NBER WORKING PAPER SERIES

OPTIMAL FLOOD INSURANCE IN A SECOND-BEST WORLD: FISCAL SPILLOVERS, RECLASSIFICATION RISK AND MORAL HAZARD

Jonathan Gruber
Adam Solomon

Working Paper 35408
http://www.nber.org/papers/w35408

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2026

We thank Eileen Sadati for excellent research assistance. We are grateful to Ben Collier, Rebecca Diamond, Cameron Ellis, Kerry Emanuel, Nathan Hendren, Larry Katz, Philip Mulder, Jim Poterba, Anna Russo, Ali Sarhadi & Pari Sastry for helpful comments, Ana-Maria Tenekedjieva for a great discussion, and to Katherine Wagner, Hannah Farkas and Seung Min Kim for generously sharing data. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Optimal Flood Insurance in a Second-Best World: Fiscal Spillovers, Reclassification Risk and Moral Hazard
Jonathan Gruber and Adam Solomon
NBER Working Paper No. 35408
July 2026
JEL No. H22, H23

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 52%, comparable to the level of subsidization before the recent reform.

Jonathan Gruber
Massachusetts Institute of Technology
Department of Economics
and NBER
gruberj@mit.edu

Adam Solomon
New York University
Leonard N. Stern School of Business
Department of Finance
adam.solomon@nyu.edu

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 \$173-\$216 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 — hurricane hazards from Gori et al. (2022) and inland (fluvial and pluvial) hazards from the Inter-Sectoral Impact Model Intercomparison Project (ISIMIP2b; ISIMIP 2024). 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 across the combination of eight hurricane and five inland climate-model projections under each of two aggregate warming scenarios (SSP2-4.5 and SSP5-8.5) chosen to plausibly bracket the trajectory the world is currently on. 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. 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'). The idiosyncratic component - the diversifiable reclassification risk a subsidy can insure - has a standard deviation of $1,842 per household. We compute households' WTP for insurance only against this idiosyncratic reclassification risk. We find that a 50% premium subsidy provides approximately $207 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 0.77% 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 7 to 11% more likely to elevate than those exposed to a lower premium jump. Since Florida's housing stock is unusually amenable to elevation, the implied national mitigation-elasticity is about 22% as large as the Florida-estimated baseline.

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, 52%, is similar to the pre-RR2.0 level of subsidization.

Even more starkly, we estimate that a premium subsidy up to 8% has infinite marginal value of public funds (MVPF); it pays for itself. Importantly, our headline results assume FEMA generosity and spillovers are fixed, and that households make rational, fully-informed flood insurance decisions. Were FEMA to cover less ex post damage as it has historically, this would weaken the case for ex ante subsidies; conversely, allowing for documented under-perception of flood risk or over-perception of FEMA generosity strengthens it substantially.

We study how the sufficient statistics, and hence the optimal subsidy, vary with household characteristics: the optimal subsidy decreases in income because larger FEMA spillovers in lower-income tracts more than offset their stronger mitigation responses to insurance prices, whereas it is much lower (19%) for non-primary residences as they are not eligible for FEMA aid. 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 at least 14%.

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 nor that construction meets minimum flood-proof standards[3], 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.

We emphasize that our subsidy prescription is a market-wide one: a subsidy on the annual flood-insurance premium that applies to NFIP and private policies alike, not a differential subsidy

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

[3]To be precise, there is no mandate at the federal level. However, some communities, particularly coastal and barrier-island ones, voluntarily adopt floodplain construction standards extending into non-SFHA zones via local ordinance.

to the NFIP. The historical concentration of subsidies in the NFIP reflects the program's status as the sole flood-insurance provider for most of its history, not what we recommend or evaluate. Nevertheless, Appendix B.24 quantifies the welfare loss from an NFIP-only subsidy that crowds out a nascent private market.

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[4] 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 on the government through Fannie Mae or Freddie Mac.

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[5] 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. Households 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

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

[5]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)).

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 15%, even absent any fiscal spillovers.

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 $m_i \in [0,\bar{m})$.[6] 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.[7]

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.

[6]Households make these choices once, in expectation over the climate state. This rules out additional self-protection in response to a bad realized future risk. Appendix B.20 relaxes this assumption and studies the optimal subsidy when households respond to future flood-risk realizations with additional mitigation and relocation. We show there that the optimal subsidy falls only modestly under this dynamic extrapolation, from 52% to 47–48%, a 4–5 percentage-point decline.

[7]In writing $\mathbb{E}_\omega$ throughout, we are assuming households' beliefs over the future climate state are correct. That is, they share the analyst's prior over the climate-projection ensemble underlying Section 6. This is distinct from households' beliefs over their level of risk and the generosity of FEMA. Those two misperceptions are studied theoretically in Appendix B.22 and empirically in Section 9.6.2.

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: $\bar{P}_i = E_\omega\left[P_i(\omega)\right]$.

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 the income gradient in optimal subsidies is modest: low-income tracts have somewhat higher optimal subsidies, implying an undifferentiated subsidy is mildly regressive.

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 realization 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.[8] 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.

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)[9], we arrive at the simplified optimality condition that we will estimate.

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

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

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

solves:

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

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

The left-hand side measures the benefits of an expanded subsidy.[10] 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 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

[10]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.

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 (2024a) 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.[11] 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 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

[11] 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.21, 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.

the incident type, date, and geographic scope. For our analysis, we restrict to the counties where assistance under the IHP was authorized.[12][13]

To estimate disaster-specific damages, we use the National Oceanic and Atmospheric Administration (NOAA) Storm Events Database. We extract all events classified as floods – hurricane-driven flooding (storm surge and rainfall), fluvial (riverine) flooding, and pluvial (surface-water) flooding – and aggregate the total estimated property damage for each event in each county. These county-level damage figures are then matched to a PDD by county and date, with a one-week buffer.[14]

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 about \$4.9 billion.[15] About one in five households have

[12]A concern is that IHP is more likely to be authorized in areas with lower NFIP coverage. Appendix B.25 shows this is not the case.

[13]NFIP claim payouts also flow to county-disasters outside this sample — approximately 7% of total payouts go to events without a presidential disaster declaration. Appendix B.26 characterizes this distribution.

[14]One concern with the Storm Events Database is that reporting is incomplete (Gallagher, 2023). Appendix B.10 re-runs the spillover IV using two alternative damage measures whose construction has no insurance-claim or local-reporting component: a hydrological model of flood intensity (Farkas and Kim, 2025) and FEMA's realized public infrastructure damage repair costs. The spillover coefficient is robust across all specifications.

[15]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 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.

NFIP insurance; the damage done to a house by flooding averages $29,000. On average, across all houses in the county, NFIP payouts are $3,562. Given that only 18% of houses are insured, if insured houses are representative, this suggests that NFIP offsets 68% 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 $736 per house[16], while on average another $967 is provided in SBA loans, and $152 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 $996 on a flood-weighted basis. Adding this to the $3,562 in NFIP payouts, we find ex ante and ex post payments amount to 16% of total flood costs - meaning that about 84% 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[17] and estimating a distinct quadratic function within each bin.[18] 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.

[16]Our measure of uninsured public expenditure captures FEMA's Individuals and Households Program (IHP) grants only, excluding FEMA Public Assistance (PA). PA represents about 90% of Disaster Relief Fund obligations (USAFacts, 2025) and funds repair of public infrastructure (roads, utilities, schools, debris removal) rather than household property damage — a margin not directly affected by NFIP take-up. The remaining about 10% is mostly the household-facing IHP that substitutes for NFIP and which is the focus of our study.

[17]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.

[18]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.

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 $173 reduction in ex post spending, which is highly significant, but is quite small relative to the total costs imposed by flooding. We estimate that 8% of houses make a FEMA claim. Assuming that these are the uninsured houses, and assuming that they also benefit from other spending (e.g., SBA loans), then this amounts to roughly $2,200 per house. If the average damaged house gets this assistance, this is only about one-thirteenth 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. 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.[19]

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 salience 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 $216 reduction, which is statisti-

[19]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.

cally significant at the 1% level. As we decompose in the Appendix Tables A2 - A6, the majority of this – $183 – 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%, broadly consistent with 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 -$177 and an IV of -$236. Moreover, in Appendix B.14, we show that our findings are robust to excluding floods in which the own-county was recently flooded.

Additionally, we test for other threats to our primary IV strategy. We show that our results hold if, instead of a binary instrument for a flood in the last five years, we instead use a continuous measure of flood damage or FEMA payouts (Appendix B.11). There is a concern that counties that have local political representatives or preferences that are more or less aligned with the state could be treated differently. To check this, Appendix B.15 shows that the primary results are robust to a rich panel of political controls including focal- and neighbour-county presidential vote share, congressional representatives, the Governor's political party, and interactions between all these. Also, there is a concern that depletion of state or federal disaster relief funds could cause a mechanical relationship between floods in neighbouring counties and the funds available to assist the primary county. We show our results persist when accounting for this state (Appendix B.13) and federal (Appendix B.12) spillover. We also show in Appendix Table A7 that our primary estimates are robust to a wide variety of damage controls, and in Appendix B.10 that they are robust to entirely different (non-SED) damage measures.

Finally, 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.

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.[20]

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

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

We additionally estimate two versions of (2) with additional heterogeneity terms. One with an interaction between the semi-elasticity and subsidy level, one with an interaction between the semi-elasticity and an indicator for a house being in an SFHA. The results are in Table 3.

Table 3 about here

We estimate a semi-elasticity of demand of -0.32.[21] This is roughly in the middle of the range estimated by the literature. Less elastic estimates include -0.15 to -0.17 (Mulder (2024a)), -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

[20]Appendix B.5 formalizes the identifying assumption as a within-policy conditional-independence condition, discusses three potential endogeneity threats, and reports an event study of pre-2021 renewal rates by future treatment intensity as a check.

[21]Appendix B.4 shows this demand elasticity estimate is essentially unchanged when differentiated by whether a county or its neighbor has had recent-flood exposure.

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.3 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.2 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 (in absolute value), 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.

Several major flood-risk information releases occurred during the RR2.0 sample. First Street Foundation released Flood Factor data in June 2020 (First Street Foundation, 2020), and these were integrated into consumer-facing apps on Realtor.com (August 2020) and Redfin (February 2021) (Realtor.com, 2020; Redfin, 2021). These rollouts were national with no geographic staggering. Therefore, the informational change might affect our within-property elasticity estimate of 0.32 (in absolute value). In a completely different setting, Mulder (2024b) identifies the price elasticity of NFIP demand separately from the informational channel and obtains demand elasticities between 0.25 and 0.29. Our between-county estimate of 0.2, in Appendix B.2, is unaffected by the national informational roll-out. Our robustness analysis (Table 6) covers demand elasticities well below either 0.25 or 0.2.

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.[22]

While aggregate risk is largely unavoidable, this idiosyncratic risk is, in principle, diversifiable across the nation.[23] 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 synthetic flood-hazard projections from Gori et al. (2022) for hurricane hazards and ISIMIP2b for inland flooding, each applying multiple climate models. Each climate model makes a different prediction about flood hazard, fixing the aggregate amount of atmospheric warming. Hence, our measure of flood-risk uncertainty captures between-model differences rather than within-model initial-condition sensitivity. Importantly, even if within-model variation were available,

[22] Premium volatility is economically meaningful over and above actual flood losses. If households drop coverage in response to premium increases, they will be left unprotected with substantial financial (Billings et al., 2022) and health (Deryugina and Molitor, 2020) consequences should a flood occur. If flood insurance must be maintained, Ge et al. (2025) find that rising premiums increase mortgage and credit-card delinquency and induce distressed relocation, with the delinquency burden falling disproportionately on financially constrained, high debt-to-income households. It is precisely this exposure to costly, involuntary adjustment as flood insurance premiums double over the next decades, that makes insurance against reclassification risk valuable.

[23] Products that insured the aggregate component of climate risk would face the same problem as, for example, long-term care insurance (LTCI): aggregate exposure cannot be hedged, and issuers can end up insolvent or having to retroactively raise premiums (Solomon, 2024). The closest theoretical analog - guaranteed-renewable health insurance with front-loaded premiums (Pauly et al., 1995; Cochrane, 1995) - has never been widely deployed; the closest empirical analog (LTCI itself) failed in the way described. Even a product designed to insure only the idiosyncratic component carries aggregate model risk over a 50-year horizon that insurers may not wish to take on. If issued, the product would have to move money toward households whose realized flood risk — residualized on the climate scenario and house fixed effects — turns out to be high and away from those whose residual is low. This is extremely hard to explain: many households who experience flood-risk increases driven by aggregate-scenario or household-level shifts would receive no payout, and might not be happy about it. Long-term rate locks and premium guarantees face the same constraint: to remain solvent they too must hedge only the idiosyncratic component, since any aggregate exposure leads to LTCI-style problems. In either case, the product is highly unfamiliar. We think a subsidy that implicitly provides this insurance, without any household having to actually understand what it is they are buying, is far more plausible.

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 these projections 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 projections of future hurricane-driven and pluvial/fluvial flood hazards. Data on future hurricane-driven flooding comes from 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) develops projections of future hurricane risk 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. We complement these hurricane-driven hazards with inland (fluvial and pluvial) flood hazards from ISIMIP2b, each likewise estimated across multiple CMIP models.

To ensure a consistent basis for comparison, each climate model is applied to the same aggregate climate scenario. We use SSP2-4.5 as our primary scenario, but we also report results under SSP5-8.5.[24] These plausibly upper- and lower-bound the trajectory that global warming is currently on.[25] The optimal subsidy is only mildly sensitive to this choice: the reclassification-risk insurance value is about 10% lower under SSP2-4.5 than SSP5-8.5 (Table 5), which moves the headline optimal subsidy by less than a percentage point.

This procedure produces, for each model and for each location, measures of future hurricane-driven and fluvial/pluvial flood hazards. 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.

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.

[24] ISIMIP2b simulates fluvial and pluvial flooding hazards under both SSP2-4.5 and SSP5-8.5. Gori et al. (2022) develop hurricane predictions only for SSP5-8.5 directly. We scale these to SSP2-4.5 using simulated differences in aggregate hurricane frequency and severity from Xi et al. (2023).

[25] 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. SSP2-4.5 is a middle-of-the-road scenario with 4.5 W/m^2 radiative forcing. The best estimates of 2100 warming relative to 1850-1900 are about 4.4 degrees Celsius under SSP5-8.5 and 2.7 degrees Celsius under SSP2-4.5 (IPCC (2021)). Current emissions trajectories under existing policies and pledges sit between these two scenarios, closer to SSP2-4.5 (Hausfather and Peters (2020)).

We estimate, for property i's RR2.0 full-risk premium,

$$\text{Full Risk Premium}_i = \beta(\text{Hurricane Hazards, Fluvial Hazards, Pluvial Hazards})_i + \gamma(\text{Building Controls})_i + \epsilon_i$$

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 hazard predictors comprise three hurricane variables from Gori et al. (2022) (100-year rainfall depth, 100-year storm tide level, and joint exceedance probability) and eight ISIMIP2b inland hazards: annual-maximum mean, maximum, and minimum daily river discharge; annual-maximum total, surface, and subsurface runoff; and annual-maximum rainfall and snowfall flux. We fit the mapping by Lasso with linear and squared hazard terms, tuned by 5-fold cross-validation. Because future hazards are much larger than current ones, the Lasso's terms would otherwise extrapolate well outside their training range. We therefore cap each hazard input at $1.25\times$ the training 99.9^{th} percentile before predicting future premiums. The fit yields 63 non-zero coefficients. We report sensitivity to the model (OLS vs LASSO), censoring level, interaction terms, and squared-term choices in Appendix B.17. In general, our choices are conservative in terms of the implied insurance value against reclassification risk.

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. 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, 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 - approximately 50% - 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 \$944. Panel B shows the effect of RR2.0 on premiums (they rise to \$1,739 on average) and the risk factors estimated for each property using current hazard levels. Our regression above is estimated on the premiums and risk factors in Panel B. Panel C then uses simulated future hurricane hazards from Gori et al. (2022) and inland fluvial/pluvial hazards from ISIMIP2b, and the model trained on Panel B to project future premiums across 40 scenarios pairing 8 Gori GCMs with 5 ISIMIP global hydrological model (GHM)–GCM combinations. Under our base-case SSP2-4.5 scenario (shown in the left-hand columns of Panels C and D), the average future premium is $3,458 in 2070 (denominated in 2023 dollars). Under the higher-warming SSP5-8.5 scenario, rainfall hazard increases by a factor of about 4.7, 100-year storm tide level by about 48%, and joint storm–rainfall probability by roughly a factor of 19, and the analogous future premium is $3,531. While hurricane-driven flooding is substantially lower under SSP2-4.5 than 5-8.5, fluvial and pluvial hazards are similar or occasionally higher. This pattern is in the raw data, unrelated to our modeling choices, and comes from the well-known non-monotonicity of inland flooding with respect to overall temperature (Sharma et al. (2018), Wasko and Sharma (2017)).

