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

Jonathan Gruber
MIT & NBER

Adam Solomon
NYU Stern

February 16, 2026

Abstract

Intensifying climate change makes protection against natural disaster risk – either through ex ante insurance or ex post aid – a first-order public policy issue. Flood risk protection in the U.S. traditionally featured ex post aid through FEMA and heavily subsidized insurance through the National Flood Insurance Program (NFIP). In an effort to correct subsidy-induced overbuilding and under-mitigation in flood-prone areas, the NFIP recently moved to actuarially fair insurance premiums. But this change has two unintended consequences: a fiscal spillover onto FEMA disaster aid as insurance coverage declines in flood-prone areas, and household exposure to uninsurable reclassification risk from uncertain climate projections. We develop a model of optimal flood insurance that incorporates the static and dynamic tradeoffs of subsidizing premiums. We then use a variety of identification strategies and data to estimate the five key parameters needed to implement this model. We find that the fiscal spillover and reclassification risk protection benefits significantly outweigh the moral hazard costs, and that the optimal subsidy to flood insurance is 46%, comparable to the level of subsidization before the recent reform.

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, Rebecca Diamond, Cameron Ellis, Nathan Hendren, Larry Katz, Philip Mulder, Jim Poterba, Anna Russo & Pari Sastry for helpful comments, and to Katherine Wagner for generously sharing data.

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 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)), and over 64% of claims were disaster related (LexisNexis Risk Solutions (2025)).

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; CBO, 2019). 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 (FEMA, 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 when 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 cover themselves incompletely ex ante with insurance. Efforts to resolve this incomplete take-up 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 combating 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 insufficiently remediate – in flood-prone areas.

[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)

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 changing 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 static risk only, we also incorporate the dynamic reclassification risk that arises from future climate uncertainty. 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. Importantly, throughout our analysis we hold the generosity of ex post assistance (e.g. FEMA payouts) constant.

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 per house in the county, 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. Additionally, we estimate a negligible response along the intensive margin of coverage choice.

Third, we measure the increased reclassification risk from risk-based pricing. As we document, there is substantial uncertainty in both the mean path of climate change and 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 use the variation across models as a proxy for the cross-sectional climate uncertainty a household is exposed to.

We find that there is substantial uncertainty in future projections. Across the eight models, the

average household's future premium has a standard deviation of $3,254. This premium risk has an aggregate component (all premiums are expected to rise) and an idiosyncratic one (some properties' expected risk rises far more than others'). We compute households' WTP for insurance against idiosyncratic reclassification risk. We find that a 50% premium subsidy provides approximately $64 per house in insurance value.

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 0.75% 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. Although properties within the high-risk flood zone were already subject to elevation requirements, RR2.0 introduced premium-reduction incentives for elevation beyond the minimum. Properties outside the high-risk zone previously had no requirement to elevate, and RR2.0 introduces clear incentives to do so. Both inside and outside the high-risk flood zones we find a strong increase in mitigation as a result of price increases. Properties exposed to a higher premium jump were 10% more likely to elevate than those exposed to a lower premium jump.

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, 46%, is similar to the pre-RR2.0 level of subsidization. Even more starkly, we estimate that a premium subsidy up to 5% has infinite marginal value of public funds (MVPF); it pays for itself. We find that redistributive impacts are not significant and that allowing for heterogeneity in our estimated costs and benefits only strengthens our conclusion. We also consider alternative policy counterfactuals motivated by existing variations in regulation - such

as mandated elevation or insurance coverage - as well as incorporating behavioral extensions to our model and administrative cost differentials between FEMA and the NFIP. Our findings are generally robust and even under a set of assumptions most adversarial to subsidies, the lower-bound subsidy rate is over 20%.

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: mitigation and location. 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 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 25% in certain coastal locations. The NFIP issues a single type of policy: the Standard Flood Insurance Policy (SFIP). The SFIP pays the replacement value of damage 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.

The low coverage is particularly surprising since historically the NFIP has offered heavy subsidies (Wagner (2022)). Indeed, we calculate that, prior to RR2.0, the average policy received approximately a 50% subsidy. In addition to this ’level’ subsidy, there was a ’twist’ toward community rating: low-risk policies paid relatively more, and high-risk policies paid relatively less.

RR2.0, which was rolled out from 2021 through 2023, undid both of these price distortions by setting premiums according to each house’s actuarially fair risk. The aim was to provide accurate risk signals so that households would make efficient mitigation and location decisions. However, RR2.0 has also been criticized[2] as likely to cause a drop in insurance coverage, raising reliance on FEMA should a disaster occur.

This concern is especially relevant given the existing mandate structure. Outside of high-risk Special Flood Hazard Areas (SFHAs) - high-risk flood zones defined by FEMA as areas with a 1% or higher annual chance of flooding - there is neither a mandate that households hold flood insurance

[2]”[R]ising costs encourage policy lapses — shifting risk back to taxpayers when disasters strike” Wicker and Hyde-Smith (2025)

nor that construction meets minimum flood-proof standards, despite these areas accounting for one third of NFIP claims (Federal Emergency Management Agency (FEMA) (2025)). Inside SFHAs, which account for approximately 6-7% of all U.S. homes (Insurance Journal (2023)), there is a mandate that federally insured mortgages carry flood insurance. This mandate, in addition to affecting a very small portion of households (and only about a quarter of risk (risQ, Inc. (2021))), is imperfectly enforced. Compliance is greater than 90% at origination, but this drops off over the life of the mortgage (U.S. Government Accountability Office (2021b)). As a result, overall compliance rates are estimated at between 50% and 80% (2M Research (2020)).

In contrast, within SFHAs, there is a distinct mandate to elevate newly constructed buildings above the 1-in-100 year-flood level that has 89% compliance (Mathis and Nicholson (2006)). The elevation mandate only needs to be enforced at the moment of construction approval; the insurance mandate needs to be enforced at every annual policy renewal. These twin mandates, and their differential enforcement, are key to our findings: raising insurance prices dramatically reduces insurance take-up without a substantial reduction in mitigation.

Finally, historically, there has been no 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 (in flood-dollar-weighted terms) to the period we study (2010 to 2022). 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 four 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[3] of typical rebuilding costs. Second, there are Small Business Administration (SBA) disaster home loans, which let owner-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. Fourth, mortgage default offers partial implicit insurance to households while imposing a cost to the government through Fannie Mae or Freddie Mac.

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

Because federal law bars duplication of benefits, each of these programs is explicitly offset by any insurance proceeds. The primary interaction is between NFIP and FEMA's IHP. In practical terms, a homeowner with an NFIP policy may see their IHP grant reduced to zero. Nevertheless there is evidence[4] of non-trivial double-payment, sometimes nefariously but often related to the lag with which insurance payouts are assessed and paid.

Relatedly, 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)). This incomplete overlap – IHP benefits sometimes duplicate NFIP payments; IHP benefits often cover only a portion of what an NFIP contract would – means that the relationship between increased insurance payouts and reduced ex-post assistance is not mechanical, and motivates our regression analysis in Section 4.

In addition to the IHP interaction, NFIP coverage should reduce reliance on other forms of ex-post assistance. They will borrow less—or nothing—from SBA. They will default less on their mortgages, many of which are federally backed. The HMGP formula explicitly allocates more money to counties with higher IHP outlays; as the latter are reduced so too will be the former.

To the extent that owners view ex post aid as reliable, it can crowd out ex ante insurance purchase, a point discussed in the literature review section below. Throughout our analysis, we hold ex post assistance fixed at its current levels and ask whether that rationalizes an ex ante subsidy. That is, 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.

A distinct question that we do not speak to is whether ex post aid should be more or less generous than it currently is. Institutionally, ex post aid has persistently been a feature of disaster assistance, largely because its salience creates strong political incentives (Healy and Malhotra (2009)). In the absence of complete ex ante insurance, some ex post assistance can be optimal in our setting due to reclassification risk. Because households that also get reclassified as high risk are, by definition, more likely to flood, ex post flood assistance indirectly provides insurance against reclassification risk. However, the variance in future premiums is substantially smaller than the variance in flood realizations, such that ex post aid, which targets the latter, turns out to insure less than 1% of the former. Conversely, reclassification risk provides a direct rationale for NFIP subsidies even if there were no ex post assistance provided by the government. An implication of Section 6 is that reclassification risk directly justifies an NFIP premium subsidy of approximately 10%, even absent any fiscal spillovers.

[4]For example, following Hurricane Sandy, a FEMA review found many households who initially received IHP assistance later received NFIP payouts. HUD explicitly announced it would not require repayment of earlier assistance (U.S. Department of Housing and Urban Development (2015)).

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); Xu and Box-Couillard (2024)). 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 ex post 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 coverage. 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. Their headline estimate — an elasticity of insurance expenditure with respect to expected ex post aid of 0.2 — would imply that FEMA, which compensates households for approximately 10% of ex post damages, depresses insurance take-up by only 2%. 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 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 (e.g., 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 present a model of optimal insurance subsidies in the presence of moral hazard, dynamic reclassification risk, and the crowding 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

A household i faces uncertainty about their future climate risk, and about the realization of specific natural disasters. We label the future climate state by ω. Given a future climate state, the household faces an uncertain flood loss $L_i(\omega)$, the distribution of which depends on ω.

Let $P_i(\omega)$ be the baseline, actuarially fair premium that would prevail in that climate state. 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 $m_i \in [0, \bar{m})$. For expositional clarity, we model all three as continuous. In Section 9, 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 $\Gamma_i(x_i)$. We denote the elasticity of the household's choice of coverage q with respect to the price of insurance by η_i^q, and similarly for η_i^m and η_i^x.

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 fixed portion $f > 0$ is covered by the government ex post. Finally, the government budget is balanced by some (climate state-specific) lump-sum tax $T(\omega)$. Hence, household consumption, after climate state ω and idiosyncratic flood shock ξ are realized, is given by

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

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, where V_i is their 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 in the cross section.

This model, although motivated by climate risk, applies to many social insurance settings where risk dynamically evolves over time. This includes health insurance with medical risk increasing as one ages (Handel et al. (2015)), workers' compensation insurance where claims today impact risk-rated premiums tomorrow, and unemployment insurance as job-loss risk due to artificial intelligence increases.

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 household-specific fixed effect, and $\varepsilon_{i\omega}$ is the residual with $\mathbb{E}_\omega[\varepsilon_{i\omega}] = 0$. Moreover, we define a household's expected future premium: $\bar{P}_i = E_\omega [P_i(\omega)]$.

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. In the empirical section, we show that in fact there is no clearly progressive or regressive redistributive effect of insurance subsidies.

The Planner's Problem

The planner chooses s to maximize

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

There are two important observations. First, we have assumed budget balance within each climate state ω. Since ω represents a possible path that the global climate takes – that is, a real-

ization of aggregate climate risk – there is no way to move resources between those states of the world. That is, our budget balance formulation assumes that aggregate risk is not insurable. Our estimates of the insurance benefits of a subsidy therefore only account for the idiosyncratic ($\varepsilon_{i\omega}$) portion of premium uncertainty, not the climate-state aggregate ($\bar{P}(\omega)$).

Second, 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 programs are funded by the same people that benefit; we do not analyze the distributional consequences of the average taxpayer possibly differing from the average disaster aid recipient. Since lump-sum taxes are not distortionary, our optimality condition is equivalent to finding the premium subsidy for which the MVPF is equal to 1. If subsidies were financed by alternative, distortionary taxes with an MVPF greater than 1, the optimal subsidy would be lower.

Formula for the Optimal Subsidy

The general formula for the optimal subsidy is more detailed than is relevant to this empirical setting. We derive the full formula in Appendix C.1. Here, we make a number of additional assumptions to match our empirical setting. However, there are a wide variety of settings in which this dynamic model is relevant and in such settings estimating the additional terms in Appendix C.1 might be appropriate. Specifically, we assume:

A1 $\mu(\omega) = \bar{\mu}$ for all ω; dynamic-flood risk is small relative to the economy.

A2 For each ω, $\text{Cov}_i\left(q_i^*, u_i'(\omega)\right) = 0$; no selection into insurance.

A3 $\text{Cov}_i(a_i, q_i \mathbb{E}_\omega[u_i'(\omega) - \bar{u}'(\omega)]) = 0$; 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 there is no selection into insurance based on the individual valuation of insurance. In the NFIP, premiums are risk-rated, so selection (on marginal utility) would have to arise from private information on flood risk or on risk preferences. Wagner (2022) does not find evidence for selection on risk; she studies a reform that sharply increased prices and reduced demand among some houses but found no change in flood costs.[5] Moreover, selection on risk preferences would only *strengthen* the case for subsidies, as a subsidy would incentivize insurance take-up among those with particularly high marginal utility, and therefore be even more desirable for the planner than if take-up was uncorrelated with marginal utility.

[5] In a different context, Xu and Box-Couillard (2024) also finds no evidence of adverse selection in post-disaster insurance take-up.

Assumption A3 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 A4-A5)[6], we arrive at the simplified optimality condition that we will estimate.

Proposition 1 (Optimal Subsidy:)**.** *Under Assumptions A1-A5, the optimal premium subsidy* s *solves:*

$$\underbrace{(1-s)\sum_i \text{Cov}_\omega(u_i' - \bar{u}'(\omega),\ q_i^* \varepsilon_{i\omega})}_{\textit{idiosyncratic insurance value}} + \bar{\mu}\ \underbrace{\mathbb{E}_\omega\left[\sum_i f\,\bar{P}_i\,\eta_i^q\right]}_{\textit{FEMA savings}} =$$

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

The left-hand side measures the benefits of an expanded subsidy.[7] Specifically, the benefit is the value of insurance against idiosyncratic (reclassification) risk that a subsidy provides. It is the continuous analogue of the difference in marginal utilities seen in the typical binary-outcome model, as we show in Appendix C.3. A subsidy redistributes *within state* ω, from those who received a favorable climate-risk shock (i.e. $\varepsilon_{i\omega} < 0$, leading to a low premium) to those who received an unfavorable shock ($\varepsilon_{i\omega} > 0$). This generates insurance value to the extent that the latter have high marginal utility *relative to the climate-state average* ($u_i' - \bar{u}'(\omega) > 0$). Since we assume budget balance within each state ω, there is no redistribution between future climate states. The second term on the left-hand side is the FEMA-cost saving from increased insurance. Although it is a fiscal saving, not a risk-smoothing gain, it is the primary benefit of expanded subsidies.

