{P}(\omega) = \sum_i \tilde{w}_i \, P_i(\omega), \;\; \tilde{w}_i \propto B_i,$$

where B_i is building+contents exposure and $\sum_i \tilde{w}_i = 1$. Predicted premiums $P_i(\omega)$ come from the IDW hazard mapping and simple regressions evaluated for eight scenarios (equal scenario weights $w_\omega = 1/8$).

Utility is either assumed to be of CRRA form with $\gamma = 2$ or CARA with $a = 10^{-4}$. The state average marginal utility is $\bar{u}'_T(\omega) = \sum_j \theta_j(\omega) \, u'_j(c_j(\omega))$, under (i) equal per–policy incidence $\theta_j(\omega) = 1/N_\omega$ and (ii) exposure–weighted incidence $\theta_j(\omega) \propto P_j$ (renormalized within ω).

Correlation and aggregation. For each household i, we compute

$$\mathsf{Corr}_\omega(\bar{u}'_T, \varepsilon_{i\omega}) = \frac{\mathsf{Cov}_\omega(\bar{u}'_T, \varepsilon_{i\omega})}{\sqrt{\mathsf{Var}_\omega(\bar{u}'_T)\,\mathsf{Var}_\omega(\varepsilon_{i\omega})}},$$

with equal scenario weights. We summarize these per–i correlations by the (i) unweighted mean across i, and (ii) premium–weighted mean. The results are in Table A8.

Table A8. Correlation across scenarios: $\mathsf{Corr}_\omega(\bar{u}'_T, \varepsilon_{i\omega})$

	N	Unweighted Corr (Mean)	Cash–Weighted Corr (WMean)
CRRA	609,562	0.0363	0.0726
CARA	609,562	0.0920	0.1199

Notes: Correlations are computed per policy across eight scenarios (equal scenario probabilities). “Cash–weighted” averages use premium weights. T

We find strong empirical support for a negligible correlation, justifying Assumption A2.

A.2 Robustness to Discounting Assumptions

The premium risk individuals face due to future climate uncertainty is measured in 2070. Nevertheless, we are solving for the optimal subsidy today. In the main paper we do not discount this by appealing to an offsetting growth in house prices. To test the robustness of our policy prescriptions to this assumption, here we discount the value of reclassification risk by 2% a year for 35 years. That is, the reclassification risk value is multiplied by $1/(1.02)^35 \approx 0.5$ and then the optimal subsidy

is computed as before. The analogue of Figure 6 with the discounted reclassification risk values is Figure 10, which follows.

Figure 10. Optimal Subsidy with SFHA Restrictions Imposed

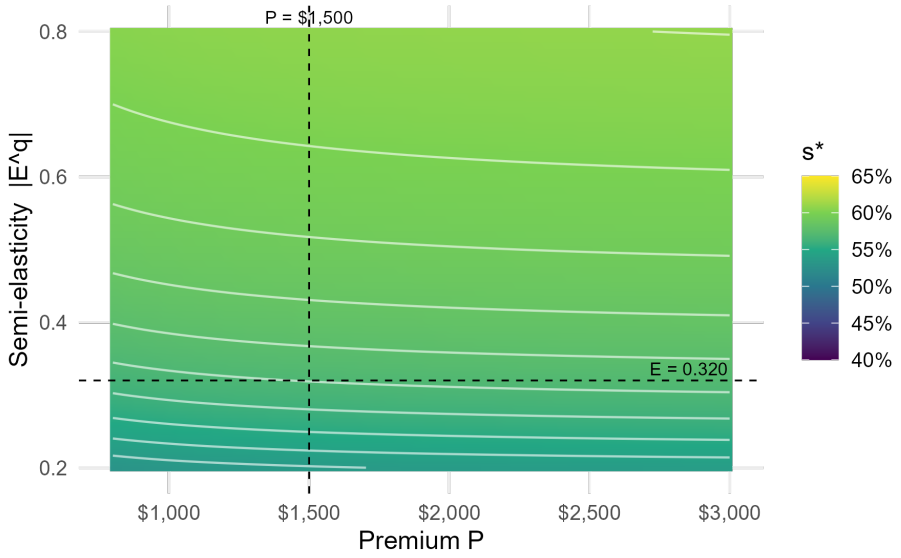


Figure 10 shows that the optimal subsidy changes falls by less than 1% at any combination of $(\mathcal{E}, \bar{P})$ considered, and by less than 0.25% at the benchmark values of $\mathcal{E} = 0.32, \bar{P} = \$1,500$.