Table 4 about here

However, future premium uncertainty is substantial even under common aggregate warming assumptions. Projected future premiums are highly dispersed across households, with a standard deviation of $\$6{,}201$ (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 $\$1{,}842$. 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.

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 50%. [26]

Recall from Section 3 the decomposition $P_i(\omega) = \bar{P}(\omega) + a_i + \varepsilon_{i\omega}$, with $\mathbb{E}_\omega[\varepsilon_{i\omega}] = 0$, and that of these three terms only $\varepsilon_{i\omega}$ represents risk that an NFIP subsidy can insure in the cross-section. An NFIP subsidy could in principle also transfer aggregate climate risk ($\bar{P}(\omega)$) to the taxpayer base, but we do not treat that transfer as insurance in our welfare model. Empirically we construct $\varepsilon_{i\omega}$ by demeaning the policy–scenario matrix of projected premiums along both margins:

$$\varepsilon_{i\omega} \;=\; P_i(\omega) - \bar{P}(\omega) - \bar{P}_i + \bar{P},$$

where $\bar{P}_i = \mathbb{E}_\omega[P_i(\omega)]$ is household i's expected future premium and $\bar{P} = \mathbb{E}_\omega[\bar{P}(\omega)]$ is the unconditional mean future premium across both households and scenarios. Subtracting both marginal means removes $\bar{P}$ twice—each marginal already contains a full copy of $\bar{P}$—so we add a single

[26] Per Table 4, the mean pre-RR2.0 premium is $944, the mean post-RR2.0 premium is $1,739.

copy back to ensure $\varepsilon_{i\omega}$ is mean-zero over ω for every household by construction. Panel D of Table 4 summarizes each component.

A proportional subsidy s scales every premium realization by $(1-s)$ and therefore scales $\varepsilon_{i\omega}$ by the same factor. Assuming CARA utility $u(c) = -\exp(-\gamma c)$, the pure insurance value at subsidy s for household i is

$$V_{\text{reclass},i}(s) \ = \ \text{CE}_\omega(\varepsilon_{i\omega}) - \text{CE}_\omega((1-s)\,\varepsilon_{i\omega})\,,$$

where $\text{CE}_\omega(X) = \frac{1}{\gamma}\log \mathbb{E}_\omega[\exp(\gamma X)]$.[27] Because $\varepsilon_{i\omega}$ is mean-zero by construction, $V_{\text{reclass},i}(s)$ is a pure risk premium with no level transfer mixed in. Table 5 reports the unweighted mean of $V_{\text{reclass},i}(0.5)$ across policies, for five values of γ and under both SSP2-4.5 and SSP5-8.5; the cross-sectional mean is the $V_{\text{reclass}}(s)$ that enters the MVPF formula in Section 9.

Table 5 about here

Table 5 indicates that RR2.0 exposed households to substantial reclassification risk.[28] Under our central risk-aversion parameter $\gamma = 5 \times 10^{-4}$, the level shift reduces insurance value by $207, about 22% of the subsidized pre-RR2.0 premium, 12% of the RR2.0 premium and 6% of the projected future premium. Under the higher-warming SSP5-8.5 scenario, the corresponding insurance value is $229.

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 15% is optimal solely in response to 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.

[27] We use $\gamma = 5\times10^{-5}$ as a lower bound estimated by Cohen and Einav (2007); $\gamma = 5\times10^{-4}$ as our primary estimate; and $\gamma = 5\times10^{-3}$ as an upper bound from Sydnor (2010) and Barseghyan et al. (2011).

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

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 $\$1{,}739$.

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.

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.[29]

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.

Premium changes under RR2.0 were overwhelmingly increases; decreases were uncommon and concentrated in pre-1970 housing stock. Appendix B.7 reports the full incidence by vintage. The mitigation response we observe, elevating a house at construction, therefore comes from those

[29]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.

that experienced a price rise under RR2.0; there is no inverse group who mitigated less because of a price decline.

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.

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.8. 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 3.

Figure 3 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 7 to 11%. 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 0.38. Outside the SFHA, the difference in price increases was \$373, against a mean full-risk price of \$1,518, implying an elasticity of 0.30.

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.[30]

Florida-to-U.S. Mitigation-Elasticity Extrapolation Our mitigation elasticity E_m is identified from Florida elevation certificates, and Florida differs from the rest of the country on several community-level characteristics that plausibly interact with the ease of elevation. For example, Florida's housing stock is younger and more heavily manufactured - characteristics that likely make Florida more amenable to mitigation than the rest of the country. We re-estimate E_m within Florida as a function of these characteristics and reweight to the national covariate distribution; under this extrapolation, the optimal subsidy we calibrate in Section 9 rises from 52% to 65%.[31] Appendix B.19 reports the details and estimates of this procedure.

[30]This elasticity is estimated from short-run RR2.0 price variation; Appendix B.20 reports robustness under a long-run extrapolation in which households respond to 2070 climate-driven price changes with the same elasticity.

[31]Generally, mitigation responses to insurance price changes appear to be small. Premium discounts produce no detectable adoption response in coastal hurricane-retrofit choices (Jasour et al., 2020), household wind insurance and mitigation choices appear to be jointly driven by risk perception rather than by premium incentives (Petrolia et al., 2015), and in the context of wildfires, direct mitigation subsidies do not clearly induce more uptake (Baylis et al., 2024).

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

Figure 4 about here

Figure 4 shows a net movement of about 0.77% of the population from census tracts treated with a large price increase under RR2.0 to those with smaller (or negative) price changes.[32] 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 0.77% 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.[33]

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). Relatedly, Qin and Voorheis (2026) find evidence using census microdata that net migration from flood-plain to non-flood-plain areas increased by approximately one percent from 2020 to 2023.

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 $911 represents 3.8% of annual U.S. rent ($23,940 per Zillow Group, Inc. (2026)). As a result, our migration response of 0.77% 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

[32] 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 0.77% change in occupied housing even without any supply side response that we might expect in the longer term.

[33] Housing stock turns over slowly, so the population response we observe within four years of RR2.0 may understate the long-run response to a permanent price change; Appendix B.20 quantifies how the optimal subsidy changes under a long-run extrapolation in which households respond to realized future climate states with further mitigation and relocation.

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 4 that measures the change in FEMA spending with respect to coverage increases *conditional on a flood*. The term in the equation is unconditional. Hence, we need to estimate the probability of a flood. We take the expected damage per house *given a flood* to be $L_f = \$29{,}267$ (see Table 1). Under actuarial fairness, an average premium $\bar{P}$ implies a flood probability $p_f = \sigma_{\text{PDD}}\,\bar{P}/L_f$ each year, where $\sigma_{\text{PDD}} = 0.93$ is the PDD share of NFIP-insured damages.[34]

Therefore, the change in *unconditional* expected FEMA spending per house for a 1 percentage point (0.01) increase in coverage is $\beta_{\text{uncond}} = \beta_{\text{cond}} \cdot p_f = (-\$215.70) \cdot \sigma_{\text{PDD}} \frac{\bar{P}}{L_f}$. 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 $\Delta\text{FEMA}(s) = (\beta_{\text{cond}})\, \sigma_{\text{PDD}} \frac{\bar{P}}{L_f} \cdot 100 \cdot \eta^q \cdot \Delta \ln \tilde{P}$.

For example, plugging in $\Delta \ln \tilde{P} = 0.01$, $L_f = 29{,}267$, the average full-risk premium $\bar{P} = \$1{,}739$, and $|\eta^q| = 0.32$, yields $\Delta\text{FEMA}(0.01) \approx \$3.84$ per house-year in response to a 1% price change.[35]

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 $\$9.28$ per insured household. At $s = 0.5$ this falls to $\$2.32$, 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 $s = 0.5$.

[34] Because L_f is calibrated on damage conditional on a presidentially-declared disaster (PDD), whereas $\bar{P}$ is the actuarial premium covering all flood losses, we scale $\bar{P}$ by 0.93 – the PDD share of total NFIP-insured damages – so that the implied flood probability is on the same PDD support as L_f. Appendix Table A34 documents this share: presidentially-declared disasters account for 93% of NFIP-insured claim dollars.

[35] In our primary calculations, we abstract from timing differences between NFIP and FEMA. The former are paid faster, which pushes in favor of a higher subsidy. We estimate that an NFIP dollar is worth up to 30% more than a FEMA dollar given a payout-delay differential of 16 weeks (Federal Emergency Management Agency, 2024; Kousky, 2023) and a household shadow value of liquidity of 100% APR (Bhutta et al., 2015; Kaplan et al., 2014). This would increase the optimal subsidy to $s^* = 64.6\%$.

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 $\Delta \ln \tilde{P} = 0.01$, $\bar{P} = \$1{,}739$, and $\eta^q = 0.32$, this coverage externality is zero at $s = 0$, about \$2.80 per house-year at $s = 1/2$, and about \$5.61 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. Note that the higher NFIP claim payouts from induced enrollment do not themselves enter as a fiscal cost: under actuarial fairness, each induced policy's expected claims are balanced by its gross-of-subsidy premium. The only fiscal externality is the subsidized portion of the premium that households do not pay.

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.[36] 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

[36] 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.

semi-elasticities of 0.38 (SFHA) and 0.30 (non-SFHA); our calculations use an average of 0.34.

In summary, 1% subsidy yields a mitigation fiscal externality of $s \times 0.4 \times 0.34 \times \bar{P} \times 0.01 = 0.00136 \times s \times \bar{P}$ At $\bar{P} = \$1{,}739$, this is equal to $\$1.18$ 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 is the terminal coefficient of this event study: high-premium tracts experience a net out-migration of -0.77%, so $\eta^x \approx -0.0077$. This gives a location fiscal externality of

$$\Delta\text{Location FE} = s \times 0.0077 \times \Delta P \times \Delta \ln \tilde{P}.$$

For a 1% log price cut, the location fiscal externality is zero at $s = 0$, $\$0.041$ at $s = 1/2$, and $\$0.08$ at $s = 1$. This is an order of magnitude smaller than the fiscal externality from mitigation since the underlying elasticity (0.0077 vs 0.34) 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 compute the optimal s^* at the central calibration; Table 6 reports the single-channel robustness of s^* across the policy levers we consider (misperception, demand semi-elasticity, FEMA generosity, and a compounded worst case).

At our central estimates – $\eta^q \approx 0.32$ and $\bar{P} \approx \$1{,}739$ – the optimal subsidy is approximately **52%**.

9.6.1 Heterogeneity in the Optimal Subsidy.

Our optimal subsidy is evaluated at overall average levels of the five sufficient statistics. To probe how the optimum varies across observable household and tract characteristics, we re-estimate how each sufficient statistic varies above and below the median value for each observable. We adjust the sufficient statistic and recompute the optimal subsidy along an observable only when the heterogeneity (i.e. the indicator variable for being above-median) is statistically different from zero at the 90% confidence level. Figure 5 displays the resulting heterogeneous optimal subsidies; the underlying regressions are in Appendix B.9.

Figure 5 about here

The strongest gradient is on primary residence: non-primary residences have an optimal subsidy of $s^* = 19\%$ vs $s^* = 53\%$ for primary residences, driven by the fact that non-primary residences are not eligible for FEMA Individual Assistance and therefore have no ex post spillover. Tract-income is also important - $s^* = 61\%$ in below-median income tracts vs $s^* = 43\%$ above - because larger FEMA spillovers in low-income areas more than offset their stronger mitigation responses. Other dimensions (tract tenure, age, building age, mobility, historical claims) do not move the optimal subsidy substantially.

9.6.2 Joint Robustness Across Policy Levers.

Table 6 illustrates the robustness of s^* as policy settings and calibration assumptions vary. Panel A varies behavioral misperceptions; Panel B varies the demand semi-elasticity η; Panel C varies counterfactual FEMA generosity λ_β; and Panel D compounds a sequence of adverse assumptions to produce a worst-case lower bound.

Table 6 about here

Behavioral Misperceptions. As discussed in Section 3.1, there is substantial evidence that households underperceive their flood risk and also overperceive the generosity of FEMA disaster aid, both of which lower demand for flood insurance below its welfare-relevant value. We assume the household buys insurance according to its (downwardly biased) perceived WTP, whereas the planner evaluates welfare according to the household's unbiased WTP under correct beliefs. This internality makes a premium subsidy more attractive: the marginal household induced into flood insurance has a first-order welfare gain (the difference between welfare-relevant and decision-relevant WTP) in addition to the costs and benefits modeled throughout.

Following Bakkensen and Barrage (2022), we set the underperception of flood risk at 43% and assume households overperceive FEMA generosity by 50%; full details of how we calculate biased and unbiased WTP are in Appendix B.22. Panel A reports the resulting optimal subsidies.

Incorporating risk misperception alone shifts the optimal subsidy from our baseline of 52% to 89% at $\gamma = 5 \times 10^{-5}$ and to 100% at $\gamma = 5 \times 10^{-4}$; layering FEMA misperception on top adds a further 4–7 percentage points. Appendix B.22 reports the full robustness table including the cross-term between the two wedges. In general, once misperceptions are accounted for, the optimal subsidy sharply rises.

Demand Robustness. Panel B varies the demand semi-elasticity η. The first alternative, $\eta = 0.25$, reflects an SFHA-style imperfectly enforced mandate per Table 3.[37] $\eta = 0.15$ is the lower bound from Mulder (2024a); $\eta = 0.059$ is the full price-path anticipation rescale (Appendix B.6). Across this $5\times$ range, s^* moves only from 52% to 39%. This insensitivity arises because the primary benefit (FEMA savings) and primary cost (ΔCoverage FE) are both linear in η.

FEMA Generosity. Panel C varies the counterfactual FEMA generosity λ_β. If FEMA were to be half as generous going forward as it was historically ($\lambda_\beta = 0.5$), the optimal subsidy falls by 19 percentage points, from 52% to 33%; eliminating FEMA entirely ($\lambda_\beta = 0$) lowers s^* further to 15%. This is the policy variable that has the most downward impact on s^*. Note, however, that Panel C should not be read as a recommendation to "solve" the Samaritan's-dilemma channel by simultaneously cutting FEMA generosity and lowering the insurance subsidy. That joint change is welfare-improving only if households are also making rational risk and FEMA-generosity decisions; with realistic misperception (Panel A), the optimal subsidy remains very high even when FEMA is removed. Rather, should FEMA become less generous over time — due to legislative changes or capped funding — this would have implications for the spillover of ex-ante insurance subsidies.

Worst Case. Panel D layers on adverse assumptions one row at a time. Starting from the 52% baseline, we first rescale the mitigation elasticity upward to $E_m = 0.62$ (the no-anticipation upper bound from Appendix B.6), lowering s^* to 44%.[38] Adding a relocation response nine times larger than baseline (calibrated to the relocation effect of Hurricane Maria, (Aron-Dine, 2025)) lowers it to 41%. The next adverse assumption is administrative costs: NFIP pays approximately 30% of its premium income as administrative expenses, primarily broker commissions,[39] while FEMA administrative costs are approximately 13% of assistance delivered,[40] so an increased subsidy pushes households from a low-overhead program (FEMA) to a higher-overhead one (NFIP). We treat administrative spending as a fully wasted net government-budget cost,[41] lowering s^* to 32%.

The upper-bound private-market crowdout from Appendix B.24 ($\theta = 0.20$, $\text{WTP}^{\text{cap}} = \$1{,}022$) lowers s^* further to 27%; an SFHA-restricted demand elasticity ($\eta = 0.25$) lowers it to 26%; and

[37]Within the SFHA, we have assumed the returns to marginal mitigation are the same as outside. This is likely an overestimate since elevation up to the BFE is mandated there, so a full accounting would push in favor of a higher subsidy.

[38]In the opposite direction, the Florida-to-U.S. mitigation-elasticity extrapolation discussed in Section 7 implies $E_m \approx 0.075$ and lifts the optimal subsidy to 65%; Appendix B.19 reports details.