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. 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

[6] See the proof of Proposition 1 for formal definitions.

[7] Almost all of the costs and benefits — flood damage, NFIP premium, premium subsidy, savings from mitigation and location — accrue to welfare immediately and annually. They are measured in terms of full-risk premiums denominated in today's dollars. Hence, these costs and benefits are contemporaneous and measured in dollars per year. The exception is reclassification risk, in which we are measuring household WTP for a reduction in future premium uncertainty. For reasons we describe in Section 6, we do not discount those future benefits since they scale with home values and therefore premiums. Moreover, in Appendix B.2, we demonstrate our results are not qualitatively affected by this discounting assumption.

not the cost of the subsidy itself. Importantly, because a premium subsidy affects location and mitigation choices only among those that remain (fully) insured, we assume that there is no impact of mitigation or location on the level of FEMA spending.

In sum, we have derived a formula for the optimal subsidy with dynamic reclassification risk and multiple fiscal spillovers that nests the standard insurance-moral hazard tradeoff. To solve the formula in Proposition 1, 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 change. In the next five sections, we estimate each of these quantities.

3.1 Misperception of Flood Risk

There is evidence that individuals misperceive their flood risk (see, e.g., Mulder (2024) and Bakkensen and Barrage (2022)). This could occur in two ways: households could misperceive their level of flood risk, and they could misperceive the responsiveness of their risk to mitigation. The evidence suggests that households *underperceive* both.

In Appendix C.2 we extend our model to incorporate misperceptions in the level of risk. When the individual chooses coverage, mitigation and location according to misperceived flood risk, yet the planner evaluates welfare according to the correct (i.e. experienced) risk level, this generates an 'internality'.

To understand how these internalities change optimal policy, consider the households who underperceive their risk and who are marginal to a slightly increased subsidy. Those induced into insurance have (mis)perceived benefit equal to cost, yet true household benefit greater than cost. Thus, correcting the coverage internality through a higher subsidy has first-order welfare benefits.[8] We analyze the implications of these risk misperceptions for optimal policy in Section 9.

In contrast, households might underperceive the returns to mitigation. Ideally, the NFIP would directly incorporate this into a clear premium schedule in which certain improvements transparently reduce premiums. However, if households still misunderstand, and a premium subsidy further depresses the mitigation salience, this would create a negative internality. To guard against this, in Section 9, we test the robustness of our policy prescriptions to a greater mitigation elasticity, which proxies for welfare if the slope-internality were corrected.

Finally, we reemphasize that our analysis takes the generosity of the ex post safety net as fixed. There is evidence that households mistakenly believe that FEMA is more generous than it actually is (Kousky and Shabman (2012)). This would make insurance subsidies even more attractive to

[8]Technically, this is only true if coverage is complete - with incomplete coverage, underperception of losses would mean households also understate the benefits of risk reduction through mitigation and location choice. However, as we show in Appendix B.9, coverage (among the insured, on the intensive margin) is effectively complete in this market. As a result, there is no internality from mitigation or location; insured households already ignore their level of risk, and so misperceiving that risk is inconsequential.

a planner: if the marginal household is indifferent between the NFIP and an erroneously inflated perception of FEMA, they would experience a first-order utility gain when induced into insurance.

4 Spillovers to Ex Post Disaster Aid

Data on presidentially declared (flood) disasters (PDDs) come 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 IHP was authorized.[9]

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 IHP, NFIP policies and claims, and the HMGP come from the OpenFEMA data portal. Similarly, disaster loan information for the 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 from 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.[10] About one in five households

[9] 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.

[10] FEMA assistance only activates for disasters large enough to warrant a Presidential Disaster Declaration; the NFIP covers all floods, even with much less total damage. Our sample, which restricts to PDDs, therefore underestimates

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 $3,440. 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. Ex post assistance from FEMA totals $733 per house, 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 (Lindsay and Getter (2023); U.S. Government Accountability Office (2021a)), 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 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:

$$\text{ExPost Spending}_{c,f,t} = \alpha_c + \gamma_t + \beta \cdot \text{NFIP Coverage}_{c,t-2m} + g\left(\text{Flood Damage}_{c,f}\right) + \epsilon \quad (1)$$

where $\text{ExPost Spending}_{c,f,t}$ 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, g is the function we use to control flexibly for flood damage, α_c are county fixed effects, and γ_t are calendar-year fixed effects. We control for flood damage by binning[11] and estimating a distinct quadratic function within each bin.[12] 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,

the benefits of an NFIP subsidy. Conditional on a PDD, an NFIP subsidy saves FEMA money, and even if there is no PDD the household is protected by the more comprehensive NFIP policy.

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

[12] In Table A7, we show the robustness of our results to alternative functional forms for the flood damage control; in general, the alternatives all yield estimates larger than the ones we obtain below.

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.

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.[13]

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 a $203 reduction, which is statistically significant at the 1% level. As we decompose in the Appendix Tables A2 - A6, the majority of

[13]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.

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 $145 per household as coverage increases by 1%, almost identical to our estimate.

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) 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 own-county c, so that we are only identifying our effect off the timing of neighbors' floods, not general trends in flooding in the area. 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.

Moreover, one might be concerned that repeat flooding in an area is treated less generously by FEMA than a first-time flood. This might explain the relatively lower FEMA spending in repeat floods when insurance coverage is higher. To rule this out, in Appendix B.1, we analyze all *rejected* household applications to FEMA and show that less than 1% are rejected for any reason related to a previous flood, previous insurance mandate, or in general anything dynamic.

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 (7) 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 an 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 -0.15 to -0.17 (Mulder (2024)), -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. Moreover, we find lower elasticities (approximately -0.25) in the 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.

Instead of dropping coverage (i.e. the extensive margin), there is the possibility that households respond to higher prices by reducing their coverage (i.e. the intensive margin). We analyze this in Appendix B.5 and find a statistically precise but economically tiny intensive margin semi-elasticity: coverage declines by approximately $3,000 against a baseline of $300,000. This small effect, no larger than 1%, is likely because NFIP policies have relatively low coverage limits ($250,000 for building coverage, $100,000 for contents) that are binding for most households. As such, we proceed assuming that the level of coverage does not change, conditional on having insurance.

In Appendix B.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). We proceed with the within-house estimates as primary as it allows for heterogeneity by house characteristics but, as we show in Section 9, our policy

conclusions are relatively insensitive to assuming different elasticities of demand.

6 Reclassification Risk

The NFIP's transition to RR2.0 ties premiums directly to risk, and those premiums will respond to any future changes in risk. 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 variability a household faces about its premium relative to the average, driven by uncertainty 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 Sources of Climate Model Uncertainty

We proxy for uncertainty across climate paths with variability between climate models. Climate projections differ across models for epistemic reasons that go beyond computational approximation error. A core difficulty is that global climate models contain parameterizations for many unresolved processes (e.g., clouds, convection) which must be tuned to match a set of historical climate averages. This calibration is an ill-posed inverse problem with multiple distinct solutions: there are far more 'parameters' than historical climate data to fit against, so multiple distinct parameterizations can fit broad historical targets comparably well (Schneider et al., 2024; Hourdin et al., 2017). Crucially, this is uncertainty faced by the climate-modeling community itself: because the historical record provides too few statistically independent constraints to identify key parameterizations, even the scientists building these models cannot uniquely determine which model structure/tuning is correct, so multiple Coupled Model Intercomparison Project (CMIP) climate models remain simultaneously plausible as possible future climate paths given the data.

Our analysis uses the synthetic flood hazard projections of Gori et al. (2022), which apply multiple CMIP models under a common aggregate warming pathway (SSP5-8.5). These data provide one projection per CMIP model at the resolution needed for our application, so our baseline measure captures between-model uncertainty rather than within-model initial-condition sensitivity. Importantly, even if within-model variation were available, between-model uncertainty would remain central for our analysis. Meiler et al. (2025) show, at the level of global climate outcomes, that between-model ('epistemic') uncertainty is substantially larger than within-model ('aleatory')

uncertainty. We therefore interpret cross-model dispersion in Gori et al. (2022) as economically relevant uncertainty about localized future hazards that translates into idiosyncratic premium reclassification risk under RR2.0, although we recognize that it is not the full measure of uncertainty.

6.2 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 are 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 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: SSP5-8.5.[14] This produces, for each model and for each location, a variable number of flood events. 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.

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. This assumes that the relationship between risk factors and premiums is stable over time. 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 millimeter 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.

6.3 Simulating Future Premiums

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.

[14] 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 SSP5-8.5 is about 4.4 degrees Celsius relative to 1850-1900 (IPCC (2021)).

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 B.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 - 40 to 50% or so - the marginal impact from a bit more insurance against reclassification risk has decayed to be very small, regardless of additional discounting.

Gori et al. (2022) simulate 2070 rain, surge, and joint-hazard measures under eight general circulation models (GCMs) from the CMIP. For each policy and each GCM, we replace historical hazard values with the projected 2070 values for that GCM, using inverse distance weighting 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 $\$6,092$ 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 $\$3,254$ 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 $\$1{,}426$. 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 sample of policies under two of the least-correlated GCMs (IPSL on the horizontal axis; GFDL on the vertical). The dashed line represents pure aggregate risk (uniform scaling of premiums). If there was no disagreement across models in the relative riskiness of houses, we would see clustering along the dashed line.

In contrast, deviations from this line reflect diversifiable idiosyncratic reclassification risk: in one scenario house A is high risk and house B is low, in the other scenario, the inverse. These across-model relative risk disagreements are insurable across houses. Visually, we see there is substantial diversifiable risk. Moreover, the dots are colored in relation to the model disagreement among all models (i.e., the two models (IPSL and GFDL) included on the axes of Figure 3 as well as the six not included). Many of the households with relative agreement among the two graphed models (near to the diagonal line) have substantial disagreement among the remaining models (green and yellow coloring). We next turn to quantifying the benefits of insuring that diversifiable risk.

Figure 3 about here

6.4 Reclassification Risk Results

We measure the reduction in insurance coverage against reclassification risk induced by RR2.0. Our estimates imply that RR2.0 removed a level subsidy of approximately $s = 0.5$.[15] 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 compute the change in the mean premium

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

and define

$$\text{Pure Insurance Value of Undoing RR2.0} = \mathsf{CE}_i\big(\widetilde{P}_i(0) - \tau_i\big) - \mathsf{CE}_i\big(\widetilde{P}_i(0.5)\big)\,.$$

That is, the insurance value of moving from zero subsidy to a 50% subsidy is the difference in certainty equivalents of the induced premium distributions, removing the level difference in average premiums. We assume CARA utility, $u(c) = -\exp(-\gamma c)$, and report results for multiple γ.[16]

Table 5 about here

Table 5 indicates that RR2.0 exposed households to substantial reclassification risk.[17] Under our central risk-aversion parameter $\gamma = 5 \times 10^{-4}$, the level shift reduces insurance value by \$64.5,

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

[16] 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.

[17] Our reclassification risk estimates hold other behavioral responses constant (and vice versa: for example, our

about 7% of the subsidized pre-RR2.0 premium, 3.6% of the RR2.0 premium and 1% of the projected future premium.

These results imply that reclassification risk justifies a flood-insurance subsidy even absent any fiscal spillover to ex post assistance. Specifically, we calculate that a subsidy of approximately 10% is optimal solely to combat reclassification risk.

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.

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 U.S., 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 $\$1,772$.

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 the price changes each flood zone experienced due to the RR2.0 repricing.

coverage estimates hold location and mitigation constant). This is formally corrected in a Baily-Chetty framework in which the effect that subsidies have on location which then in turn impacts reclassification risk exposure is, by definition, second-order. Empirically, the very small location elasticities we find in Section 8 justify treating these second-order effects as negligible.

7.2 Policy Change

Risk Rating 2.0 affected incentives to elevate new buildings distinctly in high (i.e., SFHAs) and low flood risk areas. 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. Any additional price-induced mitigation response will be above the code-mandated elevation standard.[18]

In contrast, outside of the SFHA there is no such elevation mandate. Areas outside of the 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. 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. Moreover, there is no BFE measured since the BFE measures the depth of a 1-in-100-year flood.

7.3 Analysis

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

Inside the SFHA, we compare very high risk properties to high-risk properties. 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. Correspondingly, our standard for elevation inside the SFHA is whether a house elevates at least one foot above the mandated BFE minimum.

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. As such, we measure voluntary elevation with an indicator for whether elevation is three feet above the property's highest adjacent ground level.

[18] An audit by Mathis and Nicholson (2006) finds 89% compliance with the SFHA elevation requirements. Both the insurance mandate (for federally insured mortgages) and the elevation mandate have similar compliance rates of about 90% at origination/construction. However, while the latter is a one-off decision, the former mandate needs to be enforced every year at insurance renewal.

These measures – one foot above BFE inside the SFHA and three feet above adjacent ground level outside the SFHA – are measured relative to different baselines, but amount to similar overall levels of risk reduction. In both cases, we measure adaptation relative to the mandated minimum elevation (or lack thereof). And, because depth-damage functions that we use to measure the risk reduction of elevation are in absolute terms, the fact that BFE is already a few feet above ground level makes the absolute standard comparable. We show robustness to these particular thresholds in Appendix B.6. In particular, these thresholds generate plausible upper bounds on elevation treatment effects, and are therefore as adversarial as possible to our eventual finding of optimally large positive premium subsidies.