A.3 Flood Hazard to Premium Regression

Table A9. Regression Mapping Hazard Factors to RR2.0 Premiums

	IDW
100-yr Rainfall (mm)	3.926^{***}
	(0.104)
100-yr Storm Tide (mm)	1.068^{***}
	(0.007)
Annual Joint Probability	36763.045^{***}
	(6006.871)
Num. Obs.	879,021
R^2	0.344

Notes: The dependent variable is the RR2.0 full-risk premium. Standard errors are reported in parentheses.
Significance: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

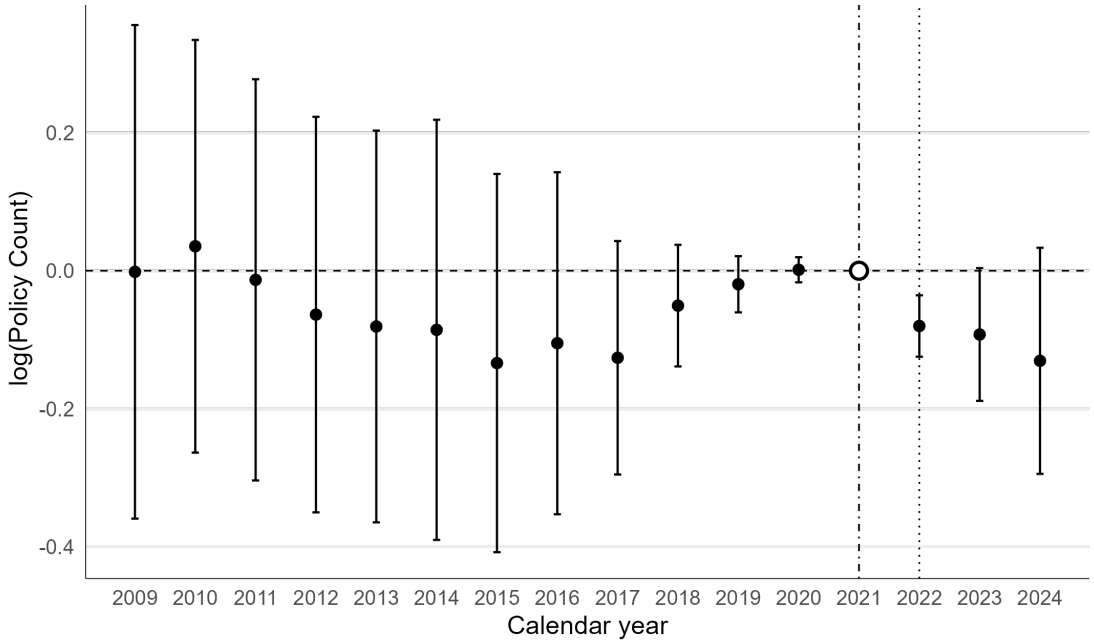
A.4 Between-County Demand Elasticity Estimates

We proceed analogously to the analysis in Section 8, except at the county level. We compare the number of NFIP policies in county c measured in year t, NFIP Policies$_{c,t}$, across counties that on average experienced a price increase versus a decrease under RR2.0. Specifically, we estimate

$$\log(\text{NFIP Policies}_{c,t}) = \alpha_c + \gamma_t + \sum_{t' \neq 2021} \beta_{t'} \Big(\text{MedianIncrease}_{ct} \times \mathbf{1}\{t = t'\} \Big) + \varepsilon_{c,t}, \tag{6}$$

where α_c are county fixed effects, γ_t are year fixed effects, and we weight by premium. The coefficients of interest are $\beta_{t'}$, which capture the relative difference in occupied residences in the more- versus less-exposed census tracts, normalized to zero in 2021. Figure 11 plots the estimated $\beta_{q'}$.