[39]U.S. Government Accountability Office (2023)

[40]U.S. Government Accountability Office (2014)

[41]Administrative spending can be split into a real-resource-cost share (genuine social waste) and a transfer share. Broker commissions, for example, simply transfer dollars from the government to brokers. If we assumed that only half of the administrative costs were wasted resources, and the other half were transfers, and that there were no redistributive reasons to favor NFIP/FEMA recipients over, say, insurance brokers, the optimal subsidy with admin would be 37% rather than 32%. We adopt the more conservative full-wasted-resources assumption as our headline.

halving FEMA generosity ($\lambda_\beta = 0.5$) lowers it to 14%. Even under these highly conservative assumptions — and, critically, without accounting for behavioral frictions — there is a case for a non-zero subsidy.

9.6.3 Marginal Value of Public Funds.

Because we assumed that premium subsidies were funded with lump-sum taxes, the optimal subsidy (52% 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 $MVPF = \frac{\text{Mechanical Benefit}+\text{Insurance Value}}{\text{Mechanical Cost}+\text{Fiscal Externality}}$. In our setting, the precise components are shown below.

$$\text{MVPF}(s) = \frac{\overbrace{q\,\bar{P}}^{\text{Mechanical Benefit}} + \overbrace{q \cdot V_{\text{reclass}}(s)}^{\text{Insurance Value}}}{\underbrace{q\,\bar{P}}_{\substack{\text{Mechanical}\\\text{Cost}}} \underbrace{- \eta^q\,|\beta_{\text{uncond}}|}_{\substack{\text{FEMA}\\\text{spillover (saving)}}} + \underbrace{\eta^q\,\bar{P}\cdot s}_{\substack{\text{Coverage}\\\text{FE}}} + \underbrace{\bar{P}\,E_m\,\Delta P_{\text{haz}}\cdot s}_{\substack{\text{Mitigation MH}\\\text{(foregone)}}} + \underbrace{\Delta P_{\text{move}}\,E_x\cdot s}_{\substack{\text{Location MH}\\\text{(foregone)}}} + \underbrace{\eta^q(\bar{P}\,A_N - A_F\,|\beta_{\text{uncond}}|)}_{\substack{\text{Net admin cost}\\\text{(NFIP minus FEMA)}}}}$$

where $\beta_{\text{uncond}} \equiv 100\,\beta_{\text{cond}} \cdot \sigma_{\text{PDD}}\,\bar{P}/L_f \approx -\$1{,}192$ is the unconditional marginal change in FEMA assistance per house-year per unit (fractional) increase in NFIP coverage, with $\beta_{\text{cond}} = -\$215.70$ the per-percentage-point conditional-on-flood spillover, $\sigma_{\text{PDD}} = 0.93$ the PDD share of NFIP-insured damages, and $L_f = \$29{,}267$ expected damage given a flood (Section 9). The results for the MVPF for each percentage of ad valorem subsidy are illustrated in Figure 6.

Figure 6 about here

For the baseline (black curve), the first 8% of subsidy have infinite MVPF: the fiscal externalities from reduced FEMA expenditure outweigh the mechanical costs and moral hazard. The MVPF then declines, crossing 1 at the optimal subsidy of 52%. Layering in administrative costs (NFIP 30%, FEMA 13%; blue curve) lowers the optimal subsidy to 40%. Assuming the subsidies are funded with the same distortion as the average federal tax dollar (an MCPF of 1.2 from Hendren and Sprung-Keyser (2020)) (upper dashed horizontal line) lowers the optimal subsidy to 34%. The remaining curves each turn off one benefit or cost at a time to illustrate the contribution of each to the headline subsidy: with no moral hazard the optimal subsidy rises to 71%, with no reclassification-risk insurance value it falls to 48%, and with no FEMA spillover to 15%.

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.

10 Conclusion

The ongoing increase in "global weirding"[42] 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 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 $216 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 $207 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 order of magnitude smaller 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 52%, 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

[42] See Friedman (2010)

approach zero.

In our primary analysis, we take FEMA policy as fixed. We show that 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.[43] While that would impact our subsidy prescription, 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.

[43] 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 $*)	4,862	3,200	10,000
Total Population	1,014,521	445,337	4,525,519
Total Housing Units	392,899	175,607	1,714,340
Panel B: Per house			
NFIP Policies per House	0.18	0.16	0.49
NFIP Payouts per House ($)	3,562.45	1,798.16	8,255.15
Total Non-NFIP Spending per House ($)	995.80	723.96	2,582.85
FEMA Assistance Claims per House	0.08	0.06	0.19
Property Damage per House ($)	29,267.06	27,454.25	81,466.06
Disaster Assistance per House ($)	736.33	440.00	1,808.91
Fannie Mae Losses per House ($)	3.19	1.17	6.14
Freddie Mac Losses per House ($)	1.10	0.29	1.96
HMGPP Expenditures per House ($)	151.52	88.94	282.36
Approved SBA Loans per House ($)	967.20	417.31	1,878.52
Fiscal Cost from SBA Loans per House ($)	125.74	54.25	244.21

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

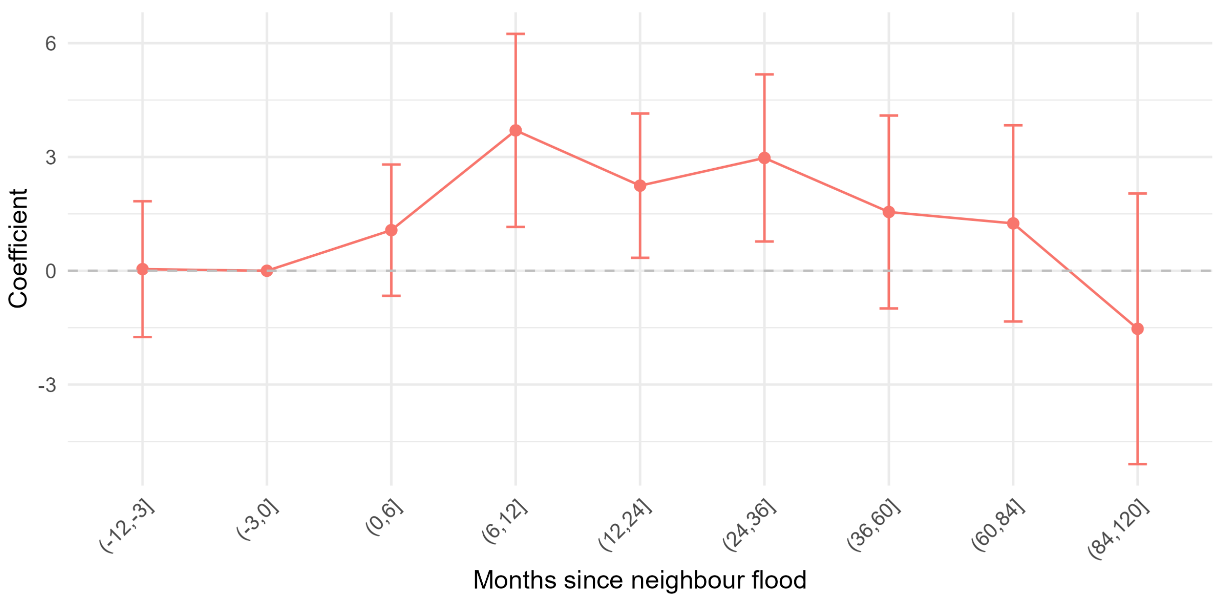

Notes: Event-study estimates from a regression of NFIP coverage per house (measured two months before flood onset) on indicators for years since the most recent flood in the closest neighboring county, controlling for flexible quadratic functions of flood damage within bins (cutoffs at $1, $10, $100, $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

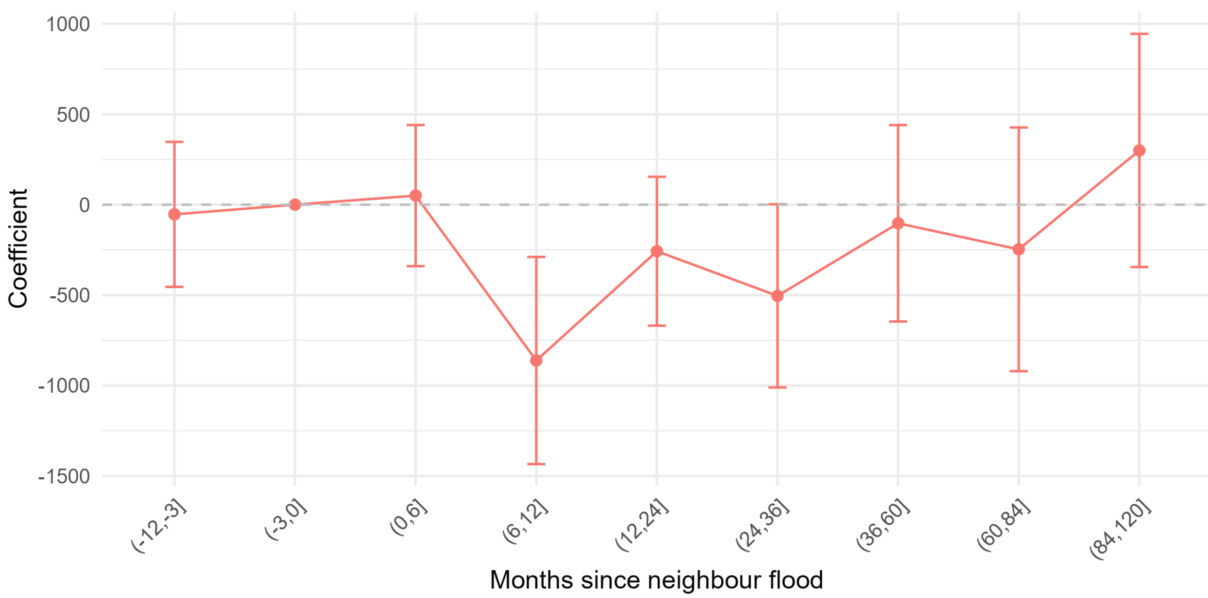

Notes: Event-study estimates from a regression of total ex-post disaster spending per house on indicators for years since the most recent flood in the closest neighboring county, controlling for flexible quadratic functions of flood damage within bins (cutoffs at $1, $10, $100, $1,000, $10,000, $100,000, 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	-172.799^{***}	-215.699^{**}
	(36.492)	(72.811)
First-stage F-statistic		81.98
Num. Obs.	886	886
Between R^2	0.332	0.424

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. 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	2,084,307	2,084,307	2,084,307

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

	SSP2-4.5				SSP5-8.5			
Metric	**Mean**	**SD**	**P10**	**P90**	**Mean**	**SD**	**P10**	**P90**
Panel A: Pre RR2.0 Premiums								
Pre RR2.0 Premium	944	1,193	356	1,799	—	—	—	—
Panel B: RR2.0 Premiums & Current Hazard Levels (IDW)								
RR2.0 Premium	1,739	2,284	391	3,777	—	—	—	—
100-year Rainfall Depth (mm)	176	56	74	236	—	—	—	—
100-year Storm Tide Water Level (mm)	2,473	940	1,204	3,656	—	—	—	—
Joint Probability (%) (x 100)	0.32	0.11	0.16	0.48	—	—	—	—
Max River Discharge (m^3/s)	622	677	75	1,545	—	—	—	—
Mean River Discharge (m^3/s)	134	153	24	278	—	—	—	—
Total Runoff (mm/day)	1.85	1.21	0.49	3.82	—	—	—	—
Surface Runoff (mm/day)	0.80	0.58	0.15	1.70	—	—	—	—
Subsurface Runoff (mm/day)	0.67	0.69	0.02	1.68	—	—	—	—
Rainfall Flux (mm/day)	4.46	1.97	2.28	6.48	—	—	—	—
Panel C: Future Premiums and Hazard Levels								
Future Premium (Scenario Avg)	3,458	6,201	1,366	4,584	3,531	6,374	1,399	4,602
100-year Rainfall Depth (mm)	688	270	238	1,035	831	309	309	1,233
100-year Storm Tide Water Level (mm)	3,503	1,387	1,634	5,434	3,663	1,426	1,740	5,776
Joint Probability (%) (x 100)	4.82	1.71	2.53	7.05	6.17	2.04	3.51	9.05
Max River Discharge (m^3/s)	839	1,109	212	1,626	791	984	214	1,483
Mean River Discharge (m^3/s)	193	323	52	349	178	207	60	332
Total Runoff (mm/day)	2.31	1.19	1.04	3.97	2.27	1.13	1.11	3.93
Surface Runoff (mm/day)	1.83	1.75	0.42	3.85	1.80	1.77	0.38	3.53
Subsurface Runoff (mm/day)	2.26	4.36	0.11	4.53	2.50	3.67	0.13	5.96
Rainfall Flux (mm/day)	6.04	2.18	3.71	9.01	6.25	2.36	3.70	9.31
Panel D: Decomposition Across Climate Scenarios								
Household Fixed Effect	0	6,201	-2,093	1,126	0	6,374	-2,133	1,071
Scenario Fixed Effect (vs current mean)	1,719	262	1,264	2,012	1,792	334	1,275	2,175
Idiosyncratic Risk	0	1,842	-902	956	0	1,978	-960	1,083

Notes: Summary statistics at the property level for $N = 3,529,254$ NFIP policies priced under Risk Rating 2.0 (RR2.0). Panels A–B report historical premiums and current (IDW-interpolated) hazard metrics from Gori et al. (2021) hurricane hazards and ISIMIP2b fluvial/pluvial hazards. Panel C reports projected 2070 premiums and hazards under both SSP2-4.5 (left) and SSP5-8.5 (right). The hurricane component of the SSP2-4.5 column scales the Gori RCP8.5 ratios by Xi–Lin (2023) ratio-of-ratios fold-changes (0.60–0.95 across the coast), and the inland component uses ISIMIP2b RCP6.0 as the CMIP5 analog of SSP2-4.5. Premium projections use a Lasso model with linear hazards, squared hazards, and projection-time censoring at 1.25$\times$ the training 99.9th percentile; see Appendix B.17 for sensitivity to specification choices. 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.

Table 5. Reclassification Risk Results

Coefficient of Absolute Risk Aversion	Reclassification Risk Insurance Value, SSP2-4.5	Reclassification Risk Insurance Value, SSP5-8.5
5×10^{-5}	\$52	\$61
1×10^{-4}	\$82	\$93
2×10^{-4}	\$122	\$136
5×10^{-4}	\$207	\$229
5×10^{-3}	\$625	\$681

Notes: Each row reports the unweighted mean across $N = 3,529,254$ policies of the CARA welfare gain (in 2023 dollars) from an $s = 0.5$ subsidy of the two-way residual $\varepsilon_{i,m} = P_{i,m} - \bar{P}_{i\cdot} - \bar{P}_{\cdot m} + \mu$ (mean-zero per policy by construction). Premium projections use a Lasso model with linear hazards, squared hazards, and projection-time censoring at 1.25$\times$ the training 99.9$^{\text{th}}$ percentile. The SSP2-4.5 column scales the Gori RCP8.5 hurricane hazards by Xi–Lin (2023) ratios of hurricane severity under SSP2-4.5 relative to SSP5-8.5. The CARA values a span literature estimates: 5×10^{-5} (lower bound, Cohen et al. 2007); 5×10^{-4} (central estimate); 5×10^{-3} (upper bound, Sydnor (2010) and Barseghyan et al. (2011)).