We estimate the probability of a house i in NFIP community j and year t elevating by 1 foot or more relative to the base flood elevation (inside the SFHA) or 3 feet or more above adjacent 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:

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

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 pre-RR2.0 subsidies comparably inside and outside the SFHAs. In both cases, RR2.0 led to elevation being priced on the margin (in the case of the SFHA, over and above the required minimum). As such, households exposed to a relatively higher price increase increased their elevation probability relative to those exposed to a lower price increase but with the same SFHA regulations (or lack thereof).

The probability of elevation increased by approximately 10% in both cases. Inside the SFHA, the difference in price increases was \$2,142 - \$1,202 = \$940, against a mean post-reform price of \$3,383. This yields an elasticity of 0.36. Outside the SFHA, the difference in price increases was \$372, against a mean full-risk price of \$1,518, implying an elasticity of 0.25.

This elasticity, in combination with the mitigation benefits of 3 feet of elevation – 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 \$1,063, to \$1,992, 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

$$\log(\text{OccupiedResidential}_{ct,q}) = \alpha_{ct} + \gamma_q + \sum_{q' \neq \text{Q1 2022}} \beta_{q'} \Big(\text{MedianIncrease}_{ct} \times \mathbf{1}\{q = q'\} \Big) + \varepsilon_{ct,q}, \quad (4)$$

where α_{ct} are census–tract 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.[19] 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, which is consistent with small estimated locational responses to actual disasters. For reference, estimates of out-migration 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).

Similarly, these effects cohere with estimates of the impact of rental increases on out-migration. For example, Autor et al. (2014) find that a 27% increase in rents caused 5% of renters to leave their properties; an elasticity of 0.19. In our case, the differential insurance increase of $921, represents 3.8% of annual U.S. rent ($23,940 per Zillow Group, Inc. (2026)). As a result, our migration response of 0.75% yields an elasticity of 0.20, remarkably similar to Autor et al. (2014).

9 Optimal Policy

In this section, we calibrate a discrete approximation of the formula in Proposition 1. We use this calibration to estimate the optimal NFIP subsidy under various assumptions.

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 η^q and the average flood insurance premium (i.e. expected risk) amongst the marginals, $\bar{P}$, general. We derive optimal subsidy rates for a grid of values for (η, $\bar{P}$).

9.1 FEMA Savings Calibration

The FEMA savings term is equal to $\sum_i f P_i \eta_i^q$. 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 η_i^q. In our baseline, we assume an elasticity and FEMA spillover that doesn't vary across people.

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

[19] In our data the vacancy rate is approximately 3%, coherent with the national averages of approximately 1.5% and 7% for, respectively, properties on the market for sale and for rent (U.S. Census Bureau (2025)). This is more than sufficient for a 0.75% change in occupied housing even without any supply side response that we might expect in the longer term.

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 $\Delta q \simeq \eta^q \Delta \ln \tilde{P}$. Hence, the FEMA saving per house-year from a small subsidy s is $\Delta\text{FEMA}(s) = (\beta_{\text{cond}}) \frac{\bar{P}}{L_f} \cdot 100 \cdot \eta^q \cdot \Delta \ln \tilde{P}$.

For example, plugging in $\Delta \ln \tilde{P} = 0.01$, $L_f = 31{,}921$, the average full-risk premium $P = \$1{,}500$, and $|\eta^q| = 0.32$, 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, define $WTP(s) = CE(\tilde{P}(s+\Delta s)) - CE(\tilde{P}(s))$ 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 $\sum_i sP_i\eta_i^q$. On average, since we assume η_i^q does not vary across people, and scaled by the price change $\Delta \ln \tilde{P}$, this becomes:

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

At our central estimates of $\Delta \ln \tilde{P} = 0.01$, $\bar{P} = \$1{,}500$, and $\eta^q = 0.32$, 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 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 $s \cdot \mathbb{E}[\partial P_i/\partial m] \cdot \eta^m$, so a small policy change that induces a price change $\Delta \ln \tilde{P}$ contributes $s \cdot \mathbb{E}[\partial P_i/\partial m] \cdot \eta^m \cdot \Delta \ln \tilde{P}$ 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. This is exactly the outcome in our event studies outside of the SFHA, and an over-estimate inside the SFHA.[20] 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 calculated in Section 7 mitigation semi-elasticities of 0.36 (SFHA) and 0.25 (non-SFHA); our calculations use an average of 0.3.

In summary, 1% subsidy yields a mitigation fiscal externality of $s \times 0.4 \times 0.3 \times \bar{P} \times 0.01. = 0.0012 \times s \times \bar{P}$ At $\bar{P} = 1500$, this is equal to $\$0.90$ per affected house at $s = 0.5$.

9.5 Location Fiscal Externality Calibration

For the location term, we focus on the intensive-margin response among the insured. Identically to the prior subsection, we set $q = 1$ and $f = 0$. Under this restriction, the location component on the right-hand side reduces to $s \cdot \mathbb{E}[\partial P_i/\partial x] \cdot \eta^x$.

When an insured person moves, they 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 η^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, $\eta^x \approx (-0.0075)/(+0.874) \approx -0.0086$. This gives a location fiscal externality of

[20] The Base Flood Elevations, although hard to measure comprehensively, are almost always above 2 feet. Our SFHA event study shows a marginal foot of elevation. Thus, an assumption of 40% risk reduction inside the SFHAs is conservative (and adversarial to subsidies) in two ways: 1) the marginal change is only one foot, not three; 2) the returns to elevation are decreasing and inside the SFHA there is already some elevation occurring.

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

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 $s = 1$. This is an order of magnitude smaller than the fiscal externality from mitigation since the underlying elasticity (0.0086 vs 0.3) is correspondingly smaller.

9.6 Optimal Subsidy

The optimal subsidy is the s^* that solves the equation in Proposition 1. 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 η^q. The results are illustrated in Figure 6.

Figure 6 about here

At our central estimates – $\eta^q \approx 0.32$ and $\bar{P} \approx \$1,500$ – the optimal subsidy is approximately **46%**.

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 40% - 50% 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 $\bar{\eta^q}$. This is because both terms are similar functions of the number of marginals (hence η^q) and 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 policyholders 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 Marginal Value of Public Funds.

Because we assumed that premium subsidies were funded with lump-sum taxes, the optimal subsidy (46% in our baseline) implicitly solves $MVPF = 1$. But if premium subsidies were funded by distortionary taxation that had an MVPF greater than one, the optimal level of subsidy would fall. To quantify this, we compute the MVPF for each marginal percentage of subsidy. The general formula for the MVPF (see Hendren and Sprung-Keyser (2020)) is $MVPF = \frac{\text{Mechanical Benefit}+\text{Insurance Value}}{\text{Mechanical Cost}+\text{Fiscal Externality}}$. The results are in Figure 8.

Figure 8 about here

We find that the first 5% of subsidy have infinite MVPF: the fiscal externalities from reduced FEMA expenditure are so large so as to outweigh the mechanical costs and moral hazard. Subsidies from 5% to about 20% have positive, but non-infinite MVPFs above 3. As the subsidy continues increasing, the MVPF declines until it reaches 1 at 46%.

We can compare this to the MVPFs of other environmental policies analyzed by Hahn et al. (2024). They find MVPFs consistently above 4 for subsidies to wind and solar production, at about 1.5 for electric vehicle subsidies, and commonly at 1 or below for appliance rebates and hybrid vehicle tax credits. As a result, absent redistributive considerations, an expanded premium subsidy funded by a reduced hybrid tax credit would be efficiency increasing.

9.6.3 SFHA Regulations - Mandated Mitigation and Insurance Purchase

An SFHA designation's main effect is to mandate insurance coverage (albeit with imperfect enforcement). Per Table 3, this leads to a lower semi-elasticity of demand: 0.25 instead of 0.32. When demand is less elastic because of an imperfectly enforced SFHA-style insurance mandate, the optimal subsidy falls from 46% to 42%. This is primarily because the main benefit of a subsidy – reduced FEMA expenditure through increased insurance demand – is dulled. The effect is minimal for the same reasons as above: both the primary benefit and the primary cost of an expanded subsidy are linear in η^q. Moreover, we have assumed that the returns to marginal mitigation inside the SFHA are the same as outside. This is likely an overestimate: since elevation up to the BFE is mandated inside the SFHA, an additional foot of elevation likely has lower benefit than the first foot of elevation outside the SFHA. A full accounting would push in favor of a higher subsidy.

9.6.4 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 B.7, 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 to approximately 35%.

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%. As noted, this increases both FEMA benefits and subsidy costs, and therefore has a minimal impact on the optimal subsidy, which increases to 50%.

9.6.5 Administrative Costs

Both ex post FEMA assistance and ex ante NFIP insurance have administrative costs. NFIP pays approximately 30% of its premium income as administrative expenses, primarily broker commissions (U.S. Government Accountability Office (2023)). FEMA administrative costs are approximately 13% of assistance delivered (U.S. Government Accountability Office (2014)). As a result, an increased premium subsidy pushes households from a low overhead program (FEMA) to a higher overhead program (NFIP). This would make a subsidy relatively less attractive.

To model this we need to make an additional assumption about what proportion of the administrative fees are simply transfers versus actual resource use. Absent strong evidence on this, we assume a half and half breakdown for both FEMA and the NFIP. We view this as conservative (i.e. bias our optimal subsidy result downward), since the bulk of the NFIP overhead comes from broker commissions, which are a transfer.

With these differential administrative costs, the optimal subsidy falls to 38%.

9.6.6 '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 0.75%. 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 50% larger than in our baseline (instead of about 90 cents per house year at $s = 1/2$, we assume \$1.35). As we discussed in Section 3.1, this proxies, in particular, for the effects of correcting household underperception of the returns to mitigation.

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.

Under these very high estimates for the costs of a subsidy, and very low estimates for the benefits, the optimal subsidy is between 20 and 30%. For our baseline demand elasticity of 0.25 and premium of \$1,500, the optimal subsidy is 22%. 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.

9.6.7 Rolling back RR2.0 – Level vs Twist

Our primary counterfactual policies are ad-valorem subsidies. However, in addition to these subsidies that are equal (in proportion) across households, RR2.0 involved a relative 'twist' in the prices: some households received relatively larger subsidies than others. We study the extent to which twists in the premium schedule are welfare-enhancing, and how this interacts with the original level subsidies.

Table 3 shows that demand semi-elasticities are lower among those with a relatively big pre-RR2.0 subsidy. Similarly, in Appendix B.7, we show that the FEMA spillover benefit increases in the pre-RR2.0 subsidy. That is, a twist in the direction that undoes RR2.0 transfers from high demand elasticity households to lower demand elasticity households, and from lower FEMA-spillover households to higher. The latter unambiguously makes a twist a good idea; the former is less certain. A lower demand elasticity makes the coverage response to a subsidy lower, but also reduces the fiscal cost from the subsidy itself.

Since the lower fiscal cost from the subsidy scales with the level subsidy s, we find that the welfare impact of a twist is, as the level subsidy increases from zero, initially negative and subsequently positive. In particular, the optimal level subsidy with this heterogeneity is approximately 50%, but an additional 20% twist (calibrated from RR2.0) adds approximately 3% to welfare over

and above the level subsidy. Full details and results are in Appendix B.8.

This implies that RR2.0's negative effect is compound: it removes the level subsidy, and it untwists the premium schedule. Put differently, conditional on a level subsidy, a twist is also optimal; with no level subsidy, a twist is less consequential but marginally welfare decreasing.

9.6.8 Behavioral Misperceptions (Internalities).

As discussed in Section 3, there is substantial evidence that households underperceive their flood risk and that this contributes to lower demand for flood insurance.

We study how the optimal subsidy changes in the presence of these misperceptions. We assume that the household buys insurance according to downwardly biased perceived WTP, whereas the planner evaluates welfare according to the household's unbiased WTP as if they held correct beliefs about their risk.

This internality makes a premium subsidy more attractive. The marginal household induced into flood insurance has a first-order welfare gain (the difference between their welfare-relevant unbiased WTP and their decision-relevant downwardly-biased WTP) in addition to the costs and benefits modeled throughout. Following Bakkensen and Barrage (2022), we assume an underperception of risk of 43%. Full details of how we calculate biased and unbiased WTP for flood insurance are in Appendix B.10.

We find that incorporating internalities shifts the optimal subsidy from our baseline of 46% to 86% when the coefficient of absolute risk aversion is $\gamma = 5 \times 10^{-4}$, and to 100% when $\gamma = 5 \times 10^{-5}$. This reflects the fact that, with a misperception so large, there is a huge additional benefit to inducing the marginal household into insurance.

10 Conclusion

The ongoing increase in "global weirding"[21] has raised attention to the risks posed to households from natural disasters. One of the most salient and discussed issues 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 SFHAs, averaging only 23% (Amornsiripanitch et al. (2024)).

One response to this lack of coverage was extensive subsidies offered for the purchase of an NFIP policy – both in terms of a reduced average premium, and muted variation by risk level of the area. But this has long been criticized as promoting development – and preventing mitigation – in our most flood-prone areas. As a result, in 2021 the government introduced a new actuarially fair

[21] See Friedman (2010)

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 sizable, 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 46%, 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.[22] 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.