Figure 11. Between-County Demand Estimates



We see, by the end of 2024, a drop in demand that is 13% higher in the treated counties than the control. The treated counties have average pre- and post-RR2.0 premiums of $\$918$ and $\$1,808$ respectively, while the control counties have $\$688$ and $\$904$. In other words, the relative price increase is 96% and 31% in the treatment and control counties respectively. Hence, a 13% greater demand drop relative to a 65% greater price drop implies an elasticity of 0.2. This is similar to our within-policy estimates, and to the extant literature.

A.5 Heterogeneity in the Fiscal Spillover

In the main paper we assumed that the spillover from more insurance coverage onto ex-post assistance was constant (in the level of subsidy s). In other words, the identity of the marginal insured did not affect the spillover. Here we loosen that restriction by estimating:

$$\begin{aligned}\text{ExPost Spending}_{c,f} = \alpha_c + \gamma_t + \beta_0\ &\text{NFIP Coverage}_{c,f} \\ &+ \beta_1\ \text{NFIP Coverage}_{c,f} \times \frac{\text{PreRR2.0 Price}}{\text{PostRR2.0 Price}}_c \\ &+ f(\text{Flood Damage}_{c,f}) + \epsilon\end{aligned}$$

The coefficients of interest are β_0, the fiscal spillover per percentage of NFIP coverage with a 100% subsidy, and β_1 the fiscal spillover term interacted with a change of the subsidy from 100% to zero. The results are in Table A10 below.

Table A10. Heterogeneity Results: Total Fiscal Cost

	OLS	**IV**
NFIP Policies per 100 Houses	-228.102^{*}	-463.722
	(115.3)	(421.0)
NFIP Policies per 100 Houses $\times$ Pre/Post RR2.0 Price Ratio	112.960	645.903
	(181.50)	(974.61)
Damage Control	Binned	
First-stage F-stat.		128.44
Num. Obs.	783	783

Notes: Standard errors in parentheses. Significance: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$.

Table A10 shows that the spillover is more pronounced a higher subsidy levels. In a (hypothetical) county with a 100% subsidy, every additional percentage of NFIP coverage reduces spillovers by about $228. In contrast, if there is no subsidy, an additional percentage of coverage reduces ex-post spillovers by only $115. These effects are relatively imprecisely estimated, since they are estimated from only 783 floods. The IV estimates ae similar in sign but very imprecisely estimated.

B Proofs

Proof of Proposition 1

Recall the public outlays

$$\mathsf{Out}(\omega; s) = \sum_i s\, q_i\, P_i(\omega) + \sum_i f\,(1 - q_i)\, P_i(\omega).$$

The planner's per–state budget is $\sum_i T_i(\omega) = \mathsf{Out}(\omega; s)$, and the Lagrangian is

$$\mathcal{L}(s, \{\mu(\omega)\}) = \mathbb{E}_\omega \mathbb{E}_\xi \Big[\sum_i u(c_i(\omega, \xi; s)) \Big] + \mathbb{E}_\omega \Big[\mu(\omega) (\sum_i T_i(\omega) - \mathsf{Out}(\omega; s)) \Big],$$

with the normalization $\mu(\omega) = \bar{u}'(\omega) \equiv \sum_i \mathbb{E}_\xi u'(c_i(\omega, \xi))$. We use the decomposition

$$P_i(\omega) = \bar{P}(\omega) + a_i + \varepsilon_{i\omega}, \qquad \mathbb{E}_\omega[\varepsilon_{i\omega}] = 0, \quad \sum_i \varepsilon_{i\omega} = 0 \ \ \forall \omega.$$