Figure 3. Mitigation Changes Under RR2.0

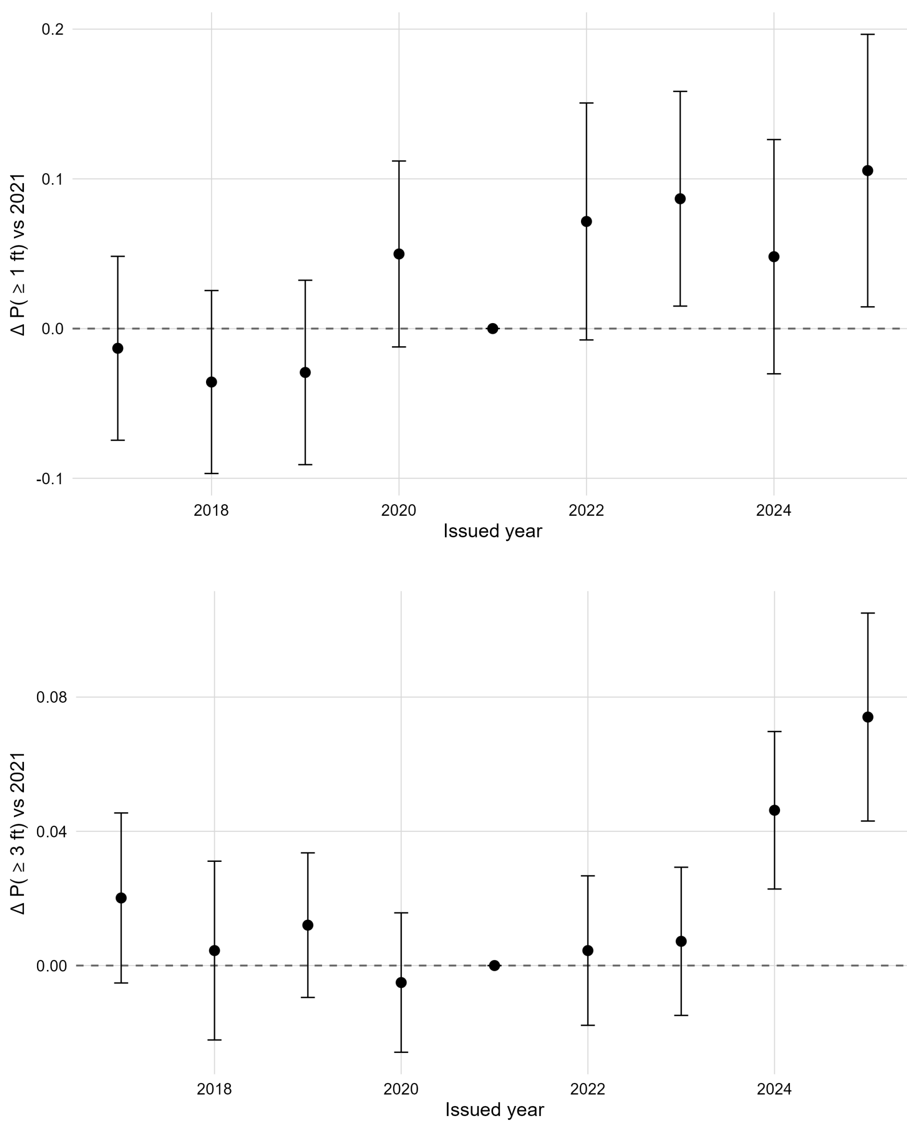

Notes: Event-study estimates from regressions of an indicator for elevating the lowest floor by 1+ feet (relative to base flood elevation inside the SFHA) or 3 feet (relative to local adjacent ground level outside the SFHA) on indicators for years relative to RR2.0 implementation (2022 as reference), 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 2017–2025. The estimating equation is (3).

Figure 4. Location Changes Under RR2.0

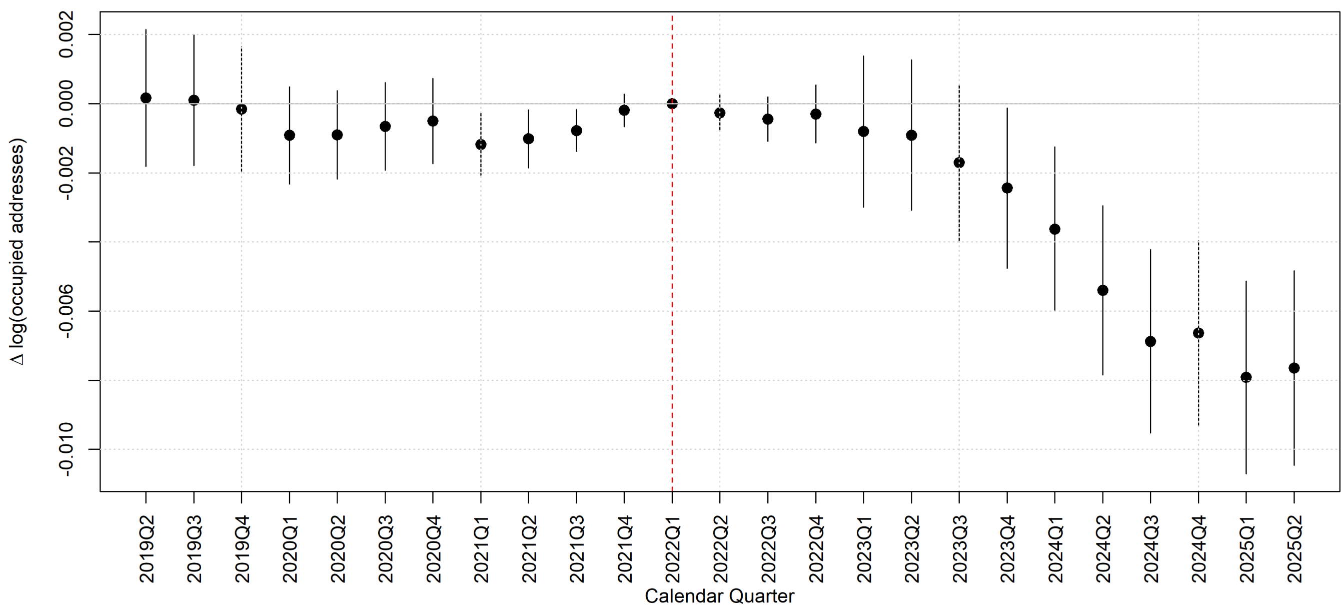

Notes: Event-study estimates of occupied addresses by census-tract $\times$ quarter. Treatment is defined by the median NFIP policy in a census tract experiencing a premium increase under RR2.0. Standard errors are clustered at the census-tract level and 95% confidence intervals are shown.

Figure 5. Optimal Subsidy s^* by Household Characteristics

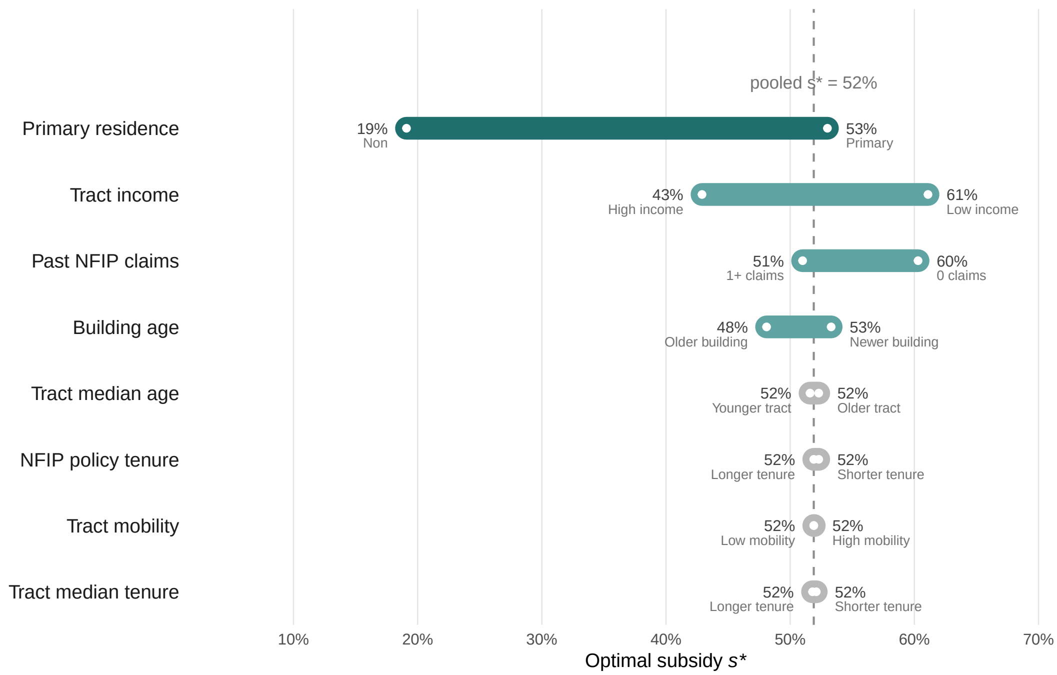

Notes: Endpoints show the optimal subsidy s^* when the sufficient statistics vary based on eight observables. The underlying heterogeneity in the sufficient statistics is reported in Appendix B.9. The dashed line marks the baseline $s^* = 52\%$.

Table 6. Optimal Subsidy – Robustness

	s^*
Panel A: Behavioural Misperceptions Robustness	
Baseline (no internality)	52%
FEMA internality only ($k_F = 1.5$)	100%
Risk internality only ($k_R = 0.57$)	100%
Risk and FEMA internalities	100%
Panel B: Demand Robustness	
Baseline ($\eta = 0.32$, paper-pool RR2.0)	52%
SFHA-restricted mandate enforcement ($\eta = 0.25$)	50%
Literature low ($\eta = 0.15$, Mulder 2024)	45%
Full price-path anticipation ($\eta = 0.059$)	39%
Panel C: FEMA Generosity Robustness	
Baseline ($\lambda_\beta = 1.0$)	52%
Half generosity ($\lambda_\beta = 0.5$)	33%
Zero generosity ($\lambda_\beta = 0.0$)	15%
Panel D: Worst-Case Build-Up (compounding)	
Baseline	52%
+ Mitigation response $\times 1.833$	44%
+ Location response $\times 9$ (Hurricane Maria)	41%
+ Admin costs (NFIP 30%, FEMA 13%, full net-budget rate)	32%
+ Private-market crowdout 2070 upper ($\theta = 0.20$, WTP$^{\text{cap}} = \$1{,}022$)	27%
+ SFHA-restricted demand ($\eta = 0.25$)	26%
+ FEMA generosity halved ($\lambda_\beta = 0.5$)	14%

Notes: Each row reports the optimal subsidy s^* at the paper-baseline calibration ($\beta = -215.70$, $\eta = 0.32$, $\gamma = 5 \times 10^{-4}$, CARA anchor from SSP2-4.5), holding every lever at baseline except the one named in the row. Panel A misperception calibrations: $k_R = 0.57$ from Bakkensen and Barrage (2022); $k_F = 1.5$. Panel B η varies the demand elasticity: $\eta = 0.25$ is the SFHA-restricted estimate (imperfectly enforced mandate); $\eta = 0.15$ is the lower bound from Mulder (2024); $\eta = 0.059$ is the full-anticipation lower bound. Panel D compounds adverse assumptions in the order listed: mitigation response $\times 1.833$ (no-anticipation rescale of E_m), location response $\times 9$ (Hurricane Maria; Aron-Dine 2025), Section 9.4 admin costs (NFIP 30%, FEMA 13%, fully wasted resources), private-market crowdout at 2070-upper calibration ($\theta = 0.20$, WTP$^{\text{cap}} = \$1{,}022$); SFHA-restricted demand ($\eta = 0.25$), and FEMA generosity halved ($\lambda_\beta = 0.5$).

Figure 6. MVPF of a Marginal Subsidy, by Welfare Channel

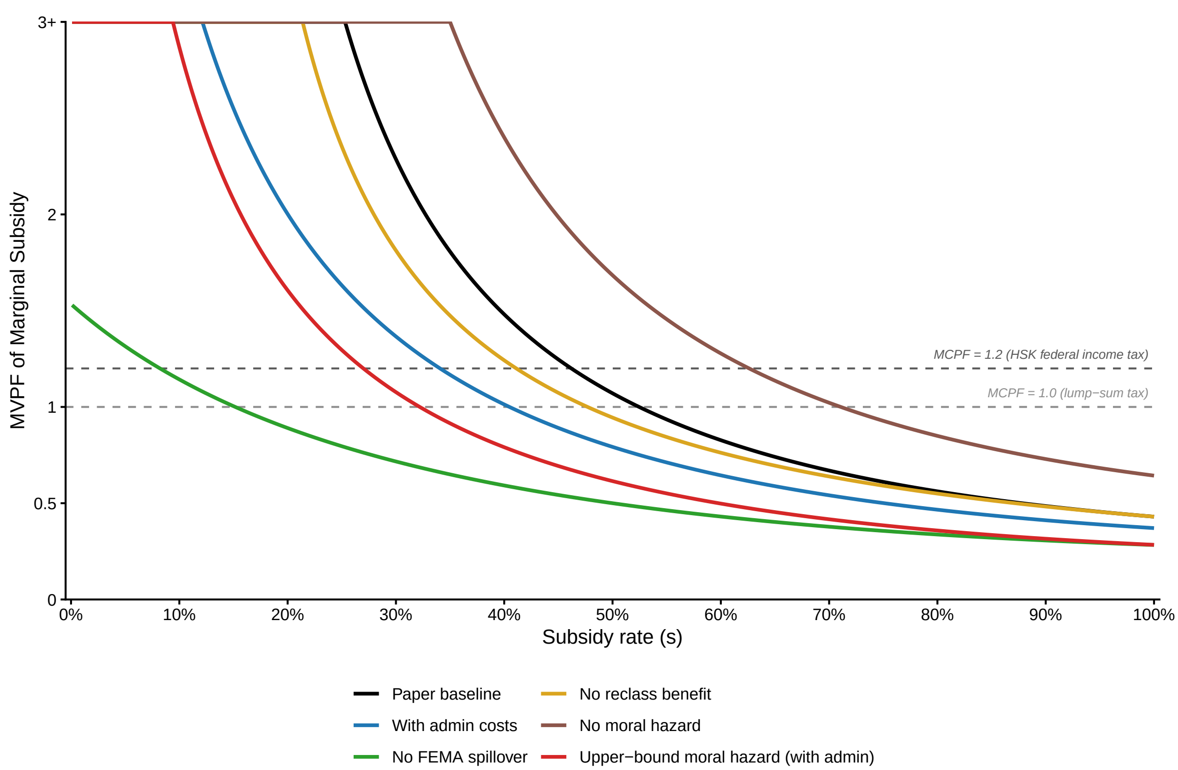

Notes: Marginal value of public funds (MVPF) of a marginal flood-insurance premium subsidy, plotted against the subsidy rate s and decomposed into welfare channels. The black curve uses the baseline assumptions; the blue curve adds administrative costs (NFIP 30%, FEMA 13%) as a net government-budget cost; the remaining curves each turn off one benefit or cost, with the red curve also using upper-bound moral hazard. The horizontal lines mark an MCPF of 1.0 (lump-sum tax) and 1.2 (the Hendren and Sprung-Keyser (2020) estimate for the federal income tax). The MVPF is top-coded at 3 (the "3+" tick); where it is infinite (the FEMA-spillover benefit exceeds all costs) the curve runs flat along this ceiling. For the baseline this holds over the first 8% of subsidy, after which the MVPF becomes finite and declines.

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.

Baylis, P., Boomhower, J., and Horton, B. (2024). Incentivizing self-protection from wildfires. Working paper; AEA RCT Registry AEARCTR-0011188; `https://www.patrickbaylis.com/pdf/dspace-rct.pdf`.

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

Bhutta, N., Skiba, P. M., and Tobacman, J. (2015). Payday loan choices and consequences. *Journal of Money, Credit and Banking*, 47(2-3):223–260.

Billings, S. B., Gallagher, E. A., and Ricketts, L. (2022). Let the rich be flooded: The distribution of financial aid and distress after hurricane harvey. *Journal of Financial Economics*, 146(2):797–819.

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.

Deryugina, T. and Molitor, D. (2020). Does when you die depend on where you live? evidence from hurricane katrina. *American Economic Review*, 110(11):3602–3633.

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

Farkas, H. and Kim, S. M. (2025). Bureaucracy and political bias: Evidence from floods. Working Paper.

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 (2024). National flood insurance program claims handbook. FEMA F-687.

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

Federal Emergency Management Agency, Office of the Chief Financial Officer (2024). Immediate needs funding fact sheet. Under Immediate Needs Funding, FEMA pauses new non-lifesaving Public Assistance and mitigation obligations while Individual Assistance payments to survivors continue.

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.

First Street Foundation (2020). First street foundation releases new data disclosing the flood risk of every home in the contiguous U.S.

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.

Gallagher, J. (2023). Retrospective voting and natural disasters that cause no damage: Accounting for the selective reporting of weather damage. Working Paper.

Ge, S., Johnson, S., and Tzur-Ilan, N. (2025). Climate risk, insurance premiums and the effects on mortgage and credit outcomes. Working Paper 2505, Federal Reserve Bank of Dallas.

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.

Hausfather, Z. and Peters, G. P. (2020). Emissions – the 'business as usual' story is misleading. *Nature*, 577(7792):618–620.