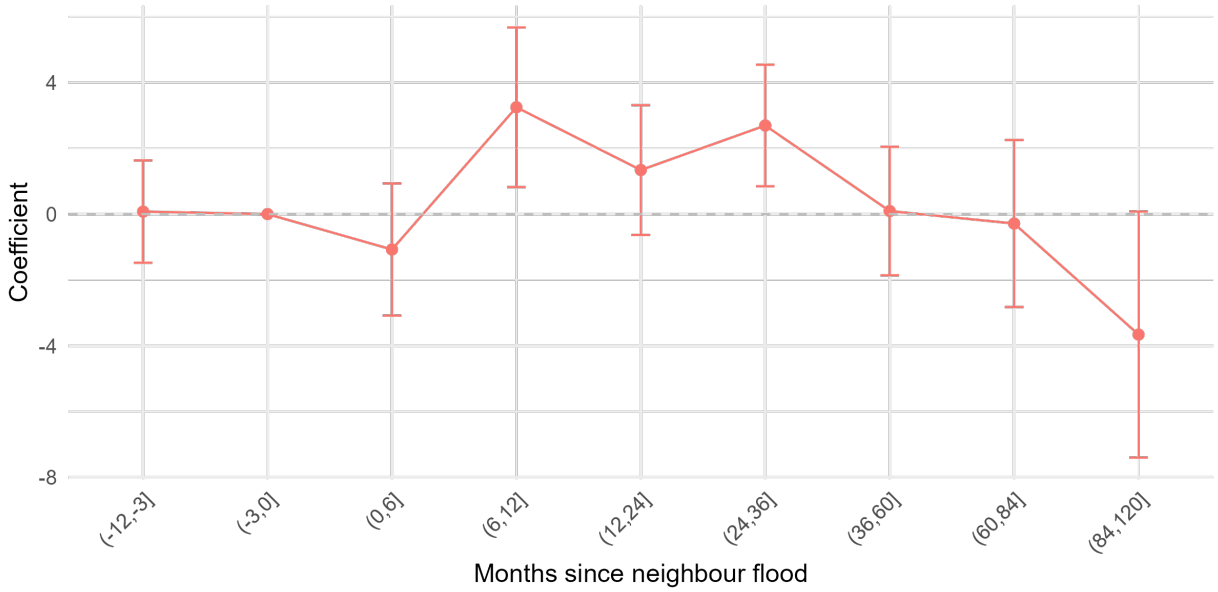
[22] See: Reuters (2025)

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

Variable	Weighted Mean	Weighted Median	90th Percentile
Panel A: Aggregate (levels)			
Property Damage (*millions $*)	4,700	3,200	10,000
Total Population	980,771	345,995	4,525,519
Total Housing Units	380,678	142,115	1,714,340
Panel B: Per house			
NFIP Policies per House	0.18	0.16	0.49
NFIP Payouts per House ($)	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 ($)	1.71	0.29	2.53
HMGPP Expenditures per House ($)	147.42	88.94	282.36
Approved SBA Loans per House ($)	703.57	301.30	1,404.00
Fiscal Cost from SBA Loans per House ($)	91.47	39.17	182.52

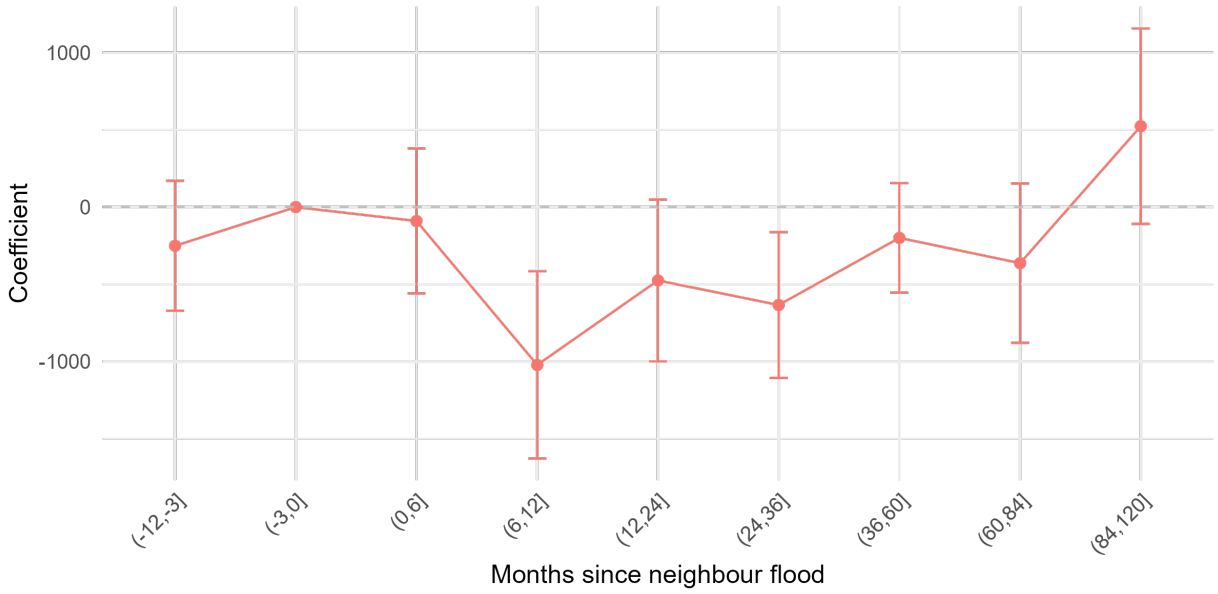
Notes: Sample statistics are from counties with presidentially declared disasters from 2010 to 2024 where assistance under the Individuals and Households Program (IHP) was authorized. All statistics are *damage-weighted* using county-level flood property damage from the NOAA Storm Events database. Panel 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. We assume that each dollar loaned out by the SBA costs the government 13c through interest rate subsidies and defaults (Lindsay and Getter (2023); U.S. Government Accountability Office (2021a)).

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	-171.509^{***}	-203.181^{**}
	(35.478)	(75.133)
First-stage F-statistic		77.28
Num. Obs.	815	815
Within R^2	0.915	0.907

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) + Rel. 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) × Price/Full-Risk		-0.338^{***}	
		(0.007)	
Log(Price/Full-Risk) × 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	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)	6092.00	3254.0	3494.00	8270.00
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	3254.0	-2598.00	2178.00
Scenario Fixed Effect (vs current mean)	4321.00	3185.0	1899.00	8898.00
Idiosyncratic Risk	0.00	1426.0	-1239.00	1521.00

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



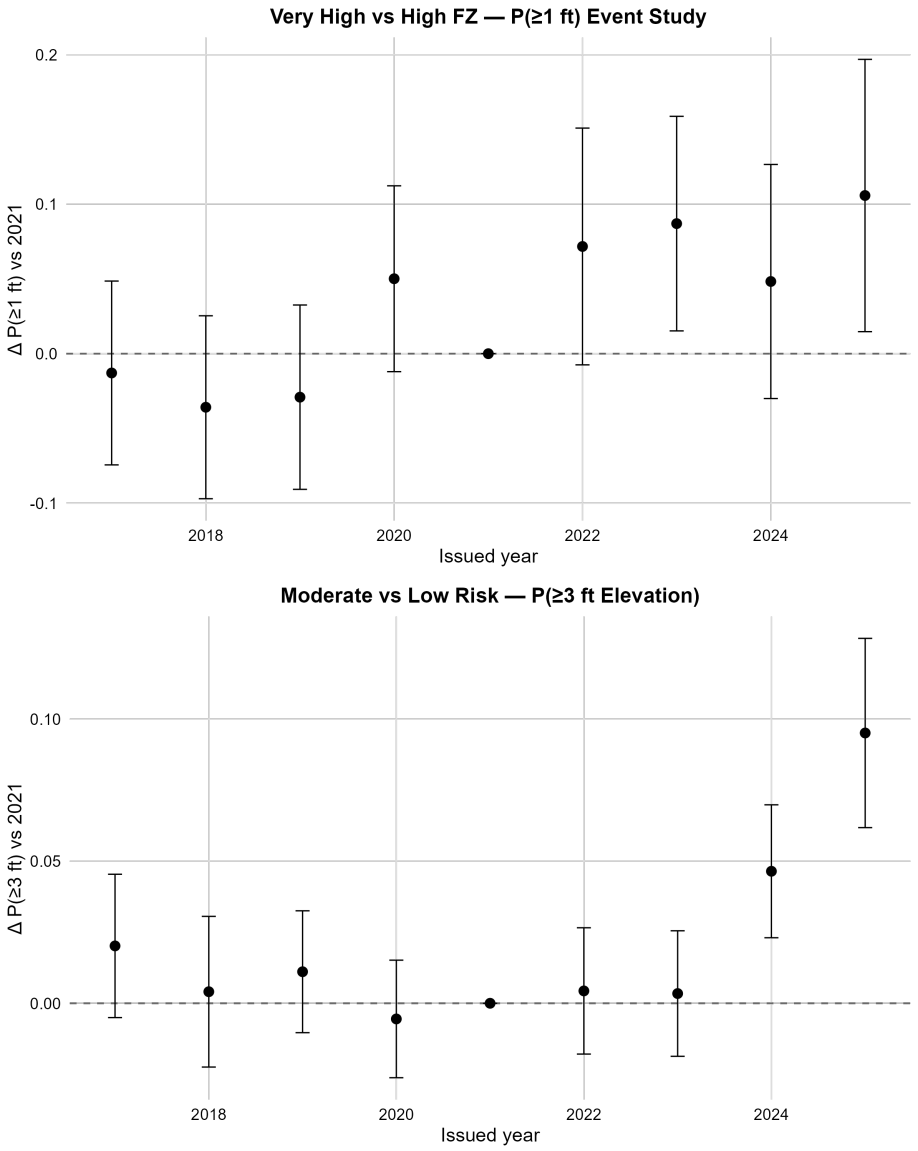
Notes: Scatter plot of projected 2070 percentage premium changes (relative to current RR2.0 levels) for a random sample of NFIP policies under two climate models (MPI on horizontal axis; IPSL on vertical), using synthetic flood data from Gori et al. (2021) under SSP5-8.5 warming scenario. Dashed 45-degree 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

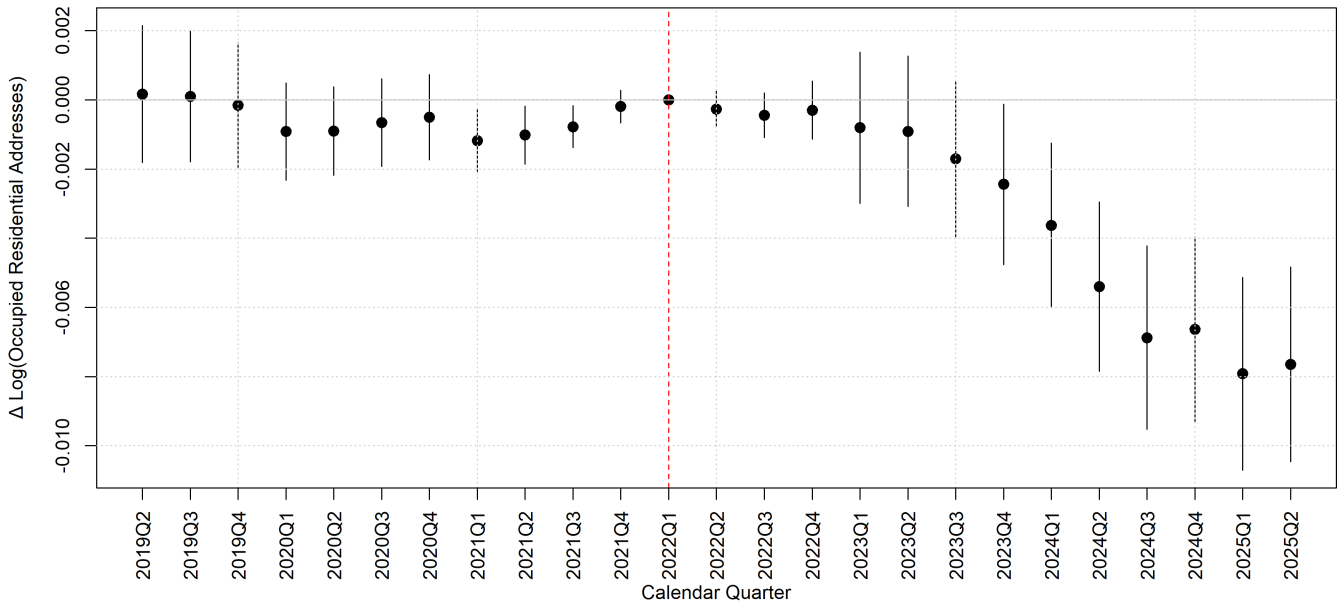
Notes: Each row reports the pure insurance value (in 2023 dollars) of reinstating a 50% average subsidy against idiosyncratic reclassification risk, net of mechanical transfers (c_i), under constant absolute risk aversion (CARA) utility $u(c) = -\exp(-\gamma c)$. γ values span literature estimates: 5×10^{-5}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

Figure 4. Mitigation Changes Under RR2.0



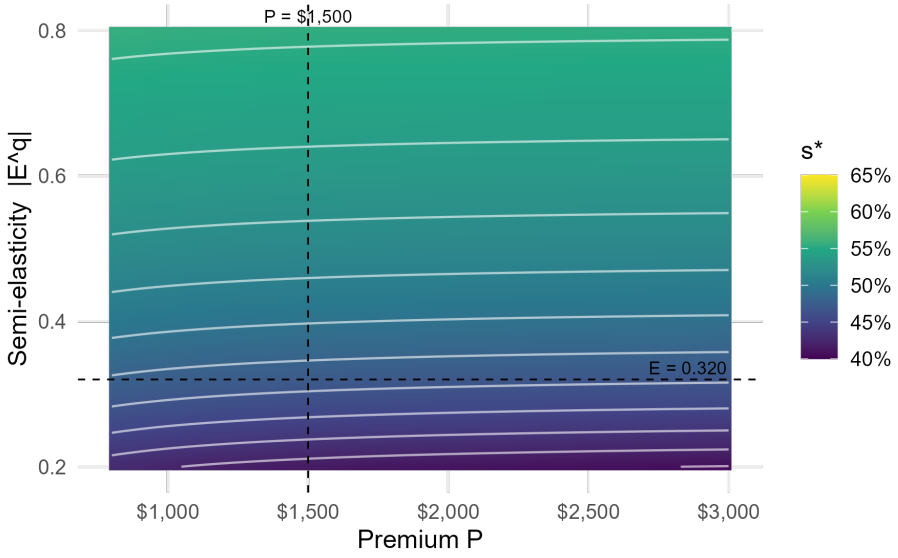
Notes: Event-study estimates from regressions of an indicator for elevating the lowest floor by 1+ feet (relative to base flood elevation inside the SFHA) or 3 feet (relative to local adjacent ground level outside the SFHA) on indicators for years relative to RR2.0 implementation (2022 as reference), 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