Within-state actuarial fairness holds for each ω:

$$P_i(\omega) = \mathbb{E}_\xi[L_i(\omega, \xi) \mid \omega], \qquad \frac{\partial P_i}{\partial a}(\omega) = \mathbb{E}_\xi \Big[\frac{\partial L_i}{\partial a}(\omega, \xi) \Big| \omega \Big], \ a \in \{m, x\}.$$

We denote semi–elasticities by

$$\mathcal{E}_i^q \equiv \frac{\partial q_i^*}{\partial \ln \mathbb{E}_\omega[\tilde{P}_i]}, \qquad \mathcal{E}_i^m \equiv \frac{\partial m_i^*}{\partial \ln \mathbb{E}_\omega[\tilde{P}_i]}, \qquad \mathcal{E}_i^x \equiv \frac{\partial x_i^*}{\partial \ln \mathbb{E}_\omega[\tilde{P}_i]}.$$

Note $d \ln \mathbb{E}_\omega[\tilde{P}_i]/ds = -1/(1 - s)$. For tax/premium neutrality: Fix a state ω. For any perturbation that changes only $T_i(\omega)$ by $\Delta T_i(\omega)$ (holding choices, losses, and outlays fixed), the first-order welfare change in state ω is zero:

$$\sum_i \mathbb{E}_\xi u_i'(\omega, \xi)\, \Delta(-T_i(\omega)) + \mu(\omega) \sum_i \Delta T_i(\omega) = 0,$$

immediate from $\mu(\omega) = \sum_i \mathbb{E}_\xi u_i'(\omega, \xi)$.

Differentiate $\mathcal{L}$ w.r.t. s. Write $\dot{(\cdot)} \equiv \frac{d}{ds}(\cdot)$.

Step 1: Direct (holding choices fixed). Using $c_i(\omega, \xi) = y_i + \Gamma_i(x_i) - C_i(m_i) - q_i \tilde{P}_i(\omega; s, x_i, m_i) - (1 - f)(1 - q_i)\, L_i(\omega, \xi; x_i, m_i) \ - \ T_i(\omega)$, noting that the only direct dependence on s is through $\tilde{P}_i$ and T_i. In particular, $\tilde{P}_i = (1 - s)P_i$, hence $\partial \tilde{P}_i / \partial s = -P_i(\omega)$ and we keep $\dot{T}_i(\omega) \equiv \frac{d}{ds}(T_i(\omega))$ general:

$$\frac{d}{ds} \mathbb{E}_\omega \mathbb{E}_\xi[u(c)] = \mathbb{E}_\omega \mathbb{E}_\xi \Big[\sum_i u_i'(\omega, \xi)\, (-q_i)\, (-P_i(\omega)) \Big] + \mathbb{E}_\omega \mathbb{E}_\xi \Big[\sum_i u_i'(\omega, \xi)\, (-\dot{T}_i(\omega)) \Big].$$

The budget term contributes

$$\frac{d}{ds}\mathbb{E}_\omega[\mu(\sum T_i - \text{Out})] = \mathbb{E}_\omega\Big[\mu(\omega)\sum_i \dot{T}_i(\omega)\Big] - \mathbb{E}_\omega\Big[\mu(\omega)\,\frac{\partial \text{Out}}{\partial s}\Big]_{\text{choices fixed}}.$$

Holding choices fixed, since $\text{Out}(\omega; s) = \sum_i s q_i P_i(\omega) + \sum_i f(1-q_i)P_i(\omega)$, we have $\partial\text{Out}/\partial s = \sum_i q_i P_i(\omega)$.