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.

Hirabayashi, Y., Mahendran, R., Koirala, S., Konoshima, L., Yamazaki, D., Watanabe, S., Kim, H., and Kanae, S. (2013). Global flood risk under climate change. *Nature Climate Change*, 3(9):816–821.

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.

ISIMIP (2024). ISIMIP2b simulation data, `water_global` sector. `https://data.isimip.org/`. Accessed via the ISIMIP Repository API.

Jasour, Z. Y., Yahyazadeh, Z., Lindell, M. K., Reuther, A., Murray-Tuite, P. M., Ge, Y., Wei, H. L., and Sadri, A. M. (2020). Hurricane mitigation adoption: A stated-preference choice experiment with u.s. atlantic and gulf coastal homeowners. *Weather, Climate, and Society*, 12(1):83–98.

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.

Kaplan, G., Violante, G. L., and Weidner, J. (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity*, 2014(1):77–153.

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. (2023). Written testimony before the u.s. senate committee on banking, housing, and urban affairs, hearing on reauthorization and reform of the national flood insurance program: Improving community resilience. May 2, 2023.

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

Mulder, P. (2024b). Mismeasuring risk: The welfare effects of flood risk information. Working paper.

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

Musselman, K. N., Lehner, F., Ikeda, K., Clark, M. P., Prein, A. F., Liu, C., Barlage, M., and Rasmussen, R. (2018). Projected increases and shifts in rain-on-snow flood risk over western north america. *Nature Climate Change*, 8(9):808–812.

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.

Painter, W. L. (2026). Disaster relief fund state of play: In brief. Congressional Research Service Report R47676.

Pauly, M. V., Kunreuther, H., and Hirth, R. (1995). Guaranteed renewability in insurance. *Journal of Risk and Uncertainty*, 10(2):143–156.

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.

Petrolia, D. R., Hwang, J., Landry, C. E., and Coble, K. H. (2015). Wind insurance and mitigation in the coastal zone. *Land Economics*, 91(2):272–295.

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.

Qin, S. and Voorheis, J. L. (2026). Flood risk, insurance, and housing in the united states. NBER Working Paper 35204, National Bureau of Economic Research.

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.

Realtor.com (2020). Realtor.com integrates flood factor data, helps home shoppers understand flood risk.

Redfin (2021). Redfin publishes flood risk data from flood factor for every u.s. home.

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.

Sharma, A., Wasko, C., and Lettenmaier, D. P. (2018). If precipitation extremes are increasing, why aren't floods? *Water Resources Research*, 54(11):8545–8551.

Solomon, A. (2024). Bundling in insurance markets: Theory and an application to long-term care. Working paper. Available at SSRN: `https://ssrn.com/abstract=4212650`.

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).

USAFacts (2025). How much does FEMA spend on disaster response? Five-year cumulative FEMA Disaster Relief Fund obligations, sourced from OpenFEMA.

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

Wasko, C. and Sharma, A. (2017). Global assessment of flood and storm extremes with increased temperatures. *Scientific Reports*, 7(1):7945.

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.

Xi, D., Lin, N., and Gori, A. (2023). Increasing sequential tropical cyclone hazards along the US East and Gulf coasts. *Nature Climate Change*, 13(3):258–265.

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	-177.1^{***}	-235.7^{**}
	(35.2)	(73.8)
Damage Control	Binned	
First-stage F-stat.		71.32
Num. Obs.	886	886
Between R^2	0.341	0.448

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

Figure 7. Response of NFIP coverage to neighboring-county floods (with own-county flood control)

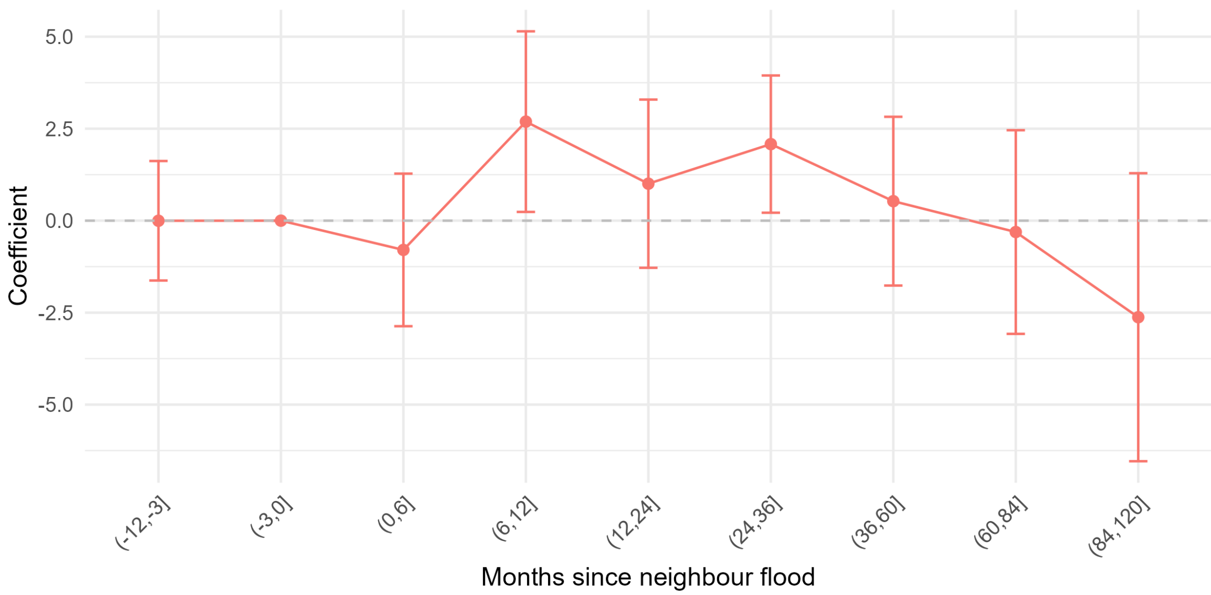

Notes: This figure reproduces Figure 1 with additional bin controls for years since the most recent own-county flood. Event-study estimates from a regression of NFIP coverage per house (measured two months before flood onset) on indicators for years since the most recent flood in the closest neighboring 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 8. Response of ex-post spending to neighboring-county floods (with own-county flood control)

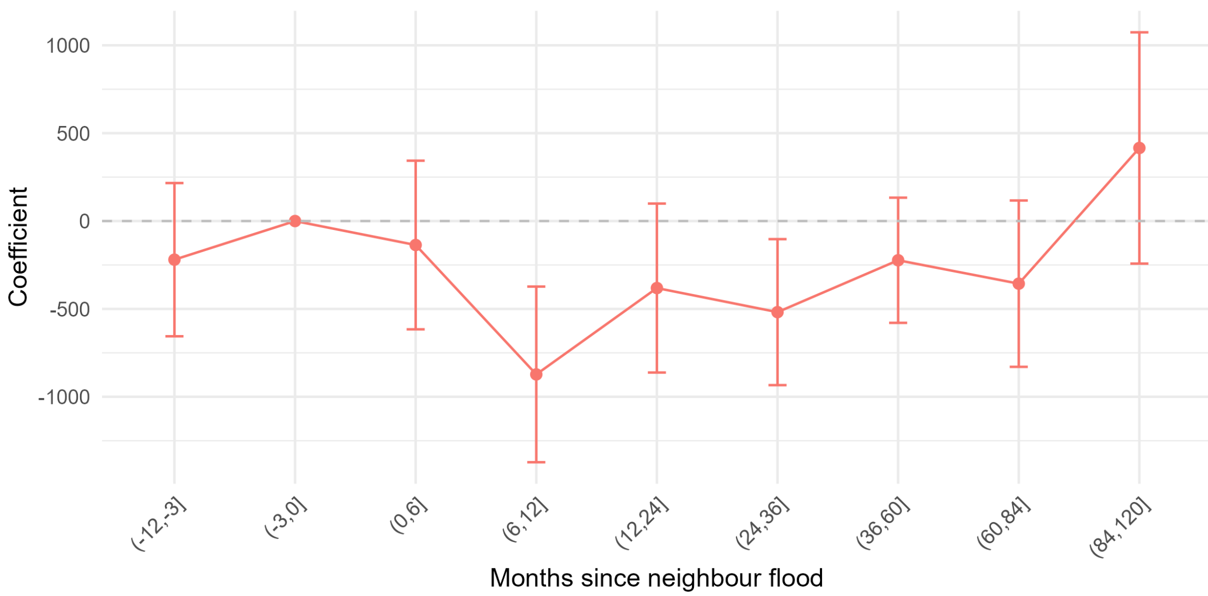

Notes: This figure reproduces Figure 2 with additional bin controls for years since the most recent own-county flood. Event-study estimates from a regression of total ex-post disaster spending per house on indicators for years since the most recent flood in the closest neighboring county, controlling 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 A2. Fiscal Cost Decomposition: FEMA Disaster Assistance

	OLS	**IV**
NFIP Policies per 100 Houses	-151.661^{***}	-183.427^{**}
	(30.464)	(58.743)
Damage Control	Binned	
First-stage F-statistic		81.98
Num. Obs.	886	886
Between R^2	0.449	0.512

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	-18.361^{*}	-9.315
	(8.789)	(17.582)
Damage Control	Binned	
First-stage F-statistic		81.98
Num. Obs.	886	886
Between R^2	0.016	0.000

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	-3.354	-21.423
	(10.066)	(18.163)
Damage Control	Binned	
First-stage F-statistic		81.98
Num. Obs.	886	886
Between R^2	0.017	0.097

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	0.606^{*}	0.541
	(0.305)	(0.849)
Damage Control	Binned	
First-stage F-statistic		75.78
Num. Obs.	886	886
Between R^2	0.009	0.004

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	0.182	-1.663
	(0.184)	(1.130)
Damage Control	Binned	
First-stage F-statistic		74.99
Num. Obs.	886	886
Between R^2	0.046	0.002

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	-172.8^{***}	-215.7^{**}	-221.4^{*}	-310.6^{**}	-167.7^{***}	-218.4^{**}
	(36.492)	(72.811)	(85.993)	(97.700)	(41.905)	(73.016)
Num. Obs.	385	385	385	385	385	385
Between R^2	0.332	0.424	0.414	0.509	0.312	0.422

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

	Cont. (Quintic)		Log-Linear		Hybrid	
	OLS	**IV**	**OLS**	**IV**	**OLS**	**IV**
NFIP Policies per 100 Houses	-190.3^{*}	-299.6^{**}	-324.1^{***}	-384.8^{***}	-158.9^{**}	-236.8^{**}
	(87.828)	(100.995)	(28.027)	(20.729)	(55.239)	(82.140)
Num. Obs.	385	385	385	385	385	385
Between R^2	0.392	0.515	0.550	0.550	0.274	0.442

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	-246.5^{***}	-351.8^{***}	-230.9^{**}	-313.9^{***}	-166.3^{**}	-233.0^{**}
	(30.925)	(28.562)	(88.204)	(77.821)	(49.906)	(79.263)
Num. Obs.	385	385	385	385	385	385
Between R^2	0.527	0.547	0.360	0.436	0.317	0.458

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 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 9 plots the estimated $\beta_{t'}$.

Figure 9. Between-County Demand Estimates

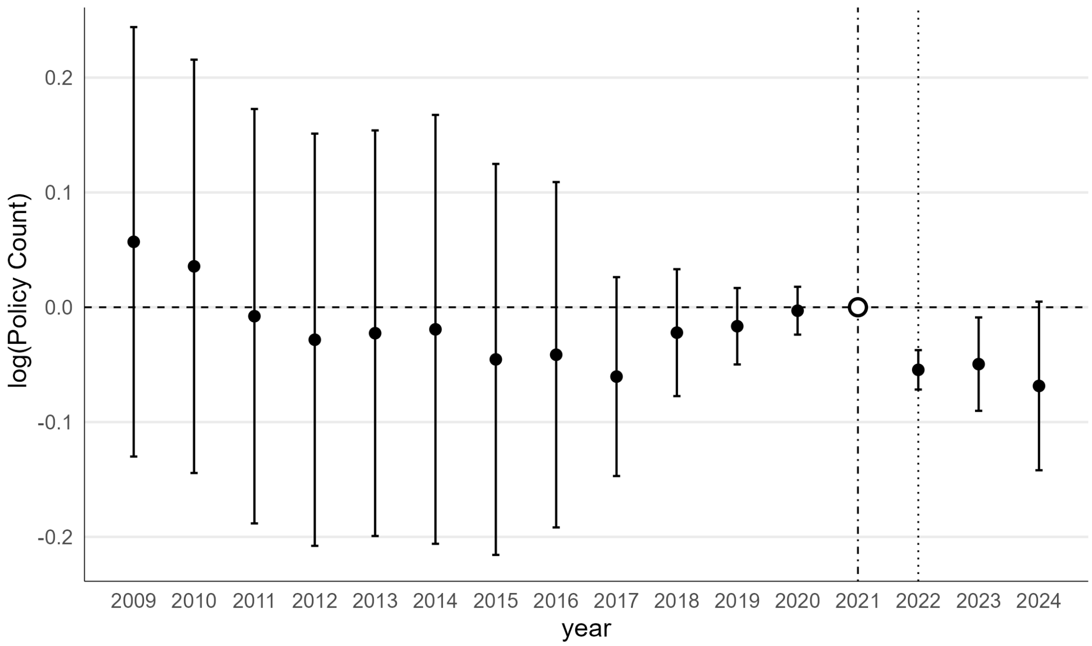

We see, by the end of 2024, a drop in demand that is 13% higher in the treated counties than the control. The treated counties have average pre- and post-RR2.0 premiums of $\$918$ and $\$1{,}808$ respectively, while the control counties have $\$688$ and $\$904$. In other words, the relative price increase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

B.3 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_t 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 A9.

We see a statistically precise, but economically negligible effect decrease along the intensive margin. The average level of coverage per policy is \$315,000. 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.4 Demand Elasticity by Recent Flood Exposure

Table A10 re-estimates the within-policy renewal semi-elasticity of Equation 2, interacting log price with an indicator for a recent flood — in the property's own (focal) county and in its nearest neighbor — at 1-, 3-, and 5-year windows. The elasticity barely moves with recent-flood exposure. Against the pooled baseline of -0.32, the interaction term is uniformly tiny — between -0.005 and 0.007 across all six focal- and neighbor-county specifications. These interactions are precisely estimated (the sample has over two million policy-years) but economically negligible: households respond to the RR2.0 price change at essentially the same rate whether or not their area has recently flooded.

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.

Table A9. 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	6,134,818	6,134,818	6,134,818

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: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$.

Table A10. Renewal Semi-Elasticity by Recent Flood Exposure

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Baseline		*Focal county*			*Nearest neighbour*	
	—	1y	3y	5y	1y	3y	5y
$\hat{\eta}$ (pooled)	-0.322^{***}	—	—	—	—	—	—
	(0.002)	—	—	—	—	—	—
$\hat{\eta}$ (no recent flood)	—	-0.320^{***}	-0.319^{***}	-0.306^{***}	-0.322^{***}	-0.321^{***}	-0.302^{***}
	—	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Interaction $\hat{\gamma}$	—	0.007^{***}	0.005^{***}	0.007^{***}	-0.005^{***}	0.006^{***}	0.006^{***}
log$(\tilde{P}) \times \mathbb{1}\{$recent$\}$	—	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Policy FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	2,084,307	2,084,307	2,084,307	2,084,307	2,084,307	2,084,307	2,084,307

Notes: Each column reports a within-policy linear probability model of NFIP renewal on log $\tilde{P}$ (log price relative to the full-risk premium). Column (1) replicates the paper's baseline pooled estimate $\hat{\eta}$. Columns (2)–(4) and (5)–(7) add an interaction with an indicator for a recent flood — within 1, 3, or 5 years — in the focal county and in its nearest neighboring county, respectively; $\hat{\eta}$ (no recent flood) is the price coefficient for the non-exposed subgroup and $\hat{\gamma}$ is the interaction (a test of equality across the two subgroups). Standard errors clustered at the policy level. Significance: $+\, p < 0.10$, $*\, p < 0.05$, $**\ p < 0.01$, $***\, p < 0.001$.