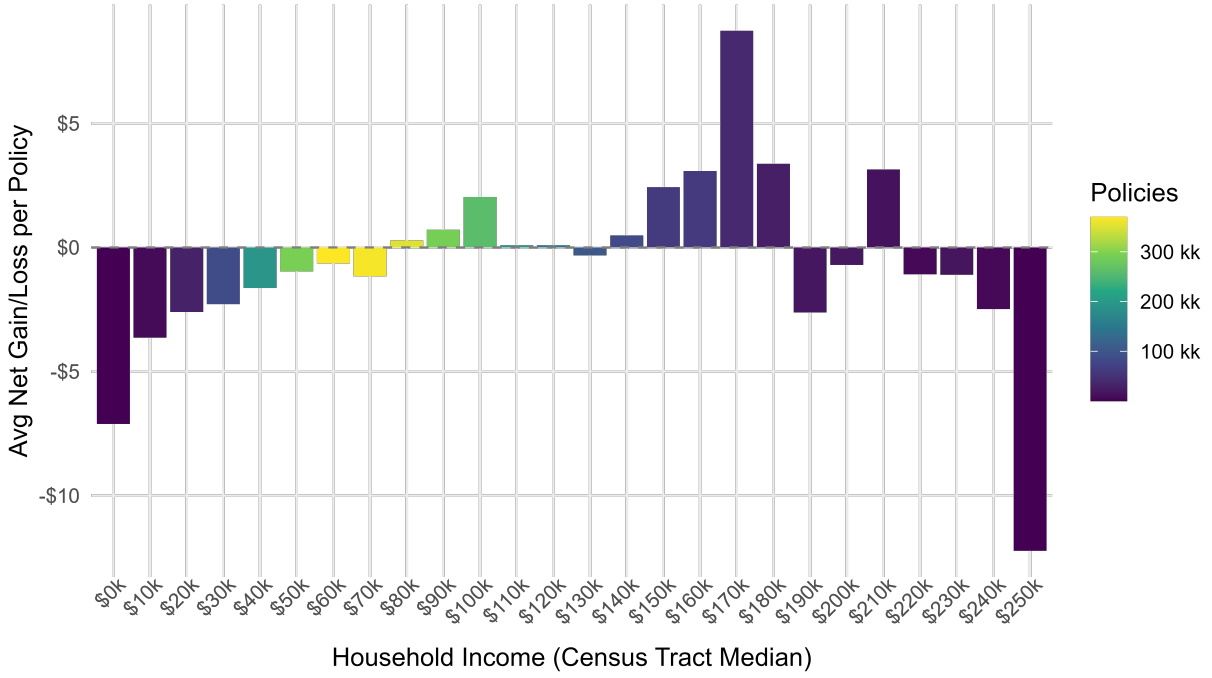
Notes: Event-study estimates of occupied addresses by census-tract × quarter. Treatment is defined by the median NFIP policy in a census tract experiencing a premium increase under RR2.0. Standard errors are clustered at the census-tract level and 95% confidence intervals are shown.

Figure 6. Optimal Subsidy



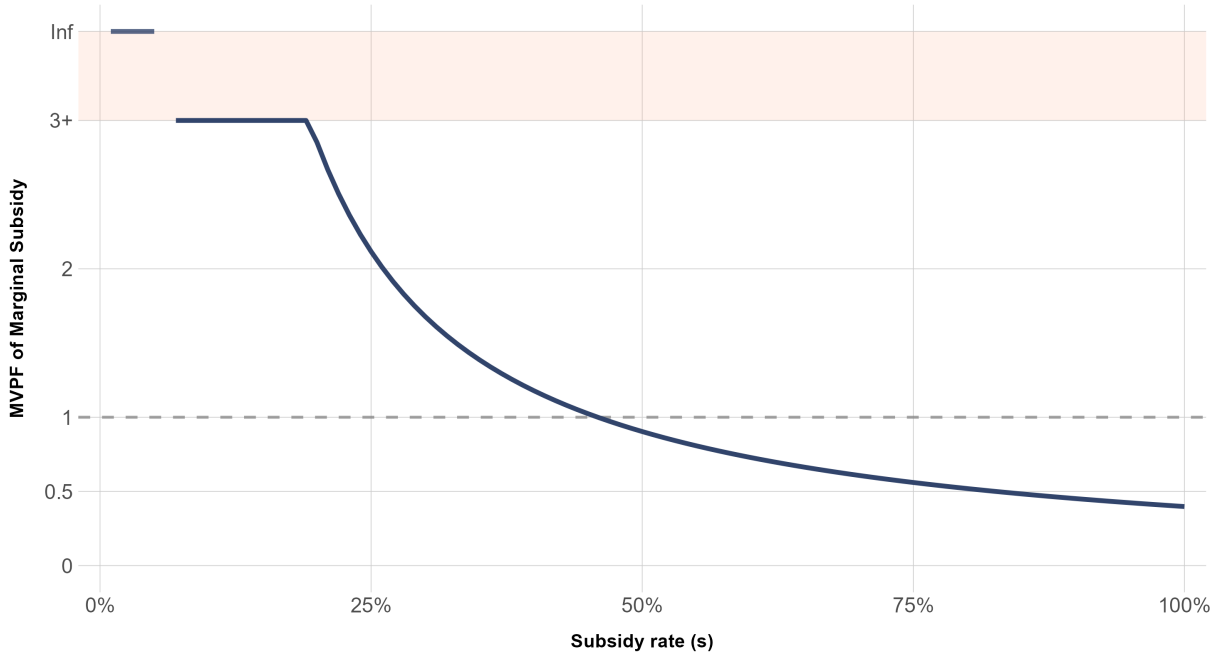
Notes: Heatmap of the optimal subsidy rate s^* for different levels of the average premium and the semi-elasticity of demand. The optimal subsidy solves Proposition 1. Discrete household choices (insure/uninsure; elevate/not; high-/low-risk location) are assumed. The dashed lines represent our baseline assumptions for the average premium ($1,500) and the semi-elasticity of demand (0.32).

Figure 7. Redistributive Implications of RR2.0



Notes: Bar chart of average net gain/loss by income bin (2023 dollars) from a simulated 1% subsidy on NFIP premiums funded by lump-sum tax on all insured households. The color scale represents the number of policies (sample size) per bin. The simulation assigns tract-median income from the ACS to each NFIP policy.

Figure 8. MVPF of an Incremental Subsidy



Notes: Illustration of the MVPF of a marginal expansion of the premium subsidy. When the denominator is negative, we code this as $MVPF = \infty$. We censor positive non-infinite MVPFs at 3. When premium subsidies are funded by lump-sum taxes, the optimal subsidy solves $MVPF = 1$ (horizontal dashed line).

References

2M Research (2020). Flood insurance coverage of federal housing administration single-family homes. Research report, U.S. Department of Housing and Urban Development, Office of Policy Development and Research, Washington, DC. Multidisciplinary Research Team: Task Order 7.

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)*.

Autor, D. H., Palmer, C. J., and Pathak, P. A. (2014). Housing market spillovers: Evidence from the end of rent control in cambridge, massachusetts. *Journal of Political Economy*, 122(3):661–717.

Baily, M. N. (1978). Some aspects of optimal unemployment insurance. *Journal of Public Economics*, 10(3):379–402.

Bakkensen, L. A. and Barrage, L. (2022). Going underwater? flood risk belief heterogeneity and coastal home price dynamics. *The Review of Financial Studies*, 35(8):3666–3709.

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.

CBO (2019). Expected costs of damage from hurricane winds and storm-related flooding. Technical Report 55019, Congressional Budget Office, Washington, DC. April.

Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics*, 90(10):1879–1901.

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 (2014). Fema individual housing program–housing assistance codes legend. Microsoft Excel spreadsheet. FEMA Program Legends (spreadsheet; date listed as 04/28/2014). Accessed 2026-01-12.

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 Emergency Management Agency (FEMA) (2025). Floodsmart: Official site of the national flood insurance program. `https://www.floodsmart.gov/`. Accessed: 2025-10-20.

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

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

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.

Hahn, R. W., Hendren, N., Metcalfe, R. D., and Sprung-Keyser, B. (2024). A welfare analysis of policies impacting climate change. Working Paper 32728, National Bureau of Economic Research.

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

Healy, A. and Malhotra, N. (2009). Myopic voters and natural disaster policy. *American Political Science Review*, 103(3):387–406.

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.

Hendren, N. and Sprung-Keyser, B. (2020). A unified welfare analysis of government policies*. *The Quarterly Journal of Economics*, 135(3):1209–1318.

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.

Hourdin, F., Mauritsen, T., Gettelman, A., Golaz, J.-C., Balaji, V., Duan, Q., Folini, D., Ji, D., Klocke, D., Qian, Y., Rauser, F., Rio, C., Tomassini, L., Watanabe, M., and Williamson, D. (2017). The art and science of climate model tuning. *Bulletin of the American Meteorological Society*, 98(3):589–602.

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.

Insurance Journal (2023). More than 17 million properties face substantial flood risk: Study. *Insurance Journal*. Reports on First Street Foundation's 8th National Risk Assessment.

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.

Kousky, C. and Shabman, L. (2012). The realities of federal disaster aid: The case of floods. RFF Issue Brief 12-02, Resources for the Future, Washington, DC.

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.

LexisNexis Risk Solutions (2025). Lexisnexis u.s. home insurance trends report highlights rising peril severity as catastrophic claims hit seven-year high. Press release (Atlanta, GA).

Lindsay, B. R. and Getter, D. E. (2023). Sba disaster loan interest rates: Overview and policy options. Technical Report R46963, Congressional Research Service, Washington, DC. CRS Report R46963.

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).

Meiler, S., Kropf, C. M., McCaughey, J. W., Lee, C.-Y., Camargo, S. J., Sobel, A. H., Bloemendaal, N., Emanuel, K., and Bresch, D. N. (2025). Navigating and attributing uncertainty in future tropical cyclone risk estimates. *Science Advances*, 11(16):eadn4607.

Mulder, P. (2024). Mismeasuring risk: The welfare effects of flood risk information. SSRN Working Paper. SSRN ID 4966795.

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

NAIC (2024). 2024 property/casualty market share. Technical report, National Association of Insurance Commissioners. Corporate author is the National Association of Insurance Commissioners (NAIC). Report tables note data based on filings received by 03/31/2025.

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.

OMB (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.

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*.

risQ, Inc. (2021). Economic and racial inequality in FEMA SFHA flood zone designations. *Medium*.

Schneider, T., Leung, L. R., and Wills, R. C. J. (2024). Opinion: Optimizing climate models with process knowledge, resolution, and artificial intelligence. *Atmospheric Chemistry and Physics*, 24:7041–7062.

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. Census Bureau (2025). Rental vacancy rate in the united states [rrvrusq156n]. FRED, Federal Reserve Bank of St. Louis. Retrieved on 2025-11-29.

U.S. Department of Housing and Urban Development (2015). Hud announces relief for victims of hurricane sandy: Department will not require repayment of “duplicate” benefits from most homeowners (hud no. 15-114). HUD Archives: News Releases. Accessed: 2026-01-13.

U.S. Government Accountability Office (2014). Federal emergency management agency: Opportunities exist to strengthen oversight of administrative costs for major disasters. Technical Report GAO-15-65, U.S. Government Accountability Office, Washington, DC. Reports $12.7B in FEMA administrative costs for FY2004–2013, equal to 13% of $95.2B DRF obligations (see pp. 1, 8, 12–18).

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 (2021a). Economic injury disaster loan program: Additional actions needed to improve communication with applicants and address fraud risks. Technical Report GAO-21-589, U.S. Government Accountability Office, Washington, DC. GAO-21-589.

U.S. Government Accountability Office (2021b). National flood insurance program: Congress should consider updating the mandatory purchase requirement. Report to Congressional Requesters GAO-21-578, U.S. Government Accountability Office, Washington, DC.

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.

U.S. Government Accountability Office (2023). Flood insurance: Fema’s new rate-setting methodology improves actuarial soundness but highlights need for broader program reform. Technical Report GAO-23-105977, U.S. Government Accountability Office, Washington, DC. Describes WYO expense allowance at about 30% of premium (p. 6/11) and 29.7% in FY2023 (p. 35/40).

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

Wicker, R. and Hyde-Smith, C. (2025). Wicker, hyde-smith demand an end to biden-era flood insurance premiums. U.S. Senate Press Release. Miss. Senators Join Colleagues in Warning FEMA of Fallout for Homeowners, Communities Without Action.

Xu, Y. and Box-Couillard, S. (2024). Social learning about climate risks. *Economic Inquiry*, 62(3):1172–1191.

Zillow Group, Inc. (2026). United states rental market trends and 2026 predictions. Zillow Rental Manager. Accessed: 2026-01-12.

A Additional Tables and Figures

Table A1. Main Results: Total Fiscal Cost (Robustness)

	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: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$. This specification adds a linear control for time since own-county flood to the main specification in Table 2

Table A2. Fiscal Cost Decomposition: FEMA Disaster Assistance

	OLS	**IV**
NFIP Policies per 100 Houses	-150.658^{***}	-169.725^{**}
	(29.865)	(58.584)
Damage Control	Binned	
First-stage F-statistic		77.28
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	$-18.537^{\star}$	-10.445
	(8.758)	(18.418)
Damage Control	Binned	
First-stage F-statistic		77.28
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: HMGP Payouts

	OLS	**IV**
NFIP Policies per 100 Houses	-2.917	-21.662
	(9.978)	(19.974)
Damage Control	Binned	
First-stage F-statistic		77.28
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 (Default Costs)

	OLS	**IV**
NFIP Policies per 100 Houses	0.295^{*}	0.855
	(0.288)	(0.835)
Damage Control	Binned	
First-stage F-statistic		72.82
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
	(0.165)	(1.020)
Damage Control	Binned	
First-stage F-statistic		70.00
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: $+\ p < 0.10, *\ p < 0.05, **\ p < 0.01, ***\ p < 0.001$.