Summing the utility and budget effects we have the direct effect D as:

$$D = \mathbb{E}_\omega\mathbb{E}_\xi\Big[\sum_i u_i'(\omega,\xi)\,(-q_i)\,(-P_i(\omega))\Big] + \mathbb{E}_\omega\mathbb{E}_\xi\Big[\sum_i u_i'(\omega,\xi)\,(-\dot{T}_i(\omega))\Big] + \mathbb{E}_\omega\Big[\mu(\omega)\sum_i \dot{T}_i(\omega)\Big] - \mathbb{E}_\omega\Big[\mu(\omega)\,\frac{\partial \text{Out}}{\partial s}\Big].$$

Note that $\mathbb{E}_\omega\mathbb{E}_\xi\Big[\sum_i u_i'(\omega,\xi)\,(-\dot{T}_i(\omega))\Big] = \mathbb{E}_\omega\Big[(-\dot{T}_i(\omega))\mathbb{E}_\xi\Big[\sum_i u_i'(\omega,\xi)\Big]\Big] = \mathbb{E}_\omega\left[(-\dot{T}_i(\omega))\bar{u}'(\omega)\right]$ and therefore since, by assumption, $\mu(\omega) = \bar{u}'(\omega) = \mathbb{E}_\xi\left[u_i'(\omega,\xi)\right]$ the tax terms cancel and we are left with

$$D \equiv \mathbb{E}_\omega\Big[\sum_i (\mathbb{E}_\xi u_i'(\omega,\xi) - \mu(\omega))\,q_i\,P_i(\omega)\Big].$$

Insert the decomposition $P_i = \bar{P} + a_i + \varepsilon_{i\omega}$ to obtain

$$\begin{aligned} D = {} & \mathbb{E}_\omega\Big[\bar{P}(\omega)\sum_i (u_i' - \mu)q_i\Big] + \mathbb{E}_\omega\Big[\sum_i (u_i' - \mu)q_i a_i\Big] \\ & + \mathbb{E}_\omega\Big[\sum_i (u_i' - \mu)q_i\varepsilon_{i,\omega}\Big], \end{aligned}$$

where $u'i = \mathbb{E}\xi u_i'(\omega,\xi)$.

Insert the decomposition $P_i = \bar{P} + a_i + \varepsilon_{i\omega}$ to obtain

$$\begin{aligned} D = {} & \mathbb{E}_\omega\Big[\bar{P}(\omega)\sum_i (u_i' - \mu)q_i\Big] + \mathbb{E}_\omega\Big[\sum_i (u_i' - \mu)q_i a_i\Big] \\ & + \mathbb{E}_\omega\Big[\sum_i (u_i' - \mu)q_i\varepsilon_{i,\omega}\Big], \end{aligned}$$

where $u_i' = \mathbb{E}_\xi u_i'(\omega,\xi)$.

We now rewrite each term using covariances over ω.

Starting with the third term, again using $\mu(\omega) = \bar{u}'(\omega)$:

$$\mathbb{E}_\omega\Big[\sum_i (u_i' - \mu)q_i\varepsilon_{i,\omega}\Big] = \mathbb{E}_\omega\Big[\sum_i (u_i' - \bar{u}'(\omega))q_i\varepsilon_{i,\omega}\Big] = \sum_i q_i\mathbb{E}_\omega\Big[(u_i' - \bar{u}'(\omega))\varepsilon_{i,\omega}\Big].$$

Since $\mathbb{E}_\omega[\varepsilon_{i\omega}] = 0$,

$$\mathbb{E}_\omega[u_i'\varepsilon_{i,\omega}] = \text{Cov}_\omega(u_i', \varepsilon_{i\omega}), \quad \mathbb{E}_\omega[\bar{u}'(\omega)\varepsilon_{i\omega}] = \text{Cov}_\omega(\bar{u}'(\omega), \varepsilon_{i\omega}).$$

Thus, the third term is

$$\sum_i \Big(\mathsf{Cov}_\omega(u_i', q_i\varepsilon_{i\omega}) - q_i\mathsf{Cov}_\omega(\bar{u}'(\omega), \varepsilon_{i\omega})\Big).$$

For the second term:

$$\mathbb{E}_\omega\Big[\sum_i (u_i' - \bar{u}'(\omega))q_i a_i\Big] = \sum_i q_i a_i \mathbb{E}_\omega[u_i' - \bar{u}'(\omega)] = \sum_i q_i a_i \Big(\mathbb{E}_\omega[u_i'] - \mathbb{E}_\omega[\bar{u}'(\omega)]\Big).$$