B.5 Potential-Outcomes Framework for Within-Policy Demand Identification

Treatment and estimand. For policy i at year t, let $D_{i,t} \equiv \ln \tilde{P}_{i,t}$ denote the log relative price faced at renewal. Let $y_{i,t+1}(d)$ denote the potential renewal indicator if the policy faced log price d. The coefficient β in Equation 2 is the population-average derivative

$$\beta \;=\; \frac{\partial \mathbb{E}\big[y_{i,t+1}(d)\big]}{\partial d},$$

identified from within-policy variation in $D_{i,t}$.

Identifying assumption. The conditional-independence assumption is that potential outcomes are independent of the realized log price, given the policy fixed effect:

$$\{y_{i,t+1}(d)\}_d \;\perp\; D_{i,t} \,\Big|\, \mu_i.$$

Within a given policy, year-to-year variation in $D_{i,t}$ — driven by the glide-path schedule of premium increases under RR2.0 — must be uncorrelated with within-policy variation in unobserved determinants of renewal. The conditional-independence assumption can fail through three channels.

Anticipation. RR2.0 was pre-announced before its October 2021 effective date. If households responded to the entire forward-anticipated path of premium increases rather than the within-year realized change, then $D_{i,t}$ understates the relevant treatment intensity and η is biased toward zero. We address this in Appendix B.6 and in the robustness to s^* in Table 6, where we rescale the demand moments under both "contemporaneous-only" and "full-path anticipation" assumptions.

Information confounds. First Street Foundation began publishing parcel-level flood-risk scores in 2020, overlapping the RR2.0 transition window. Since these estimates were incorporated simultaneously across the U.S., our between-county estimate (Appendix B.2) is not affected. Moreover, in a completely different setting, Mulder (2024b) identifies the price elasticity of NFIP demand separately from the risk-information content and obtains $\hat{\eta} \in [0.25, 0.29]$, similar to our headline of 0.32.

Private-flood-market substitution. Private flood insurance policies could be differentially available across the U.S. in a way that is correlated with RR2.0 price increases. In Appendix B.24 we find at most a 20% spillover from NFIP to private flood. Hence, in the worst case, our demand elasticity of 0.32 might be 20% too large. This is well within the elasticities we analyze in our robustness of s^* in Table 6.

Pre-treatment trends. The demand specification is identified within policy, with a policy fixed effect absorbing every time-invariant policy- and property-level characteristic. The cross-policy level differences that a conventional balance table would summarize are differenced out and do not bear on identification — a levels balance table is not informative here. What matters is whether within-policy variation in the price change is correlated with within-policy variation in unobservables.

The ideal test of that condition is a pre-trend test in the within-policy treatment intensity itself. That is not directly feasible: a policy's RR2.0 treatment intensity is observed only for policies that survive into the RR2.0 era, so binning at the policy level conditions on survival. We therefore test for pre-trends at the minimal geographic aggregation that avoids this survivor bias — an event study of pre-2021 renewal rates by census-block-group quartile of future treatment intensity, estimated on the full panel, which includes policies that exit before RR2.0. Figure 10 reports the result. Pre-period renewal trends are unfortunately not parallel, but in a way that only strengthens our eventual conclusion: following treatment, renewal rates were approximately 1.5% lower in the most treated block groups than the least.

Figure 10. Pre-Treatment Renewal Trends by Future RR2.0 Exposure

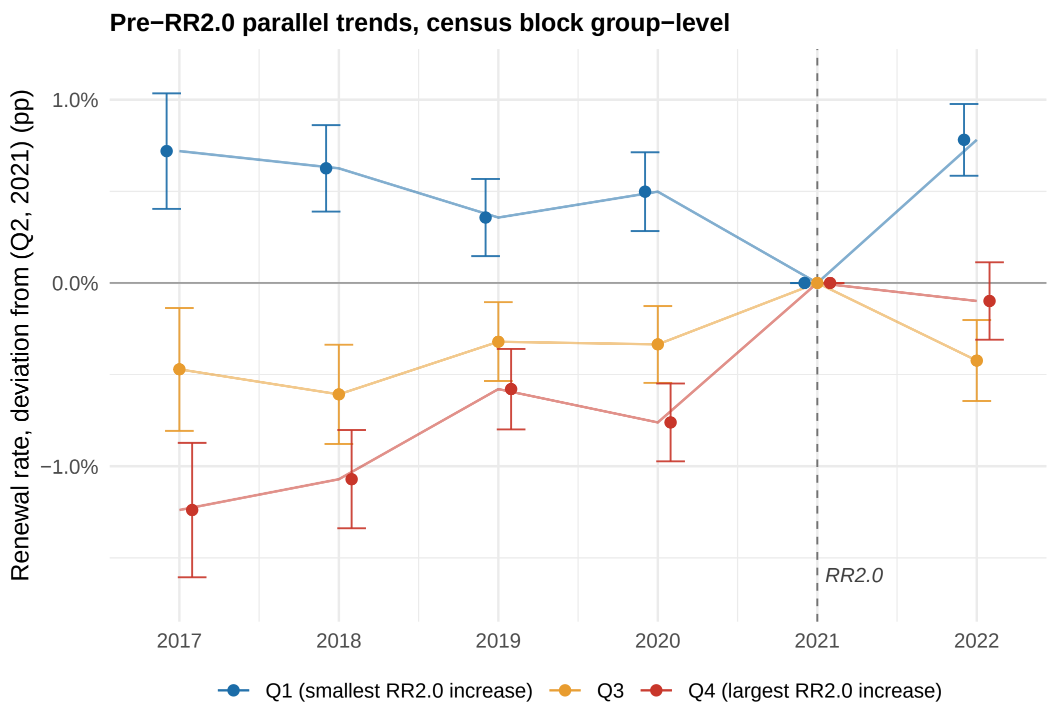

Notes: Event-study coefficients on year × block-group quartile of future RR2.0 treatment intensity, for pre-2021 NFIP renewal rates. Bars are 95% confidence intervals.

B.6 Anticipation of RR2.0 Premium Changes

Our headline estimates deliberately make different anticipation assumptions for the demand and the moral-hazard margins. The demand semi-elasticity is estimated from realized — that is, capped — price changes, while the mitigation and relocation elasticities are estimated assuming households account for the entire path of RR2.0 price increases. This is by construction. Demand elasticities concern existing policies deciding whether to renew, and a renewing policyholder faces

Table A11. RR2.0 Anticipation: Elasticities and Optimal Subsidy

	Demand $\hat{\eta}$	Mitigation $\hat{E}_m$	Location $\hat{E}_x$	Optimal s^*
No anticipation	-0.322	0.62	0.0159	43%
Full anticipation	-0.059	0.34	0.0077	39%

Notes: For each of demand and moral hazard (mitigation and relocation), columns 2, 3, and 4 show the elasticity under "no anticipation" of the RR2.0 price path and under "full anticipation"; column 5 is the optimal subsidy s^* at each row's elasticities, with all other inputs at the baseline calibration. The paper's baseline $s^* = 52\%$ pairs the top-left demand elasticity with the bottom-row mitigation and location elasticities. The remaining cells rescale by the ratio of the realized cumulative RR2.0 price change to the full Full-Risk-Premium gap.

only the realized within-year change, which RR2.0 caps at 18% per year. In contrast, mitigation occurs at new construction, and relocating households do not benefit from the 18% cap. Hence, in both cases, the full RR2.0 risk-rate price rise is experienced immediately. Each margin is therefore estimated under the anticipation assumption that actually applies to it.

To check how this design choice affects our results, Table A11 re-estimates the optimal subsidy under two uniform alternatives: all three elasticities under no anticipation, and all three under full anticipation. The optimal subsidy moves only modestly — to 43% and 39% respectively, against the baseline of 52% — and our broader robustness in Table 6 spans demand elasticities well below these.

B.7 Incidence of Premium Decreases by Construction Vintage

Premium decreases under RR2.0 turn out to be both uncommon and small. Table A12 reports their incidence by construction vintage. Nationally, only a few percent of policies saw any decrease — from 6.75% of pre-1970 homes down to 0.25% of post-2021 construction — and the median decrease is modest, between roughly \$80 and \$480. Decreases are concentrated in the oldest housing stock, which carried the most over-priced legacy premiums.

This pattern is what matters for our mitigation identification. Premium decreases barely occur for new construction, and our mitigation moral-hazard estimates are identified from new construction — which realizes the full RR2.0 risk-rated price. The premium-decrease margin therefore does not touch the population that drives our E_m estimate, so it cannot have reduced the mitigation response we measure.

B.8 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,

Table A12. RR2.0 Premium-Decrease Incidence by Construction Vintage

	National				Florida		
Construction vintage	Policies (n)	Decreases (n)	% with decrease	Median $ drop	Policies (n)	% with decrease	Median $ drop
pre–1970	1,054,448	71,204	6.75	$484	223,820	4.04	$320
1970–1979	635,587	36,344	5.72	$407	193,219	3.82	$239
1980–1989	616,536	29,308	4.75	$200	224,947	3.52	$121
1990–1999	574,039	24,277	4.23	$138	201,529	3.49	$101
2000–2009	745,137	32,167	4.32	$128	236,890	3.58	$107
2010–2014	190,057	7,564	3.98	$121	46,701	3.29	$100
2015–2020	313,030	11,010	3.52	$107	98,144	2.83	$88
2021+	145,593	365	0.25	$79	54,495	0.19	$46

Notes: For each construction-vintage bucket, the number of policies, the number and share that saw any premium decrease at the RR2.0 transition, and the median dollar drop among decrease-side policies; national figures on the left, Florida on the right.

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

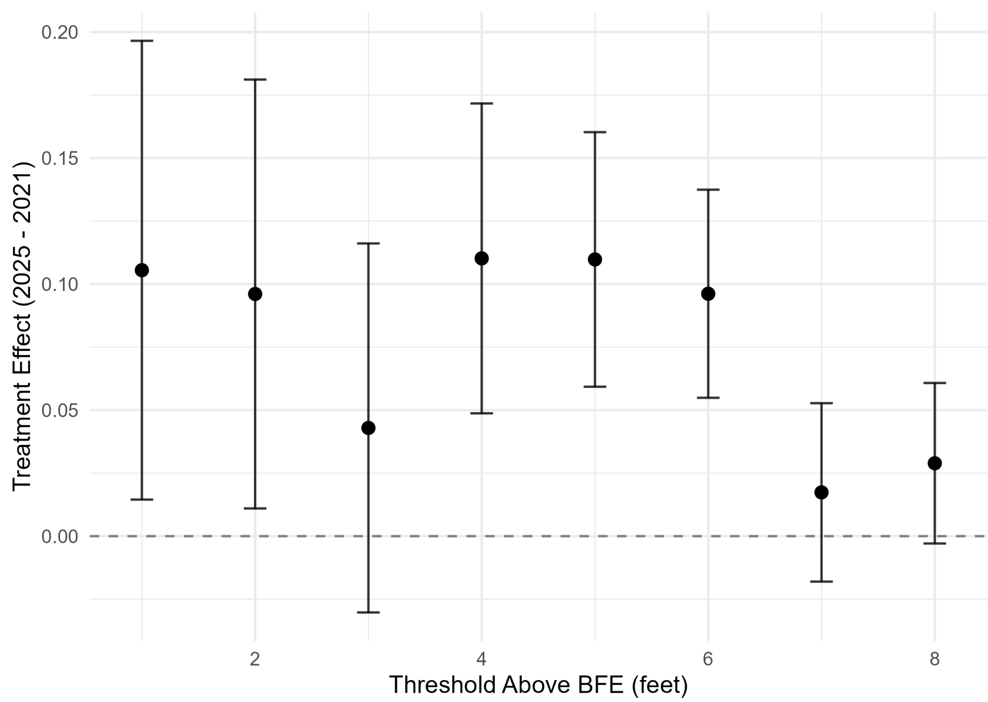

Notes: Each coefficient is the 2025 coefficient from specification (3) where the threshold is varied according to the horizontal axis. We restrict to the very-high-risk and high-risk flood zones inside the SFHA, and code the former as treated, the latter as control. 95% confidence intervals are shown.

Figure 12. Elevation Threshold Robustness Outside SFHA

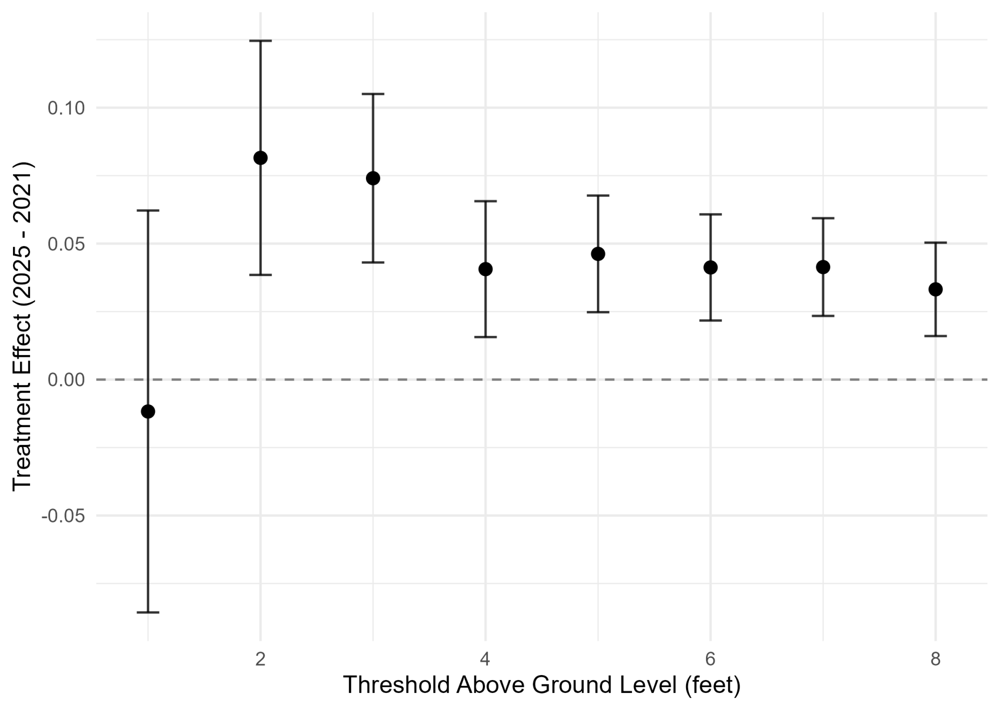

Notes: Each coefficient is the 2025 coefficient from specification (3) where the threshold is varied according to the horizontal axis. We restrict to the moderate- and low-risk flood zones outside the SFHA, and code the former as treated, the latter as control. 95% confidence intervals are shown.

Figures 11 and 12 show that the 7 to 11% 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.9 Sufficient Statistic Heterogeneity

The optimal-subsidy formula in Proposition 1 is a function of five sufficient statistics: the value of insurance against idiosyncratic reclassification risk, the FEMA fiscal spillover β, the demand semi-elasticity η, the location semi-elasticity E_x, and the mitigation semi-elasticity E_m. In the main paper we calibrate each to it's average. In this appendix we examine how each varies across observable dimensions when those dimensions: every sufficient statistic is re-estimated from a single regression that includes eight above-/below-median dummies simultaneously.

Heterogeneity dimensions and data sources. We stratify along eight dimensions: four census-tract demographics (median household income, median tenure, residential mobility rate, median age, each from the 2023 ACS five-year estimates) and four property attributes (building age, NFIP policy tenure, historical NFIP claims, and primary-residence indicator, from the 2023 NFIP policyholder snapshot). For continuous dimensions we form an above-/below-median indicator at the appropriate cross-section: policy-level for Tables A13 and A14, tract-level for Table A15, community-level for Table A16. The two categorical dimensions—claims and primary residence—map naturally to ≥ 1 claim vs. zero and primary vs. non-primary. The FEMA spillover β is only analyzed heterogeneously with respect to income; no other heterogeneity produces any precise results.

B.9.1 Reclassification Risk

For each policy we compute the two-way residual ε from the insurer's predicted loss-cost rate (Section 6), along with the per-insured CARA welfare gain from an $s = 0.5$ subsidy at absolute risk aversion $a = 5 \times 10^{-4}$, separately under SSP5-8.5 and SSP2-4.5. We regress each per-policy metric on the eight above-/below-median dummies jointly, clustering standard errors at the tract. Table A13 reports the partial effects.