B Empirical Appendix

B.1 FEMA Rejections

In this section we analyze the reasons that a household's claim to FEMA for assistance might be rejected. Specifically, we study all registrations for the Housing Assistance (HA) program in which flood damage was reported and categorize them based on Federal Emergency Management Agency (2014). We bucket together substantially similar reasons, and report the results in Table A8. Households are often rejected for multiple reasons and therefore are counted in multiple buckets.

The most common reasons for rejection are that there is insufficient damage to warrant housing assistance, or that the household has insurance (e.g. flood insurance) that already covers the damage.

Most notably, households can be rejected because they, or their NFIP community, are not in compliance with flood insurance requirements. However, this is empirically very infrequent: fewer than 1% of claims are rejected because of flood insurance noncompliance, NFIP community non-compliance or other recertification-related reasons.

B.2 Robustness to Discounting Assumptions

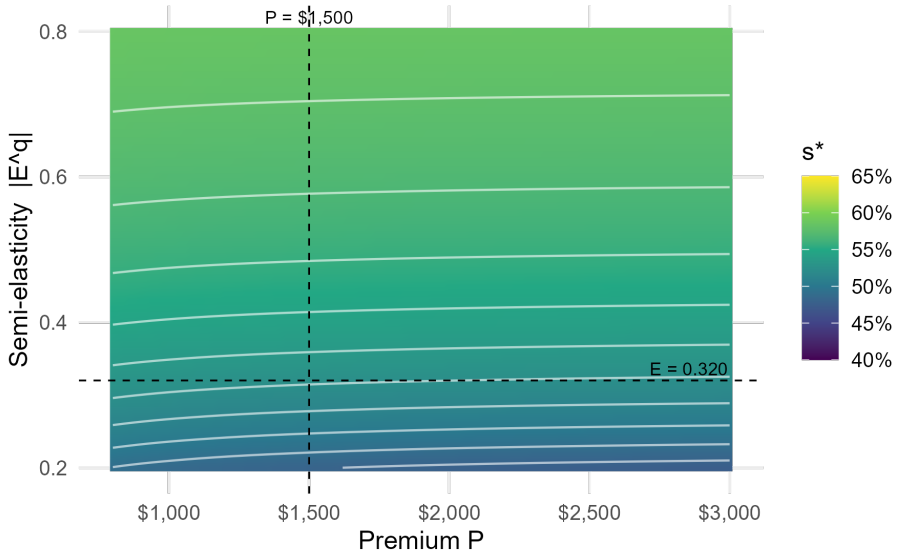
The premium risk individuals face due to future climate uncertainty is measured in 2070. Nevertheless, we are solving for the optimal subsidy today. In the main paper we do not discount this by appealing to an offsetting growth in house prices. To test the robustness of our policy prescriptions to this 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 9, which follows.

Table A8. FEMA Ineligibility Reasons

Ineligibility reason (bucket)	Count	Share of ineligibility codes
Insufficient damage	292,192	37.623%
No relocation	138,882	17.883%
Other / Unknown	89,814	11.565%
Has flood insurance / flood losses covered by policy	76,140	9.804%
Other insurance (insured / ineligible due to insurance)	72,303	9.310%
Duplicate / linked duplicate review	23,657	3.046%
Ownership not verified	21,686	2.792%
Occupancy not verified	13,910	1.791%
Identity verification failed	11,406	1.469%
Damage not disaster-related	11,402	1.468%
Other (explicit INO/IOR)	6,591	0.849%
Not primary residence	4,810	0.619%
Missing/insufficient documentation or signature	4,759	0.613%
Flood insurance noncompliance (required policy not maintained)	3,317	0.427%
Withdrawn / applicant withdrew voluntarily	3,279	0.422%
No contact (inspection/contact failure)	1,503	0.194%
Missed inspection	391	0.050%
Reported no damage	240	0.031%
NFIP/community status (sanctioned community in SFHA)	149	0.019%
CBRA/OPA restriction	148	0.019%
Recertification-related	30	0.004%
Over program maximum	20	0.003%

Notes: The table reports counts and shares of Housing Assistance ineligibility reason codes in FEMA's Individuals and Households Program (IHP) Valid Registrations data. Registrations with multiple ineligibility codes contribute once to each relevant bucket.

Figure 9. Optimal Subsidy With Discounted Reclassification Risk Estimates



Notes: Heatmap of optimal subsidy rate s^* that solving the calibrated benefit-cost condition from Proposition 1 when the reclassification risk benefits are discounted by $1/(1.02)^{35}$.

B.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: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

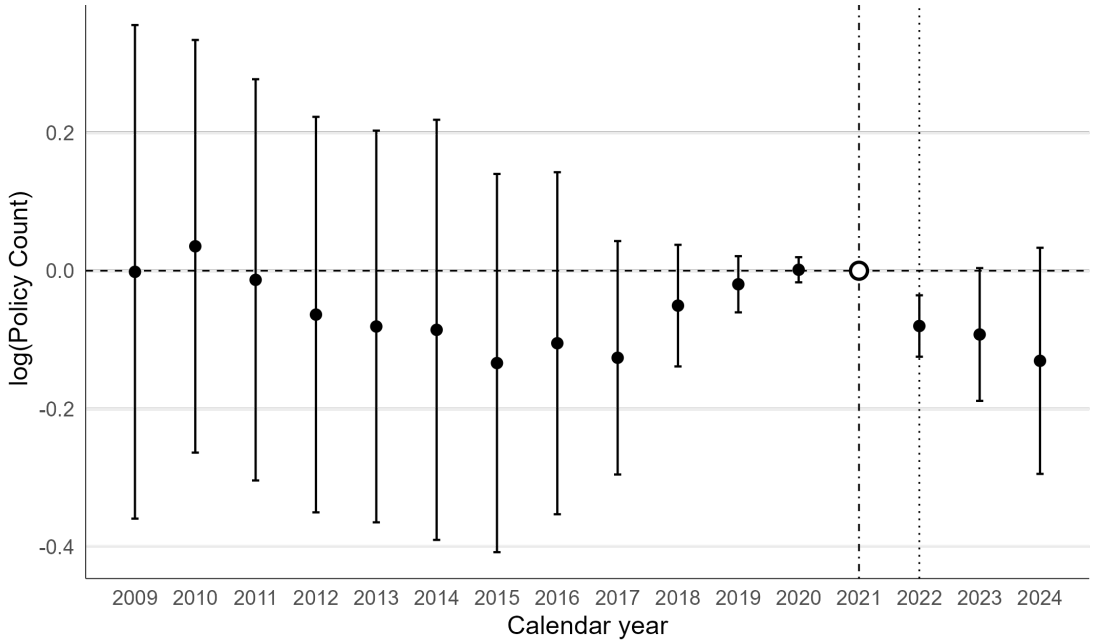
B.4 Between-County Demand Elasticity Estimates

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 10 plots the estimated $\beta_{q'}$.

Figure 10. 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.

B.5 Intensive Margin NFIP Response

Similarly to Section 5, we estimate the response of increased prices due to RR2.0 on household's choice of coverage limit. That is, their demand response along the intensive margin.

The specifications are almost identical to Section 5, except the outcome is the total coverage (building plus contents) per policy. Specifically,

$$\text{Total NFIP Coverage}_{i,t+1} = \beta \ln \tilde{P}_{i,t} + \mu_i + \delta_t + \varepsilon_{i,t+1}, \tag{7}$$

where μ_i are policy (contract) fixed effects, absorbing all time-invariant heterogeneity, and δ_i are year fixed effects. Standard errors are clustered at the policy level. Our identification comes from *within-policy* variation in the premium due to RR2.0. Specifically, among the houses that remain insured, when faced with a higher price due to RR2.0, we analyze whether they adjust their coverage limit. We additionally run the same heterogeneity specifications as in Section 5. The results are in Table A10.

We see a statistically precise, but economically negligible effect decrease along the intensive

Table A10. Coverage vs. Price Relative to Full-Risk: Semi-Elasticity and Heterogeneity (Intensive Margin)

	(1) Semi-Elasticity Only	(2) + Relative Price Interaction	(3) + SFHA Interaction
Log(Price/Full-Risk)	-3023.857^{***}	-3079.996^{***}	-4511.683^{***}
	(82.528)	(86.153)	(94.206)
Log(Price/Full-Risk) $\times$ Price/Full-Risk		382.007^{*}	
		(193.320)	
Log(Price/Full-Risk) $\times$ SFHA			3540.875^{***}
			(118.484)
Panel ID FE	Yes	Yes	Yes
Policy Year FE	Yes	Yes	Yes
N	6,852,683	6,852,683	6,852,683
Adj. R^2	0.991	0.991	0.991
Within R^2	0.001	0.001	0.001

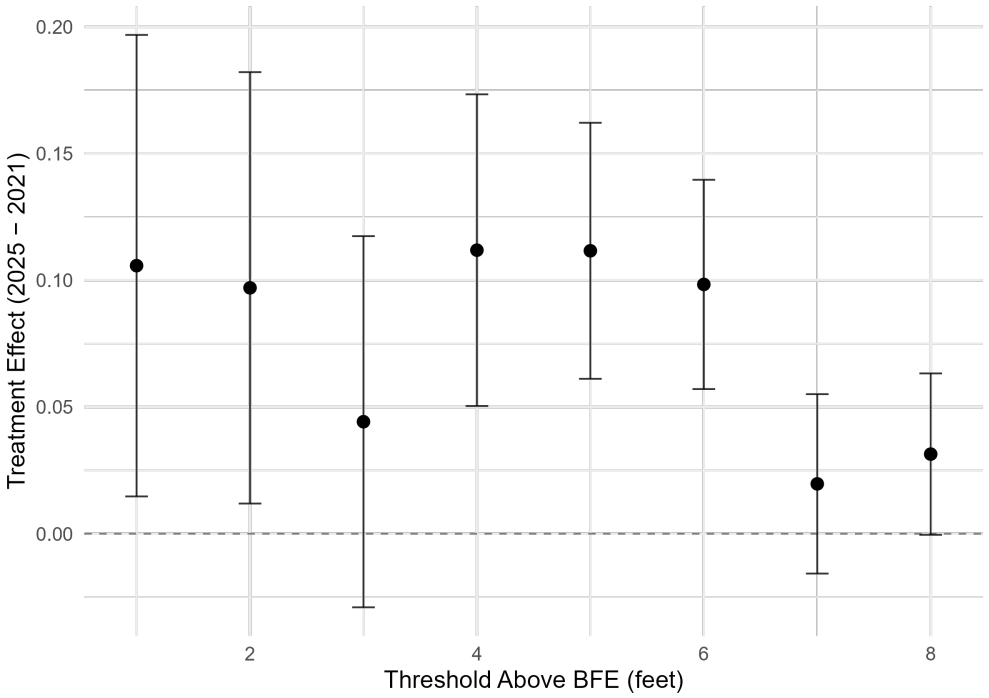
Notes: Each column reports estimates from a fixed-effects model of NFIP policy coverage amount (the intensive margin) on the log of the relative renewal price $\tilde{P}_{i,t} \equiv P^{\text{offer}}_{i,t+1}/R_{i,t+1}$, where $P^{\text{offer}}_{i,t+1}$ is the offered renewal premium and $R_{i,t+1}$ is 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 panel ID level in parentheses. Significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$.

margin. The average of level of coverage per policy is $316,000. Thus, the estimates are on the order of a 1% decline in coverage limit. This is because the NFIP coverage limits ($250,000 for building and $\$100,000$ for contents) are likely binding throughout.

B.6 Different Elevation Definitions

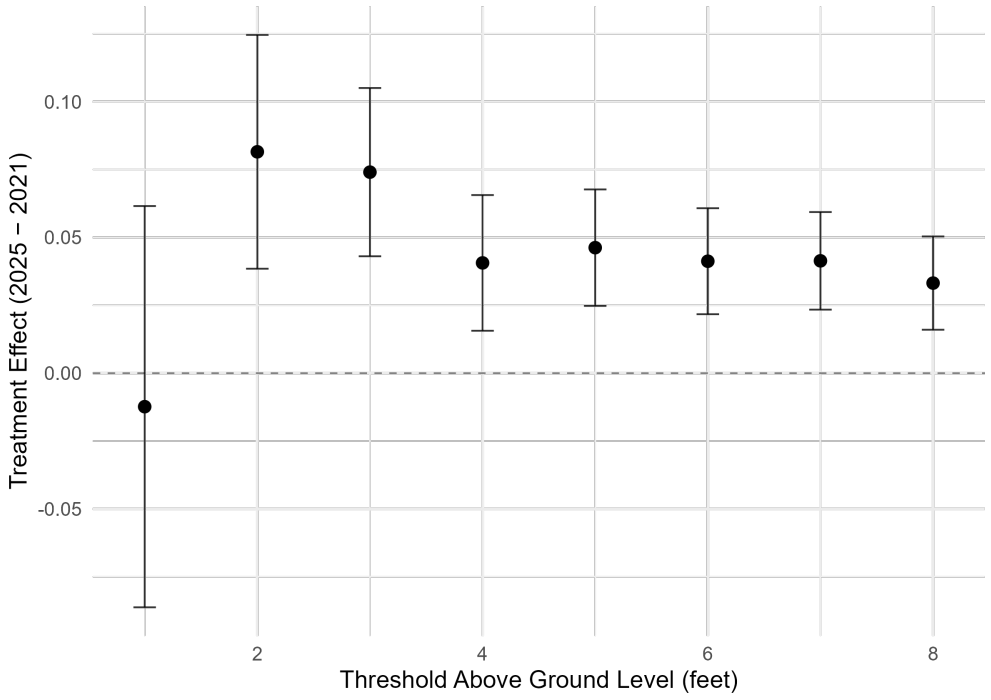
In this section we test the robustness of our mitigation findings to the threshold chosen for elevation. Specifically, we re-run equation (3) with different thresholds. In each case, we do not present the entire event study, rather the coefficient from 2025. Thus, the results below should be thought of as a difference-in-differences coefficient of 2025 against the base year, 2021. As in the main paper, inside the SFHA the very-high risk flood zones are compared to the high-risk; outside the SFHA the moderate risk flood zones are compared to the low risk. The results are in Figures 11 and 12 below.