For the first term:

$$\mathbb{E}_\omega\Big[\bar{P}(\omega)\sum_i (u_i' - \bar{u}'(\omega))q_i\Big] = \mathsf{Cov}_\omega\Big(\bar{P}(\omega),\, \sum_i q_i(u_i'(\omega) - \bar{u}'(\omega))\Big) + \mathbb{E}_\omega[\bar{P}(\omega)]\cdot\mathbb{E}_\omega\Big[\sum_i q_i(u_i' - \bar{u}'(\omega))\Big].$$

This produces the average selection ω-specific selection terms.

Combining these, we have

$$\begin{aligned} D = &\sum_i \Big(\mathsf{Cov}_\omega(u_i',\ q_i\varepsilon_{i\omega}) - q_i\,\mathsf{Cov}_\omega(\bar{u}',\ \varepsilon_{i\omega})\Big) \\ &+ \sum_i q_i a_i \Big(\mathbb{E}_\omega[u_i'] - \mathbb{E}_\omega[\bar{u}']\Big) \\ &+\ \mathbb{E}_\omega[\bar{P}(\omega)]\cdot\mathbb{E}_\omega\Big[\sum_i q_i(u_i' - \bar{u}'(\omega))\Big] \\ &+\ \mathsf{Cov}_\omega\Big(\bar{P}, \sum_i q_i(u_i' - \bar{u}')\Big). \end{aligned} \tag{A.1}$$

Setting $\frac{d\mathcal{L}}{ds} = 0$ at the optimum yields $D + B = 0$. From (A.6), $B = \frac{1}{1-s}\tilde{K}$, where $\tilde{K} = \mathbb{E}_\omega\Big[\mu(\omega)\Big(\sum_i ((s-f)\,P_i(\omega))\,\mathcal{E}_i^q + \sum_i (s\,q_i + f(1-q_i))\,\frac{\partial P_i(\omega)}{\partial m_i}\,\mathcal{E}_i^m + \sum_i (s\,q_i + f(1-q_i))\,\frac{\partial P_i(\omega)}{\partial x_i}\,\mathcal{E}_i^x\Big)\Big]$. Thus, $D = -\frac{1}{1-s}\tilde{K}$. Multiplying both sides by $(1-s)$ gives $(1-s)D = -\tilde{K}$, which matches the proposition with LHS from the covariance rewriting in (A.1) multiplied by $(1-s)$, and RHS from $-\tilde{K}$

Step 2: Behavioral (choices respond to s). Let $a_i \in q_i, m_i, x_i$. Total differentiation gives

$$\frac{d\mathcal{L}}{ds} = \underbrace{D}_{\text{Step 1}} + \sum_i \sum_{a\in\{q,m,x\}} \underbrace{\frac{\partial\mathcal{L}}{\partial a_i}}_{\text{FOC gap}} \cdot\frac{da_i}{ds}.$$

By the envelope theorem (households optimize over $\mathbb{E}_\omega\mathbb{E}_\xi[u(c)]$), $\partial\mathbb{E}_\omega\mathbb{E}_\xi[u(c)]/\partial a_i = 0$ (private terms cancel). What remains in $\partial\mathcal{L}/\partial a_i$ is the public-outlay piece valued at μ:

$$\frac{\partial\mathcal{L}}{\partial a_i} = -\mathbb{E}_\omega\Big[\mu(\omega)\,\frac{\partial\mathsf{Out}}{\partial a_i}\Big].$$

Compute the outlay derivatives (per state):

$$\frac{\partial\mathsf{Out}}{\partial q_i} = s\,P_i(\omega) - f\,P_i(\omega), \qquad \frac{\partial\mathsf{Out}}{\partial m_i} = s\,q_i\,\frac{\partial P_i}{\partial m_i}(\omega) + f(1-q_i)\,\frac{\partial P_i}{\partial m_i}(\omega),$$

$$\frac{\partial \text{Out}}{\partial x_i} = s\, q_i\, \frac{\partial P_i}{\partial x_i}(\omega) + f(1-q_i)\, \frac{\partial P_i}{\partial x_i}(\omega).$$