Table A13. Heterogeneity in Reclassification Risk by Household Characteristics

Dimension	SSP5-8.5		SSP2-4.5	
	σ_ε	Reclass. Ins. Value	σ_ε	Reclass. Ins. Value
Baseline (pool)	$1,978.5$	228.9	$1,841.7$	206.6
Marginal Differences on Observables				
Tract median household income	253.0^{*} (120.7)	34.5^{**} (10.8)	272.9^{**} (99.9)	46.6^{***} (9.8)
Tract median tenure (yrs)	-217.4^{*} (110.7)	-24.0^{*} (10.7)	-134.0 (98.2)	1.8 (10.3)
Tract mobility rate	248.3^{*} (115.3)	15.9 (10.9)	136.0 (102.2)	2.4 (10.5)
Tract median age	470.5^{***} (113.0)	62.5^{***} (10.3)	476.6^{***} (94.3)	57.9^{***} (9.0)
Building age	-23.1 (73.2)	7.8 (6.8)	27.1 (81.8)	13.3 (9.1)
NFIP policy tenure	2.0 (45.0)	2.0 (4.4)	21.0 (52.2)	7.3 (6.3)
Historical NFIP claims (1+)	283.2 (201.5)	33.3 (22.4)	317.1^{+} (192.2)	32.7 (20.9)
Primary residence	53.1 (61.8)	0.8 (7.4)	-20.6 (115.3)	-21.1 (15.3)

Notes: The top row reports the headline estimates of idiosyncratic-risk dispersion σ_ε and the mean per-policy reclassification-risk insurance value of a 50% subsidy. Each subsequent row reports how those two quantities shift with an above-median indicator on the named dimension, from a multivariate regression that holds the other seven fixed; standard errors are clustered at the tract level. The right two columns use the less extreme SSP2-4.5 warming scenario. Significance: $+\, p < 0.10$, $*\, p < 0.05$, $**\ p < 0.01$, $***\, p < 0.001$.

Median tract age is the strongest predictor, with older-tract policies carrying a \$62 higher CARA welfare gain under SSP5-8.5, holding other dimensions constant. Household income also predicts a positive gradient of about \$35, and tract tenure (conditional on mobility and age) carries a small *negative* partial effect, reversing the univariate positive correlation. Building age, NFIP tenure, and historical-claims indicators are not significant once the other dimensions are included. Under SSP2-4.5, tract median age remains the largest insurance-value gradient (\$58) and income is also significant (\$47); tract tenure and primary residence change sign relative to SSP5-8.5 but are imprecisely estimated, while the other dimensions remain statistically insignificant.

B.9.2 Demand Semi-Elasticity

We re-estimate the within-policy renewal regression of Section 5 with log price interacted with all eight above-median dummies simultaneously:

$$\text{renew}_{i,t} = \eta_0 \ln \widetilde{P}_{i,t} + \sum_{k=1}^{8} \gamma_k \cdot \ln \widetilde{P}_{i,t} \cdot \mathbf{1}\{\text{hi}_k\} + \mu_i + \delta_t + \varepsilon_{i,t+1},$$

with μ_i and δ_t policy and year fixed effects and standard errors clustered at the policy level. Table A14 reports $\hat{\gamma}_k$, the partial change in demand semi-elasticity associated with being above the median on dimension k holding the other seven at their sample means. The base semi-elasticity (all dummies switched to zero, i.e., all covariates below the meian) is $\hat{\eta}_0 = -0.271$.

Table A14. Demand Semi-Elasticity $\hat{\eta}$ Heterogeneity (Multivariate)

Dimension	$\Delta\hat{\eta}$ (interaction with log p)
Tract median household income	0.0238 (0.0051)***
Tract median tenure (yrs)	0.0220 (0.0053)***
Tract mobility rate	-0.0082 (0.0052)
Tract median age	0.0202 (0.0051)***
Building age	0.0220 (0.0051)***
NFIP policy tenure	0.0267 (0.0048)***
Historical NFIP claims (1+)	-0.0227 (0.0130)+
Primary residence	-0.1459 (0.0051)***

Notes: Feols renew $\sim$ log p + log $p \cdot \mathbf{1}\{\text{hi}_k\} \times 8$ | panel_id, cluster = panel_id. Each row reports the interaction log $p \cdot \mathbf{1}\{\text{hi}_k\}$ holding the other seven above-median dummies fixed. Sample $N = 1,151,619$ policy-years (post-singleton-drop). Base $\hat{\eta}$ (all dims below median) = -0.2713.
Significance: $+\, p < 0.10$, $*\, p < 0.05$, $**\ p < 0.01$, $***\, p < 0.001$.

Primary residence remains the dominant gradient: primary-residence policyholders are about 14.6 percentage points more elastic ($p < .001$), and this coefficient is essentially unchanged between the univariate and multivariate specifications. Income, tract tenure, median age, building age, and NFIP tenure all flatten demand by about 2 percentage points each, conditional on the other dimensions. Mobility and historical claims are not significant.

B.9.3 FEMA Fiscal Spillover

We analyze the heterogeneity of the FEMA Fiscal Spillover only with respect to tract income. All other dimensions are insignificant and underpowered jointly or univariately. The results are reported in Figure 13. Lower income households have higher FEMA spillovers. For this reason, when we compute heterogeneous optimal subsidies, these will decline in tract-income level.

Figure 13. FEMA Fiscal Spillover by County Median Income

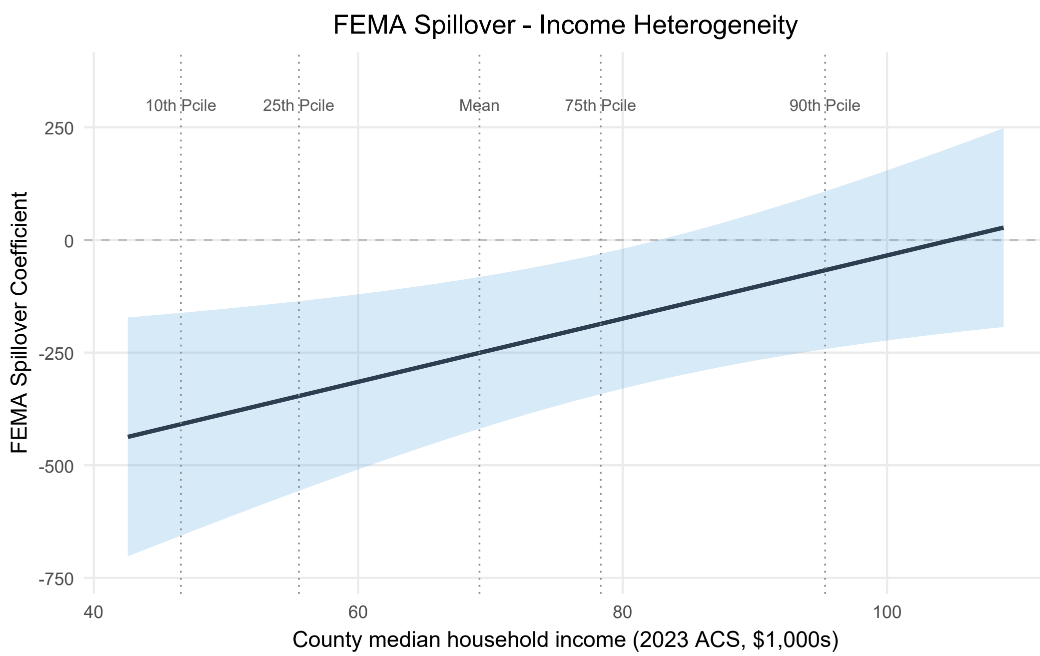

Notes: Implied FEMA fiscal-spillover coefficient $\widehat{\beta}$ per additional NFIP policy per 100 houses, as a linear function of county median household income. Univariate continuous-IV specification. Pooled point estimate $\widehat{\beta} = -215.7$ (SE 72.8). At the 10th percentile of county income ($46,600), $\widehat{\beta} = -\$410$; at the 90th percentile ($95,300), $\widehat{\beta} = -\$67$ and not significantly different from zero.

B.9.4 Location Response

We run a single tract-quarter event study with all eight tract-level above-median dummies interacted with the treatment-by-period indicator, plus per-period level controls for each hi dummy:

$$\log \text{occ\_res}_{i,t} = \sum_{\tau} \alpha_\tau \cdot T_i \cdot \mathbf{1}\{\text{rel\_q} = \tau\} + \sum_{\tau,k} \gamma_{k,\tau} \cdot T_i \cdot \mathbf{1}\{\text{hi}_{i,k}\} \cdot \mathbf{1}\{\text{rel\_q} = \tau\} + \cdots + \mu_i + \delta_t + \varepsilon_{i,t}.$$

We extract the final-period coefficient ($\tau = 12$, the last quarter in our USPS panel) for the base T and each $T \times \text{hi}_k$ interaction. Table A15 reports the resulting E_x. Building age and NFIP policy tenure are the only gradients significant at $p < .05$; tract tenure is marginally significant at $p < .10$.

Table A15. Location Event-Study Coefficient by Heterogeneity Cut (Multivariate)

Bin	Mean value	Event-study coef $\widehat{\beta}^{(12)}$ log occupied addresses
Tract median household income		
Below Median	$63,071	-0.00205 (0.00153)
Above Median	$125,515	-0.00496 (0.00203)*
Tract median tenure (yrs)		
Below Median	7.9	-0.00597 (0.00204)**
Above Median	13.8	-0.00071 (0.00171)
Tract mobility rate		
Below Median	0.072	-0.00115 (0.00174)
Above Median	0.170	-0.00565 (0.00197)**
Tract median age		
Below Median	37.1	-0.00265 (0.00140)+
Above Median	52.9	-0.00423 (0.00233)+
Building age		
Below Median	26.1	-0.01212 (0.00258)***
Above Median	52.1	0.00119 (0.00126)
NFIP policy tenure		
Below Median	7.3	-0.00450 (0.00136)***
Above Median	13.6	0.00211 (0.00234)
Historical NFIP claims (1+)		
0 claims		-0.00312 (0.00127)*
1+ claims	1.89	-0.00410 (0.00377)
Primary residence		
Non-primary	0.613	-0.00487 (0.00202)*
Primary	0.902	-0.00164 (0.00142)

Notes: Raw event-study treatment effect $\widehat{\beta}^{(12)}$ on log(occupied addresses) at the final available period (rel_q $=$ 12). The regression interacts T_{median^+} $\times$ rel_q with eight above-median tract dummies simultaneously; T_{median^+} is the main-paper baseline indicator (median NFIP policy in tract experiences a premium increase under RR2.0). Each cell is a linear contrast of MV coefficients evaluated at $\mathbf{1}\{\text{hi}_k\} \in \{0, 1\}$ with the other seven above-median dummies held at their sample means; SEs from the MV variance, clustered by tract; p from the normal approximation. Values are reported on the same log scale as the paper's baseline event study for direct comparability. No sig-based collapse to pool applied here (that is deferred to Table 13).
Significance: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

B.9.5 Mitigation Treatment Effect

We run a single Florida EC event study with all eight community-level above-median dummies interacted with the combined zone treatment (V or X-shaded, relative to A or X-unshaded) by issued-year, plus per-year level controls for each hi dummy and the SFHA indicator. We extract the final-period coefficient (issued-year = 2025) for the base and each interaction term. Table A16 reports the result; the pooled 2025 raw coefficient is used directly as the pooled E_m. Three dimensions differ significantly across bins at $p < .05$: tract median income, building age, and historical claims. The raw 2025 coefficient is lower above the income median (0.136 vs. 0.320), higher above the building-age median (0.327 vs. 0.209), and higher for communities with historical claims (0.278 vs. 0.096).

Table A16. Mitigation Event-Study Coefficient by Heterogeneity Cut (Multivariate)

Bin	Mean value	Event-study coef $\widehat{\beta}^{(2025)}$ elevation (ft above BFE / TBF)
Tract median household income		
Below Median	\$63,071	0.3202 (0.0512)***
Above Median	\$125,515	0.1358 (0.0430)**
Tract median tenure (yrs)		
Below Median	7.9	0.2205 (0.0446)***
Above Median	13.8	0.3270 (0.0809)***
Tract mobility rate		
Below Median	0.072	0.3105 (0.0818)***
Above Median	0.170	0.2162 (0.0453)***
Tract median age		
Below Median	37.1	0.1811 (0.0872)*
Above Median	52.9	0.2540 (0.0401)***
Building age		
Below Median	26.1	0.2086 (0.0437)***
Above Median	52.1	0.3272 (0.0485)***
NFIP policy tenure		
Below Median	7.3	0.2678 (0.0458)***
Above Median	13.6	0.2066 (0.0520)***
Historical NFIP claims (1+)		
0 claims		0.0957 (0.0571)+
1+ claims	1.89	0.2776 (0.0417)***
Primary residence		
Non-primary	0.613	0.2996 (0.0426)***
Primary	0.902	0.1657 (0.0691)*

Notes: Raw event-study treatment effect $\widehat{\beta}^{(2025)}$ on building elevation (ft above BFE in SFHA; TBF in non-SFHA). The regression interacts treatC $\times$ issuedYear with eight above-median community-level dummies simultaneously. Each cell is a linear contrast of MV coefficients evaluated at $\mathbf{1}\{\mathrm{hi}_k\} \in \{0, 1\}$ with the other seven above-median dummies held at their sample means; SEs from the MV variance; p from the normal approximation. Values are reported on the same scale as the paper's baseline mitigation event study. No sig-based collapse to pool applied here (that is deferred to Table 13).
Significance: $+\, p < 0.10$, $*\, p < 0.05$, $**\, p < 0.01$, $***\, p < 0.001$.

B.10 Robustness to Alternate Damage Measures

Our spillover regression controls for flood damage using the NOAA Storm Events Database (SED). The SED is subject to non-random missingness in reporting (Gallagher, 2023). Table A17 re-estimates the spillover with completely distinct damage measures without the same reporting completeness concerns. Columns (2)–(5) add the GloFAS river-discharge flood-extent measures of Farkas and Kim (2025) - a physically-based, remotely-sensed measure of water levels - alongside or in place of the SED damage bins; columns (6)–(8) instead use realized FEMA Public Assistance spending (i.e., infrastructure damage) per capita, the alternative Gallagher (2023) adopts precisely because it avoids the SED missingness. The spillover estimate is stable across all of these.

Table A17. Spillover Robustness to Alternative Damage Controls

	Baseline	+ GloFAS	+ GloFAS No Hurr/Coast	GloFAS only	GloFAS only No Hurr/Coast	+ PA	PA only	All three
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: IV estimate (Neigh. County Flood $\leq$ 5 Years)								
NFIP Policies per 100 Houses	-215.7^{**}	-203.6^{*}	-190.6^{*}	-393.6^{***}	-368.6^{***}	-200.4^{**}	-360.1^{***}	-185.0^{*}
	(72.8)	(83.9)	(77.1)	(55.0)	(59.5)	(74.7)	(42.3)	(88.0)
First-stage F-stat.	82.0	68.7	71.6	65.8	67.2	75.0	140.7	62.5
Num. Obs.	886	886	822	886	822	886	886	886
Panel B: Damage controls								
SED $\$$-damage bins ($\times$ dmg, dmg^2)	✓	✓	✓			✓		✓
GloFAS share $>$ 2-yr flood (UW)		✓	✓	✓	✓			✓
GloFAS mean log return-period (UW)		✓	✓	✓	✓			✓
PA fed. share $\$$ / housing unit						✓	✓	✓
Drop hurricane $\times$ coastal-shoreline county			✓		✓			

Notes: Each column re-estimates the spillover IV with alternative or additional damage controls. Columns (2)–(5) add the GloFAS river-discharge flood-extent measures of Farkas and Kim (2025) (alongside or instead of the baseline SED damage bins); columns (6)–(8) add realized FEMA Public Assistance federal-share spending per housing unit. Columns (3) and (5) repeat (2) and (4) on a sample that drops cells that are both hurricane disasters and coastal-shoreline counties, as GloFAS is only directly relevant to pluvial and fluvial floods. Significance: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

B.11 Robustness to Continuous Instruments

The paper's headline spillover IV uses a binary indicator: whether the nearest neighbouring county experienced any FEMA-declared flood disaster in the past five years. Table A18 re-estimates the IV with three continuous alternatives that sum, over the same five-year window, the neighbour county's (a) total flood property damage (NOAA SED dollars), (b) total FEMA-IHP applications, and (c) total FEMA-IHP-approved households.

B.12 Robustness to Disaster-Relief-Fund Depletion

The neighbor-county-flood instrument could violate the exclusion restriction if a neighbor's flood depletes FEMA's federal Disaster Relief Fund (DRF) and thereby mechanically reduces the ex-post aid available to the focal county.