Figure 11. Elevation Threshold Robustness Inside SFHA



Notes: Each coefficient is the 2025 coefficient from specification (3) where the threshold is varies according to the horizontal axis. We restrict to the very-high-risk and high-risk flood zones inside the SFHA, and code the former as treated, the latter as control. 95% confidence intervals are shown.

Figure 12. Elevation Threshold Robustness Outside SFHA



Notes: Each coefficient is the 2025 coefficient from specification (3) where the threshold is varies according to the horizontal axis. We restrict to the moderate- and low-risk flood zones outside the SFHA, and code the former as treated, the latter as control. 95% confidence intervals are shown.

Figures 11 and 12 show that the 10% treatment effects we have used as our benchmark in the main paper are relatively robust to the precise choice of threshold. Naturally, they decline as the threshold increases.

B.7 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 A11 below.

Table A11. 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 A11 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.8 Rolling Back RR2.0 Details

Setup and notation. Let P_i^{new} be the RR 2.0 full-risk premium and P_i^{old} the pre–RR 2.0 paid premium. Define the "rollback distance"

$$d_i \;\equiv\; 1 - \frac{P_i^{\text{old}}}{P_i^{\text{new}}} \in [0,1],$$

interpreted as the fraction of the new premium one would have to remove to return to the old level. We use premium weights $w_i = P_i^{\text{new}}$ and the weighted mean $\bar{d} = \sum_i w_i d_i / \sum_i w_i$, and define $z_i = d_i - \bar{d}$ (so $\sum_i w_i z_i = 0$).

We model heterogeneity in FEMA spillovers and demand elasticities as linear functions of d_i:

$$\beta_i \;=\; \beta_0 + \beta_1 d_i, \qquad E_i \;=\; e_0 + e_1 d_i, \qquad \alpha_i \;\equiv\; 100\,\frac{-\beta_i}{L_f},$$

where β_i is the FEMA dollars saved per insured home for a 1% log price cut (conditional on flood), E_i is the price semi-elasticity of coverage, and L_f is expected loss conditional on flood.

Level vs. twist. A *level* rollback applies a uniform ad valorem subsidy $s \in [0,1]$ to the new premium: each household gets $s_i = s$.

A *twist* adds a budget-neutral, cross-sectional reallocation proportional to z_i:

$$s_i^{\text{raw}}(s,t) = s + t\, z_i,$$

where $t \geq 0$ is the twist intensity. We enforce $s_i^{\text{eff}} \in [0,1]$ by bounding s_i^{raw} to $[0,1]$, then *rebalance dollars* so that the post-bounding grant bill equals the uniform-s bill:

$$\sum_i P_i^{\text{new}}\, s_i^{\text{eff}} \;=\; \sum_i P_i^{\text{new}}\, s.$$

This preserves $\sum_i w_i z_i = 0$ in dollars after caps.

Estimation/calibration. We set $(\beta_0, \beta_1, e_0, e_1)$ from reduced-form relationships (see Tables 3 and A11) with d_i; in the baseline we use

$$\beta_0 = -116,\ \beta_1 = -112, \qquad e_0 = 0.608,\ e_1 = -0.334.$$

Procedure We evaluate a grid of uniform subsidies $s \in [0,1]$ and, for each s, a grid of twist intensities $t \in [0, \bar{t}]$ (baseline $\bar{t} = 0.30$ with step 0.02). For each pair (s,t) we: (i) form $s_i^{\text{raw}} = s + tz_i$; (ii) clip to $[0,1]$; (iii) rebalance dollars to match the uniform-s bill; (iv) compute components and net welfare; and (v) record, for each s, the range of net welfare over $t \in [0, \bar{t}]$ to produce a "twist band."

Results We find that welfare is increasing in the level shift up, assuming no twist, to $s \approx 0.45$–0.46, and then declines. This is a restatement of our baseline optimal subsidy being 46%.

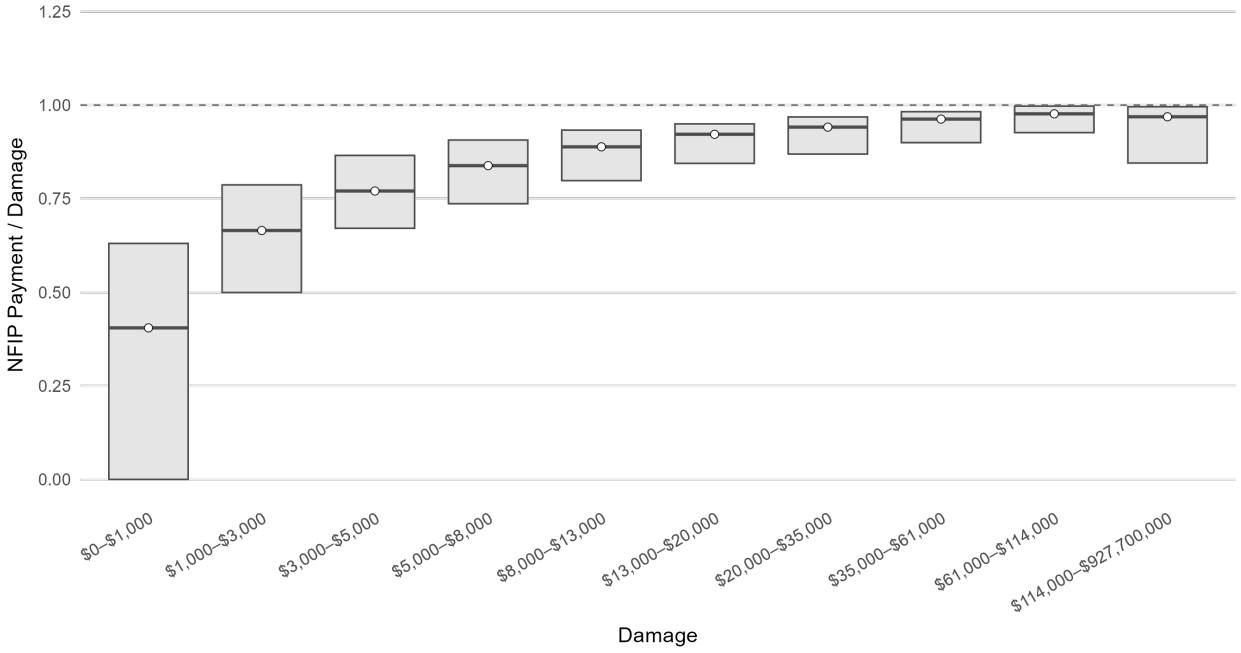
The twist reallocates from those with a small pre-RR2.0 subsidy to those with a large one. The heterogeneity estimates show that these net recipients from a twist have bigger FEMA savings and lower demand elasticities than the net payers. The larger FEMA savings makes a twist good. The latter have a cost (less uptake of insurance) but a benefit (lower cost of subsidy due to update). The cost is approximately constant in the level shift; the benefit increases in the level shift. Hence, the twist, for larger s, adds to welfare. A 20% twist, at a level subsidy of approximately 46% (our baseline optimum) adds \$3 in welfare per person (a 4% increase on the \$84 per person welfare gain of a 46% level shift).

B.9 Estimating the Completeness of NFIP Coverage

In this section we analyze the completeness of the coverage that the NFIP provides. The NFIP has per-policy coverage limits of $\$250,000$ for building damage and $\$100,000$ for contents. This raises the concern that, especially for big floods, some damage might not be insured. To study this, we compute, for every NFIP claim, the ratio of paid insurance beenfits to assessed damage. We do

this separately by damage level to understand if underinsurance covaries with flood severity. The results are in Figure 13 below.

Figure 13. Insurance Coverage Relative to Flood Damage



Notes: Each box summarizes the distribution of the claim-level insurance payment per dollar of damage. Boxes show the inter-quartile range; the bold horizontal segment marks the median. The dashed horizontal line at 1.0 indicates full coverage.

Figure 13 shows that for the most severe floods coverage is essentially complete. Over 95 cents of each dollar of damage is covered for the median large flood. The lower ratios observed for tiny floods are likely due to misreporting. Overall, most claims are well below both the NFIP coverage limits ($350,000) and the average chosen coverage ($315,000). For this reason, in the main paper we assume NFIP coverage is complete.

B.10 Behavioral Misperception Wedge (Internality Correction)

Here we derive a quantitative measure, in the spirit of Section 9, of the welfare effects of the internality caused by household underperception of flood risk and therefore of the value of flood insurance. The central object is the difference between the decision-relevant WTP for insurance that is biased downward by underperception of flood risk, and the unbiased welfare-relevant WTP.

Setup Consider a Bernoulli flood-loss risk: household i experiences loss L_f with probability p_i and zero otherwise. Full insurance eliminates this risk. Under CARA utility with coefficient γ, a

standard small-risk approximation implies that the certainty-equivalent value of eliminating this risk is approximately

$$WTP_i^* \;\approx\; \underbrace{p_i L_f}_{\text{expected loss}} \;+\; \underbrace{\frac{\gamma}{2}\mathsf{Var}(L_i)}_{\text{risk premium}} \;=\; p_i L_f \;+\; \frac{\gamma}{2}\,p_i(1-p_i)L_f^2. \tag{8}$$

Internality Wedge Suppose households misperceive the flood probability, evaluating the same insurance contract as if the probability were $\bar{p}_i$ rather than p_i.

Define the internality wedge:

$$\begin{aligned}\Delta_i^B &\equiv WTP_i^* - \widetilde{WTP}_i \\ &\approx (p_i - \bar{p}_i)L_f \;+\; \frac{\gamma}{2}\Big[p_i(1-p_i) - \bar{p}_i(1-\bar{p}_i)\Big]L_f^2.\end{aligned} \tag{9}$$

Parameterizing $\bar{p}_i = kp_i$ with $k < 1$, then (9) becomes

$$\Delta_i^B \;\approx\; (1-k)p_i L_f \;+\; \frac{\gamma}{2}\Big[p_i(1-p_i) - kp_i(1-kp_i)\Big]L_f^2. \tag{10}$$

Using $p_i L_f = P_i$ (i.e., assuming actuarial fairness) we get:

$$\Delta_i^B \;\approx\; (1-k)P_i \;+\; \frac{\gamma}{2}\left[\frac{P_i}{L_f}\left(1-\frac{P_i}{L_f}\right) - k\frac{P_i}{L_f}\left(1-k\frac{P_i}{L_f}\right)\right]L_f^2. \tag{11}$$

This internality wedge is per each household marginal to a subsidy. To convert this to be per household, note that a premium subsidy that reduces log prices by $\Delta \ln \tilde{P}$ induces additional coverage of $\Delta q \simeq \eta^q \,\Delta \ln \tilde{P}$. Hence the welfare gain per household is

$$\Delta W^B \;\approx\; \overline{\Delta^B} \;\cdot\; \eta^q \;\cdot\; \Delta \ln \tilde{P}, \tag{12}$$

where $\overline{\Delta^B}$ is the average wedge among the marginal households induced into insurance.

Calibration. We take $L_f = \$31{,}921$ and $\bar{P} = \$1{,}500$, which implies $p = \bar{P}/L_f \approx 0.047$. Following Bakkensen and Barrage (2022), we set $k = 0.57$ (a 43% underperception), and assume a coefficient of absolute risk aversion of $\gamma = 2 \times 10^{-5}$. Plugging these into (10) yields:

$$\Delta^B \;\approx\; (1{-}0.57){\times}0.047{\times}31{,}921 + \frac{2\times10^{-5}}{2}\Big[0.047(1{-}0.047){-}0.0268(1{-}0.0268)\Big]{\times}31{,}921^2 \;\approx\; \$838$$

per newly insured household-year.

Finally, for a 1% log price change ($\Delta \ln \tilde{P} = 0.01$) and our baseline semi-elasticity $\eta^q = 0.32$, the internality benefit is approximately $\$2.68$ per house-year for a 1% price cut.

C Theoretical Appendix

C.1 General Formula for the Optimal Subsidy

In this section we study the optimal subsidy level s that obtains without making any of the assumptions A1-A5. 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 2 (Optimal Subsidy: General Condition)**.** *The optimal premium subsidy s solves:*

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

The right-hand side terms are identical to Proposition 1, as is the first term on the left-hand side. The second term on the right-hand side accounts for the fact that those who have predictably higher future flood premiums might have higher marginal utility than average. Empirically, this depends on the covariance between consumption and flood risk. In the main paper we only wish to count the insurance benefits; we assume away this redistributive benefit or cost between those known to be high or low risk. Finally, the remaining terms on the right-hand side arise when there is adverse or advantageous selection. Although there is no evidence of selection in flood insurance, in other settings this is likely important.

C.2 Extended Model With Misperceptions

This appendix adds a behavioral friction layer to the model in Section 3. Specifically, households misperceive their level of flood risk. We show how this generates, in addition to the sufficient statistics of Proposition 1, internalities that the planner's optimal subsidy must account for.