Express da_i/ds via semi–elasticities:

$$\frac{da_i}{ds} = \frac{\partial a_i}{\partial \ln \mathbb{E}_\omega[\tilde{P}_i]} \cdot \frac{d \ln \mathbb{E}_\omega[\tilde{P}_i]}{ds} = -\frac{1}{1-s}\, \mathcal{E}_i^a.$$

The behavioral contribution is

$$B = \sum_i \mathbb{E}_\omega\Big[\mu(\omega)\Big(-(s-f)\, P_i(\omega) \cdot \frac{dq_i}{ds} - (s\, q_i + f(1-q_i))\, \frac{\partial P_i(\omega)}{\partial m_i} \cdot \frac{dm_i}{ds} - (s\, q_i + f(1-q_i))\, \frac{\partial P_i(\omega)}{\partial x_i} \cdot \frac{dx_i}{ds}\Big)\Big].$$

Insert (A.5):

$$B = \frac{1}{1-s}\, \mathbb{E}_\omega\Big[\mu(\omega)\Big(\sum_i \left((s-f)\, P_i(\omega)\right) \mathcal{E}_i^q + \sum_i \left(s\, q_i + f(1-q_i)\right) \frac{\partial P_i(\omega)}{\partial m_i}\, \mathcal{E}_i^m + \sum_i \left(s\, q_i + f(1-q_i)\right) \frac{\partial P_i(\omega)}{\partial x_i}\, \mathcal{E}_i^x\Big)\Big].$$

Setting $d\mathcal{L}/ds = 0$ at the optimum and multiplying by (1-s) yields the proposition, with LHS from (A.1) and RHS from -(1-s) B from (A.6).

Proof of Proposition 2

Start from Proposition 1. Under A1-A4:

- A2 removes the macro–idiosyncratic correction $-(1-s)\sum_i q_i \text{Cov}_\omega(\bar{u}', \varepsilon_{i\omega})$ on the LHS.
- A4 removes the risk-class redistribution $(1-s)\sum_i q_i a_i(\mathbb{E}_\omega[u'i] - \mathbb{E}\omega[\bar{u}'])$.
- A3 kills the state-specific selection $(1-s)\text{Cov}_\omega\Big(\bar{P}, \sum_i q_i(u'i - \bar{u}')\Big)$.
- A1 ($\text{Cov}\omega(\mu(\omega), P_i(\omega)) = 0$ and for derivatives) allows decomposing the RHS: e.g., $\mathbb{E}_\omega[\mu(\omega)(s-f)P_i(\omega)\mathcal{E}_i^q] = (s-f)\mathcal{E}_i^q \mathbb{E}_\omega[\mu(\omega)]\bar{P}_i$ (no cov), simplifying to averages.

What remains on the LHS is the idiosyncratic reclassification-insurance term

$$(1-s)\sum_i \text{Cov}_\omega(u_i', q_i^* \varepsilon_{i\omega}),$$

and the RHS reduces to the three public wedges with averages:

$$-\mathbb{E}_\omega\left[\mu(\omega)\left(\sum_i \left((s-f)\, \bar{P}_i\right) \mathcal{E}_i^q + \sum_i \left(s\, q_i^* + f(1-q_i^*)\right) \mathbb{E}_\omega\Big[\frac{\partial P_i}{\partial m}\Big]\, \mathcal{E}_i^m + \sum_i \left(s\, q_i^* + f(1-q_i^*)\right) \mathbb{E}_\omega\Big[\frac{\partial P_i}{\partial x}\Big]\, \mathcal{E}_i^x\right)\right].$$

To simplify the the behavioral term, we formally state our assumptions of the m and x margins only mattering smoothly among the insured:

A6 (*Insured-only behavioral margins*) For $q_i^*(s) = 0$, $\dfrac{\partial m_i^*}{\partial \ln \mathbb{E}_\omega[\tilde{P}_i]} = 0$ and $\dfrac{\partial x_i^*}{\partial \ln \mathbb{E}_\omega[\tilde{P}_i]} = 0$.

A7 (*No jump at switching / smooth pasting*) For marginal switchers at s, $m_i^1(s) = m_i^0(s)$ and $x_i^1(s) = x_i^0(s)$, so there is no discrete level change in m, x when q flips.[15]

Hence, by A6, $\mathcal{E}_i^m = \mathcal{E}_i^x = 0$ whenever $q_i^*(s) = 0$, so the mitigation/location pieces reduce to $s\, q_i^*(s) \times (\cdots)$ and can be re-indexed to $\mathcal{J}_1(s)$ with conditional semi-elasticities $\mathcal{E}_i^{m|1}, \mathcal{E}_i^{x|1}$. By A7, there is no additional "switching" term from discrete jumps in m, x at the coverage boundary. The coverage term remains $(s-f)\bar{P}_i\, \mathcal{E}_i^q$ (marginal entry/exit).

Thus, from $0 = \frac{d\mathcal{L}}{ds} = D + B$ and (as in the general proof) $B = \frac{1}{1-s} \times$ [RHS bracket], multiply by $(1-s)$ to obtain the stated condition:

$$(1-s)D \;=\; -\mathbb{E}_\omega\Big[\mu(\omega)\Big(\sum_i (s-f)\bar{P}_i\, \mathcal{E}_i^q \;+\; s \sum_{i \in \mathcal{J}_1(s)} \big(\mathbb{E}_\omega[\tfrac{\partial P_i}{\partial m}]\, \mathcal{E}_i^{m|1} + \mathbb{E}_\omega[\tfrac{\partial P_i}{\partial x}]\, \mathcal{E}_i^{x|1}\big)\Big)\Big].$$

Proof of Two-State Model

Proposition 3 (Two-state identity for the reclassification term)**.** *Under the example assumptions—fixed behavior at (x, m, q^*), constant aggregate price and losses across states—the covariance simplifies to:*

$$\mathrm{Cov}_\omega(u', q^*\varepsilon_\omega) = q^* \left(\frac{u'^B - u'^G}{\Delta\varepsilon}\right) \mathit{Var}_\omega(\varepsilon_\omega).$$

Proof. With fixed L and $\bar{P}$, consumption differs only via ε_ω: $c^B - c^G = -(1-s)q^*\Delta\varepsilon < 0$ (assuming $\Delta\varepsilon > 0$), so $u'^B > u'^G$. The covariance is $q^*\mathrm{Cov}_\omega(u', \varepsilon_\omega) = q^*\mathbb{E}_\omega[u'\varepsilon_\omega]$ (since $\mathbb{E}_\omega[\varepsilon_\omega] = 0$):

$$\mathbb{E}_\omega[u'\varepsilon_\omega] = (1-\pi)u'^B(\pi\Delta\varepsilon) + \pi u'^G(-(1-\pi)\Delta\varepsilon) = \pi(1-\pi)\Delta\varepsilon(u'^B - u'^G).$$

And $\mathrm{Var}_\omega(\varepsilon_\omega) = \pi(1-\pi)(\Delta\varepsilon)^2$, so:

$$\mathrm{Cov}_\omega(u', q^*\varepsilon_\omega) = q^* \cdot \frac{\pi(1-\pi)\Delta\varepsilon(u'^B - u'^G)}{\Delta\varepsilon} \cdot \Delta\varepsilon = q^* \left(\frac{u'^B - u'^G}{\Delta\varepsilon}\right) \mathrm{Var}_\omega(\varepsilon_\omega).$$

□

[15] Equivalently, the "switching term" $(z_i^1 - z_i^0)\, \partial q_i^*/\partial \ln \mathbb{E}_\omega[\tilde{P}_i]$ for $z \in \{m, x\}$ is zero.