For each household i, let $\lambda_i \in [0,1]$ capture the fraction by which the household underweights expected flood loss. That is, the household makes their choices of q, m, x as if L_i were $(1-\lambda_i)L_i$. That is, they maximize expected utility of consumption where consumption is

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

The planner, however, evaluates welfare according to the true expected loss; they care about ex-post experienced shocks, not ex-ante misperceived shocks.

we the (experienced minus decision) internality gaps, evaluated at the current choices (q,m,x):

$$\begin{aligned}
\Delta_i^q &\equiv \lambda_i(1-f)\,\bar{P}_i,\\
\Delta_i^m &\equiv \lambda_i(1-f)\,(1-q_i^*)\Big(-\tfrac{\partial \bar{P}_i}{\partial m}\Big),\\
\Delta_i^x &\equiv \lambda_i(1-f)\,(1-q_i^*)\Big(-\tfrac{\partial \bar{P}_i}{\partial x}\Big),
\end{aligned}$$

where $\bar{P}_i \equiv \mathbb{E}_\omega[P_i(\omega)]$ and we use actuarial fairness so that $\partial\mathbb{E}[L_i]/\partial\cdot = \partial\bar{P}_i/\partial\cdot$. Intuitively, with underperceived risk ($\lambda_i > 0$) households undervalue the uninsured–loss benefit of marginally more coverage and safety investments. Because a subsidy generates more coverage, the former internality pushes in favor of a larger subsidy. In contrast, since a higher subsidy dampens risk reduction through mitigation or location choice, the latter internalities push against a subsidy.

Formally, these internalities augment the optimality condition in Propositions 2 and 1.

Proposition 3. *With misperceptions, the left-hand sided of the optimality equation in Propositions 2 and 1 is augmented by:*

$$-\frac{1}{1-s}\,\mathbb{E}_\omega\left[\sum_i \mathbb{E}_\xi u_i'(\omega,\xi)\left(\Delta_i^q\,\mathcal{E}_i^q + \Delta_i^m\,\mathcal{E}_i^m + \Delta_i^x\,\mathcal{E}_i^x\right)\right],$$

Binary Coverage With binary coverage (either no insurance or full insurance), the mitigation and location internationalities are zero: $= \Delta_i^m = \Delta_i^x = 0$. We empirically verify that coverage is essentially complete in Appendix B.9.

Consequently, the only internality term that remains is from the choice of insurance. This makes subsidies even more favorable than without a misperception, and so we conclude that our estimates of the optimal subsidy are, if anything, conservative.

C.3 Relating Propositions 1 & 2 to Baily–Chetty

A novelty of our approach is that it accommodates continuous risk outcomes. In contrast, standard optimal social insurance ("Baily–Chetty") models (Chetty, 2006; Baily, 1978) use a binary outcome. We show our formulation nests the standard case.

Insurance term with macro–idiosyncratic correction (combined form). From Proposition 2, the idiosyncratic insurance and the macro–idiosyncratic correction combine to

$$(1-s)\sum_i q_i^* \,\mathsf{Cov}_\omega\big(u_i'(\omega)-\bar{u}'(\omega),\,\varepsilon_{i\omega}\big). \tag{13}$$

Example 1 (Two states; Baily–Chetty analogue)**.** *Let $\omega \in \{B,G\}$ with probabilities $1-\pi$ and π, and suppose the idiosyncratic component is mean–zero across states: $\varepsilon^B = \pi\,\Delta\varepsilon$, $\varepsilon^G = -(1-\pi)\,\Delta\varepsilon$ ($\Delta\varepsilon > 0$). For a covered household ($q^* \in \{0,1\}$), we have*

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

Proof. By definition,

$$\mathsf{Cov}_\omega\big(u'-\bar{u}',\varepsilon\big) = (1-\pi)\big(u'^B-\bar{u}'^B\big)\,\varepsilon^B \;+\; \pi\big(u'^G-\bar{u}'^G\big)\,\varepsilon^G.$$

Under *purely idiosyncratic risk*, $\bar{u}'^B = \bar{u}'^G$ is constant across states, so $\mathsf{Cov}_\omega(\bar{u}'(\omega),\varepsilon_\omega) = 0$ and hence

$$\mathsf{Cov}_\omega\big(u'-\bar{u}',\varepsilon\big) = \mathsf{Cov}_\omega(u',\varepsilon).$$

With $\varepsilon^B = \pi\Delta\varepsilon$, $\varepsilon^G = -(1-\pi)\Delta\varepsilon$ and $\mathbb{E}_\omega[\varepsilon] = 0$,

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

Since $\mathsf{Var}_\omega(\varepsilon) = \pi(1-\pi)(\Delta\varepsilon)^2$, we obtain

$$\mathsf{Cov}_\omega(u'-\bar{u}',\varepsilon) = \left(\tfrac{u'^B-u'^G}{\Delta\varepsilon}\right)\mathsf{Var}_\omega(\varepsilon),$$

and multiplying by $(1-s)q^*$ yields the stated result. □

D Proofs

We first prove the more general Proposition 2, after which the more specialized Proposition 1 naturally follows under further assumptions.

Proof of Proposition 2

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 $\sum_i T(\omega) = \mathsf{Out}(\omega; s),$ with $T_i(\omega) \equiv T(\omega)\ \forall i$. By assumption of lump-sum taxes $\mu(\omega) = \bar{u}'(\omega) \equiv \sum_i \mathbb{E}_\xi u_i'(\omega, \xi)$.

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

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

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

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

The budget term contributes

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

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:

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

But because the taxes are lump sum, and in particular $\mu(\omega) = \bar{u}'(\omega) \equiv \sum_i \mathbb{E}_\xi u_i'(\omega,\xi)$, the middle terms cancel $\mathbb{E}_\omega\mathbb{E}_\xi\Big[\sum_i u_i'(\omega,\xi)\,(-\dot{T}(\omega))\Big] + \mathbb{E}_\omega\Big[\mu(\omega)\sum_i \dot{T}(\omega)\Big] = 0$ 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

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

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):

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

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

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 $\frac{d\mathcal{L}}{ds} = 0$ at the optimum yields $D + B = 0$. From (A.6), $B = \frac{1}{1-s}\tilde{K}$, where $\tilde{K} = \mathbb{E}_\omega\Big[\mu(\omega)\Big(\sum_i \left((s-f)\, P_i(\omega)\right) \mathcal{E}_i^q + \sum_i \left(s\, q_i + f(1-q_i)\right) \frac{\partial P_i(\omega)}{\partial m_i}\, \mathcal{E}_i^m + \sum_i \left(s\, q_i + f(1-q_i)\right) \frac{\partial P_i(\omega)}{\partial x_i}\, \mathcal{E}_i^x\Big)\Big]$. Thus, $D = -\frac{1}{1-s}\tilde{K}$. Recalling that because of lump-sum taxes $\mu(\omega) = \bar{u}'(\omega)$ yields the result.

Proof of Proposition 1

Start from Proposition 2. Recall the assumptions:

A1 $\mu(\omega) = \bar{\mu}$ for all ω; dynamic-flood risk is small relative to the economy.

A2 For each ω, $\mathrm{Cov}_i\left(q_i^*, u_i'(\omega)\right)$; no selection into insurance

A3 $\mathrm{Cov}_i(a_i, q_i \mathbb{E}_\omega[u_i'(\omega) - \bar{u}'(\omega)]) = 0$; rules out redistribution on predictable individual risk.

Fix any state ω and denote the cross-sectional average marginal utility by $\bar{u}'(\omega) \equiv \mathbb{E}_i[u_i'(\omega)]$. Assumption A2 states that for each ω,

$$\mathrm{Cov}_i(q_i^*,\, u_i'(\omega)) = 0 \iff \mathbb{E}_i[q_i^* u_i'(\omega)] = \mathbb{E}_i[q_i^*]\, \mathbb{E}_i[u_i'(\omega)] = \mathbb{E}_i[q_i^*]\, \bar{u}'(\omega).$$

Rewriting this yields

$$\sum_i q_i^* u_i'(\omega) = \Big(\sum_i q_i^*\Big) \bar{u}'(\omega) \quad \Longrightarrow \quad \sum_i q_i^* (u_i'(\omega) - \bar{u}'(\omega)) = 0 \quad \text{for every } \omega.$$

Hence

$$\mathbb{E}_\omega\Big[\sum_i q_i^* (u_i'(\omega) - \bar{u}'(\omega))\Big] = 0, \qquad \mathrm{Cov}_\omega\Big(\bar{P}(\omega),\, \sum_i q_i^* (u_i'(\omega) - \bar{u}'(\omega))\Big) = 0,$$

so the "average selection" and "$\omega-$specific selection terms vanish.

By A2, for each ω,

$$\sum_i q_i^* (u_i'(\omega) - \bar{u}'(\omega)) = 0 \quad \Longrightarrow \quad \sum_i q_i^*\, \mathbb{E}_\omega[u_i'(\omega) - \bar{u}'(\omega)] = 0.$$

Using the definition of covariance, A3 implies

$$0 \;=\; \mathrm{Cov}_i\Big(a_i,\; q_i^*\,\mathbb{E}_\omega[u_i'(\omega)-\bar u'(\omega)]\Big) = \frac{1}{n}\sum_i a_i q_i^*\,\mathbb{E}_\omega[u_i'(\omega)-\bar u'(\omega)] - \bar a\;\frac{1}{n}\sum_i q_i^*\,\mathbb{E}_\omega[u_i'(\omega)-\bar u'(\omega)].$$

The second term on the right is zero by A2, so

$$\sum_i a_i q_i^*\,\mathbb{E}_\omega[u_i'(\omega)-\bar u'(\omega)] = 0.$$

That is, the "risk-class redistribution" term is zero.

A1 says $\mu(\omega) = \bar\mu$. Then for any state-varying object $g(\omega)$, $\mathbb{E}_\omega[\mu(\omega)g(\omega)] = \bar\mu\,\mathbb{E}_\omega[g(\omega)]$ and $\mathrm{Cov}_\omega(\mu(\omega), g(\omega)) = 0$. Hence the RHS becomes

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

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

A4 (*Insured-only behavioral margins*) For $q_i^*(s) = 0$, $\dfrac{\partial m_i^*}{\partial \ln \mathbb{E}_\omega[\tilde P_i]} = 0$ and $\dfrac{\partial x_i^*}{\partial \ln \mathbb{E}_\omega[\tilde P_i]} = 0$.

A5 (*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.[23]

Hence, by A4, $\mathcal{E}_i^m = \mathcal{E}_i^x = 0$ whenever $q_i^*(s) = 0$, so the mitigation/location pieces reduce to $s\,q_i^*(s) \times (\cdots)$ and can be re-indexed to $\mathcal{I}_1(s)$ with conditional semi-elasticities $\mathcal{E}_i^{m|1}, \mathcal{E}_i^{x|1}$. By A5, 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). This yields the result.

Proof of Proposition 3

As in the main proof, write

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

with $\mu(\omega) = \bar u'(\omega) = \sum_i \mathbb{E}_\xi u_i'(\omega,\xi)$ and $\mathrm{Out}(\omega;s) = \sum_i s\,q_i P_i(\omega) + \sum_i f(1-q_i)P_i(\omega)$. Differentiating w.r.t. s and using the same decomposition as in (A.2) of the main proof:

$$\frac{d\mathcal{L}}{ds} = \underbrace{D}_{\text{direct, choices fixed}} + \sum_i \sum_{a\in\{q,m,x\}} \underbrace{\frac{\partial \mathcal{L}}{\partial a_i}}_{\text{FOC block}} \cdot \frac{da_i}{ds}. \tag{14}$$

[23] 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.

To compute the direct contribution of each action to welfare, $\frac{\partial \mathcal{L}}{\partial a_i}$, note that the household solves the *perceived* problem with consumption $\tilde{c}_i$ that replaces the uninsured-loss expectation by $(1-\lambda_i)\mathbb{E}[L_i]$. Thus the household FOC sets

$$\frac{\partial}{\partial a_i}\mathbb{E}_\omega\mathbb{E}_\xi[u(\tilde{c}_i)] = 0.$$

Evaluated at the household's choice, the planner's experienced-utility marginal is

$$\frac{\partial}{\partial a_i}\mathbb{E}_\omega\mathbb{E}_\xi[u(c_i)] = \underbrace{\frac{\partial}{\partial a_i}\mathbb{E}_\omega\mathbb{E}_\xi[u(\tilde{c}_i)]}_{=0} + \mathbb{E}_\omega\mathbb{E}_\xi\big[u_i' \cdot \partial(c_i - \tilde{c}_i)/\partial a_i\big].$$

Begining with the coverage choice q_i, with misperceptions λ_i, the derivative of the household's perceived consumption is

$$\mathbb{E}_\xi\big[u_i'\,\partial\tilde{c}_i/\partial q_i\big] = \mathbb{E}_\xi[u_i'](-\tilde{P}_i(\omega)) + (1-f)\,(1-\lambda_i)\,\mathbb{E}_\xi[u_i']\,\mathbb{E}_\xi[L_i \mid \omega].$$

Hence, using actuarial fairness, the coverage internality wedge is

$$\Delta_i^q = \mathbb{E}_\xi\big[u_i'\,\partial(c_i - \tilde{c}_i)/\partial q_i\big] = (1-f)\,\lambda_i\,\mathbb{E}_\xi[u_i']\,P_i(\omega).$$

Similarly, the mitigation mitigation wedge is $\Delta_i^m = \lambda_i(1-f)(1-q_i)\,\mathbb{E}_\omega\Big[\mathbb{E}_\xi u_i'\Big(-\frac{\partial P_i}{\partial m_i}(\omega)\Big)\Big]$ and the location internality wedge is $\Delta_i^x = \lambda_i(1-f)(1-q_i)\,\mathbb{E}_\omega\Big[\mathbb{E}_\xi u_i'\Big(-\frac{\partial P_i}{\partial x_i}(\omega)\Big)\Big]$.

Thus, translating the derivatives $\frac{da_i}{ds}$ into elasticities, we have the additional term in the planner's welfare optimality condition

$$-\frac{1}{1-s}\,\mathbb{E}_\omega\left[\sum_i \mathbb{E}_\xi u_i'(\omega,\xi)\left(\Delta_i^q\,\mathcal{E}_i^q + \Delta_i^m\,\mathcal{E}_i^m + \Delta_i^x\,\mathcal{E}_i^x\right)\right],$$

as required.