When the DRF runs low, FEMA institutes "Immediate Needs Funding" (INF), which pauses obligations for new non-lifesaving Public Assistance and hazard-mitigation projects. Individual Assistance payments to disaster survivors — the Individuals and Households Program that constitutes the bulk of our spillover measure — continue regardless (Federal Emergency Management Agency, Office of the Chief Financial Officer, 2024; Painter, 2026). The DRF-depletion channel therefore cannot mechanically reduce IHP aid to the focal county.

For completeness, Table A19 re-estimates the primary spillover IV with an interaction term for the three INF episodes in the sample (2010, 2011, and 2017, covering 148 of the 886 damage-weighted observations). The interaction coefficient is -91.1 (s.e. 37.2): the FEMA ex post spillover is, if anything, larger during INF episodes. This goes against the predicted direction of a DRF-depletion confound.

B.13 Robustness to State-Fund Depletion

A similar exclusion-restriction concern applies to state-specific disaster relief funds, which could be depleted by a neighbor's flood. Comprehensive panel data on state disaster-relief fund balances are not available, so we cannot run a directly analogous test. Instead, we rebuild the IV by neighbor location in the following way to address the concern.

Table A20 reports the spillover using two alternative neighbor restrictions: column (2) instruments with the nearest neighbor in the same state as the focal county, and column (3) instruments with the nearest neighbor in a different state. The cross-state IV (column 3) is immune by construction to any state-fund-depletion confound, since the focal county and the instrumenting neighbor draw from different state pools.

The estimates are similar across columns: -241.2 within-state versus -277.6 cross-state. The cross-state spillover is, if anything, slightly larger, the opposite of what a state-fund-depletion confound would predict. The baseline IV (column 1, -215.7) need not fall between the within-state and cross-state estimates, since it identifies off the single nearest neighbor wherever located, while the

Table A18. Spillover Estimate with Continuous Instruments

Instrument	**IV coefficient**	**SE**	**First-stage F**
Any Neighbour Flood Disaster *(binary; published Table 2 col. 2)*	-215.7^{**}	(72.8)	81.98
Total Neighbour Flood Damage ($) *(NOAA SED property damage)*	-212.1^{***}	(47.5)	293.44
Total Neighbour IHP Applications *(count of households)*	-174.3^{***}	(49.4)	818.59
Total Neighbour IHP-Approved Households *(count of households)*	-173.6^{***}	(49.3)	803.28

Notes: IV estimates of the total fiscal spillover from NFIP coverage, comparing the paper's binary instrument (any FEMA flood disaster in the nearest neighboring county in the past five years) against three continuous variants that sum, over the same five-year window, the neighbor county's flood property damage, its FEMA IHP applications, and its IHP-approved households. All four use the same $N = 886$ regression sample, with two-way fixed effects, quadratic damage controls, damage weights, and county-clustered standard errors. Significance: $+\, p < 0.10$, $*\, p < 0.05$, $**\ p < 0.01$, $***\, p < 0.001$.

Table A19. Spillover IV and Disaster-Relief-Fund Depletion

	Baseline (1)	**IV $\times$ DRF depleted** (2)
Panel A: IV estimates (Neigh. County Flood $\leq$ 5 Years)		
NFIP Policies per 100 Houses	-215.7^{**}	-204.6^{*}
	(72.8)	(98.3)
$\times$ DRF depleted (INF)		-91.1^{*}
		(37.2)
First-stage F-stat.	81.98	41.16
Num. Obs.	886	886

Notes: Dependent variable is total ex-post disaster spending per house. Column (1) reproduces the paper's main spillover IV; column (2) adds an interaction with an indicator that the focal disaster fell during a FEMA "Immediate Needs Funding" episode — a period when the Disaster Relief Fund was depleted (three such episodes, in 2010, 2011, and 2017, cover 148 of the 886 damage-weighted observations). Both NFIP coverage and the interaction are instrumented by the neighbor-county flood instrument and its interaction. Significance: $+\, p < 0.10$, $*\, p < 0.05$, $**\ p < 0.01$, $***\, p < 0.001$.

restricted IVs identify off the nearest neighbor satisfying the state constraint (potentially the second- or third-nearest in the unrestricted ordering).

B.14 Robustness to Simultaneous and Recent Own-County Floods

A concern with the neighbor-county-flood instrument is that the damage data are reported only at the county level, so the binary indicator could be picking up not just a neighbor's flood but also a contemporaneous or recent flood in the focal county itself. The original paper addressed this in part by controlling for time since the most recent focal-county flood (Appendix Table A1), and the Gallagher (2014) literature treats own-county recent floods as a salience shock to NFIP coverage in their own right, suggesting these events should not be excluded from the analysis indiscriminately.

Nevertheless, to address the concern directly we re-estimate the spillover with progressively stricter drop rules and IV definitions. Table A21 reports the spillover under four such variants. The first two columns retain the paper's primary IV (any neighbor-county flood in the past five years) but drop focal county-disasters whose own county also experienced a same-type (e.g., fluvial/pluvial vs hurricane) flood within the past 1, 3, or 5 years. Column 1 applies this drop at the whole-county level; column 2 applies a more surgical rule, dropping only when the prior and focal floods affected at least two overlapping ZIP codes within the focal county. The last two columns additionally narrow the IV window itself: the neighbor-flood IV is restricted to a 6-month-to-X-year band, where X matches the row's drop window. This both rules out simultaneous events (the 6-month buffer) and ensures the window from which our IV variation is derived does not overlap with own-county floods.

The spillover estimate is robust to all of these modifications, albeit with a loss of power as more and more floods are excluded.

Table A21. Spillover IV: Robustness to Simultaneous and Repeat Floods

	Fixed 5yr IV		**Matched 6mo–Xyr IV**	
Drop window	Whole county	Shared ZIPs	Whole county	Shared ZIPs
No drop	-215.7^{**}	-215.7^{**}	-216.9^{***}	-216.9^{***}
	(72.8)	(72.8)	(59.7)	(59.7)
1 year	-157.9^{**}	-154.7^{**}	-309.0^{**}	-312.2^{**}
	(53.3)	(52.6)	(96.8)	(98.1)
3 years	-207.9^{***}	-190.7^{***}	-330.8^{***}	-271.1^{***}
	(58.6)	(53.9)	(92.5)	(79.6)
5 years	-219.8	-149.5	-227.6	-166.2
	(148.2)	(95.3)	(212.8)	(149.5)

Notes: Each cell re-estimates the spillover IV after dropping focal county–disasters whose county had a prior declared, IHP-authorized same-type (fluvial, pluvial vs. hurricane) flood within the listed window; the first two columns use the canonical 5-year neighbor-flood indicator, and the last two narrow the neighbor-flood window to a 6-month-to-X-year band matched to the row's drop window ($X = 5$ in the "No drop" row). The "Whole county" columns apply this rule at the county $\times$ disaster level; the "Shared ZIPs" columns additionally require the prior and focal disasters to share at least 2 IHP-damaged ZIP codes within the focal county. The resulting sample sizes (Whole county / Shared ZIPs) are $886/886$, $852/856$, $769/774$, and $696/701$ for the No drop, 1-year, 3-year, and 5-year rows respectively.
Significance: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

B.15 Robustness to Political Controls and Shared-Representation Interactions

The neighbor-county-flood instrument requires that a neighbor's flood affects own-county outcomes only through the salience channel we document. This could fail if a county and its nearest neighbor share political representation: a common governor, congressional delegation, or partisan lean could let a neighbor's flood route federal aid to the focal county – or move its insurance behavior – through a political channel rather than salience. To test this, Table A22 adds controls for the governor's party, presidential vote share, and congressional-representative party of both the focal and the neighbor county (columns 1–5), and then interactions between them that isolate shared versus split representation (columns 6–8). The spillover coefficient is stable across all eight specifications. Columns (4)–(8) use fewer observations because the county-to-congressional-district crosswalk begins in 2013.

Table A20. Spillover IV: Within-State vs. Cross-State Neighbor Instruments

	Baseline	**Within-state only**	**Cross-state only**
	(1)	**(2)**	**(3)**
Panel A: IV estimates (Neigh. County Flood $\leq$ 5 Years)			
NFIP Policies per 100 Houses	-215.7^{**}	-241.2^{**}	-277.6^{*}
	(72.8)	(75.8)	(117.5)
First-stage F-stat.	81.98	77.47	18.35
Num. Obs.	886	886	886

Notes: Column (1) reproduces the paper's main spillover IV; columns (2) and (3) restrict the neighbor-county flood instrument to neighbors in the same state as the focal county (within-state) or in a different state (cross-state). Significance: $+\, p < 0.10$, $*\, p < 0.05$, $**\, p < 0.01$, $***\, p < 0.001$.

Table A22. Spillover Robustness to Political Controls

	Baseline (1)	+ Gov. (2)	+ Pres. (3)	+ CD (4)	All Main (5)	+ Shared Rep. (6)	+ Unified Control (7)	All Pairwise (8)
Panel A: IV estimates (Neigh. County Flood $\leq$ 5 Years)								
NFIP Policies per 100 Houses	-215.7^{**}	-214.3^{**}	-232.2^{**}	-201.0^{**}	-211.5^{**}	-214.9^{**}	-225.6^{*}	-240.5^{*}
	(72.8)	(66.2)	(75.0)	(69.5)	(67.3)	(69.6)	(99.1)	(109.4)
First-stage F-stat.	81.98	74.66	76.26	77.82	90.38	86.97	75.86	67.81
Panel B: OLS estimates								
NFIP Policies per 100 Houses	-172.8^{***}	-171.1^{***}	-174.7^{***}	-185.5^{***}	-186.1^{***}	-187.4^{***}	-191.6^{***}	-203.4^{***}
	(36.5)	(34.7)	(34.2)	(29.1)	(27.3)	(29.1)	(45.5)	(51.0)
Main-effect controls (focal & neighbour county)								
Governor party indicator		✓			✓	✓	✓	✓
Pres. Dem. vote share			✓		✓	✓	✓	✓
Congressional District party				✓	✓	✓	✓	✓
Interaction terms								
Shared representation, focal $\leftrightarrow$ neigh. (3 terms)						✓	✓	✓
Unified control within county (6 terms)							✓	✓
Mixed cross-county $\times$ cross-level (6 terms)								✓
Num. Obs.	1,268	1,268	1,268	809	809	809	809	809

Notes: Each column reruns the main FEMA spillover IV regression with additional controls. Columns (1)–(5) progressively add political-party controls for the focal and nearest-neighbor county — governor party, presidential Democratic vote share, congressional-district party — and columns (6)–(8) layer on interactions between those measures (shared partisanship across the two counties, unified partisan control within a county, and the remaining cross-county $\times$ cross-level products). Columns (4)–(8) have a smaller sample because the county-to-congressional-district crosswalk starts in 2013. Significance: $+\,p < 0.10$, $*\,p < 0.05$, $**\ p < 0.01$, $***\,p < 0.001$.

B.16 Robustness: Spillover by Disaster Type

Our headline spillover estimate pools across all flood types in the FEMA declarations sample. Table A23 tests whether hurricane-driven spillovers differ from pluvial/fluvial spillovers by including an interaction term in the primary IV regression for hurricane-driven floods. The IV gives a pluvial/fluvial spillover of -310.7 (181.3) and an implied hurricane spillover of -273.4 (106.3); the hurricane-pluvial difference is insignificant.

B.17 Robustness of the Reclassification-Risk Specification

The reclassification-risk projection used in Section 6 maps current and future hazards to RR2.0 premiums via a Lasso regression with linear hazards, squared hazards, and censoring at $1.25\times$ the training 99.9^{th} percentile ("Model 3"). Table A24 reports sensitivity to the three main design choices behind this specification. Panel A compares the three Lasso models we considered against unregularized OLS counterparts (with and without censoring); Panels B, C, and D each vary one dimension of Model 3's specification (censoring level, interaction terms, squared terms) holding the others at the baseline. In general, the choices we made are conservative. If additional interaction terms, or the censoring definition is loosened, future climate risk and future reclassification risk rise sharply.

B.18 Hazard Inputs by Climate Scenario

Table A25 reports the projected 2070 hazard inputs that feed our reclassification-risk premiums under SSP2-4.5 and SSP5-8.5, alongside the 2023 baseline. The pattern across the two scenarios is asymmetric. Cyclone hazards (Panel A) - 100-year rainfall, storm tide, and joint exceedance probability - are markedly less extreme under SSP2-4.5 than SSP5-8.5. Fluvial and pluvial hazards (Panel B), by contrast, are essentially indistinguishable across the two scenarios, and occasionally higher under SSP2-4.5.

This is not an artifact of our prediction algorithm; it is visible in the raw hazard data, before they are passed through any of our prediction steps. It reflects a well-known non-monotonicity of fluvial and pluvial flood hazards in aggregate warming (Hirabayashi et al., 2013; Musselman et al., 2018): certain flood types, such as snow-melt flooding, decline as warming increases, and some dry regions become drier. Because the inland channel barely separates the two scenarios while only the coastal cyclone channel does, the reclassification-risk estimates are correspondingly similar - differing by about 10% across the two scenarios (Table 5) - so the optimal subsidy moves by less than a percentage point.

B.19 Florida to National Mitigation-Elasticity Extrapolation

Our mitigation elasticity E_m is identified from Florida elevation certificate data. Florida is potentially unrepresentative of the rest of the country on several community-level characteristics that plausibly

Table A23. Fiscal Spillover by Disaster Type

	OLS	**IV (Neigh. County Flood $\leq$ 5 Years)**
NFIP Policies per 100 Houses	-188.9^{***}	-310.7^{+}
	(42.2)	(181.3)
NFIP Policies per 100 Houses $\times$ Hurricane-Driven	9.9	37.3
	(19.9)	(82.3)
Implied effect on hurricane-driven events	-179.0^{***}	-273.4^{*}
	(34.6)	(106.3)
Damage control	Binned	
First-stage F-stat., `NFIP`		51.7
First-stage F-stat., `NFIP` $\times$ `Hurr.`		51.3
Num. Obs.	884	884

Notes: Each column reruns the main FEMA spillover regression with an added interaction between NFIP coverage and an indicator for hurricane-driven disasters; the "Implied effect on hurricane-driven events" row sums the main and interaction coefficients. The IV specification instruments both NFIP coverage and its hurricane interaction with the neighbor-county flood instrument (and its hurricane interaction); all specifications are damage-weighted with county-clustered standard errors.
Significance: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

Table A24. Robustness of the Reclassification-Risk Specification

Specification	R^2	Number of Non-Zero Coefficients	Idiosyncratic Risk	Ratio of Future/Current Climate Risk	Mean Insurance Value
Panel A — Model Specifications					
Model 1, LASSO, Linear Coeffs	0.167	53	$1,843	1.82	$235
Model 1, OLS, Linear Coeffs	0.167	53	$1,850	1.82	$237
Model 1, OLS, Linear Coeffs, CR 125% Censor	0.167	53	$1,005	1.48	$81
Model 2, LASSO, Linear + Squared Coeffs	0.176	63	$106,092	14.36	$24,404
Model 2, OLS, Linear + Squared Coeffs	0.176	63	$106,898	15.15	$25,509
Model 2, OLS, Linear + Squared Coeffs, CR 125% Censor	0.177	63	$1,866	2.02	$213
Model 3, LASSO, Linear + Squared Coeffs, CR 125% Censor (Baseline)	**0.177**	**63**	**$1,842**	**1.99**	**$207**
Panel B — Censoring Level					
CR 125% Censor (Baseline)	**0.177**	**63**	**$1,842**	**1.99**	**$207**
CR 150% Censor	0.177	63	$2,517	2.44	$353
CR 200% Censor	0.176	63	$4,132	3.58	$797
No Censor	0.176	63	$106,092	14.36	$24,404
Panel C — Interaction Terms					
0 Interactions (Baseline)	**0.177**	**63**	**$1,842**	**1.99**	**$207**
+3 Cross-type Interactions (Rain $\times$ Surge, etc.)	0.178	65	$1,811	2.01	$202
+6 Compound Interactions (3 Cross + 3 Hurricane $\times$ Pluvial)	0.178	69	$1,862	2.00	$201
+45 All Pairwise Interactions	0.201	108	$4,481	2.30	$1,207
Panel D — Squared Terms					
3 Intensity Squares (Rain, Surge, Max Discharge)	0.172	56	$1,184	1.82	$109
10 Squared Coeffs (Baseline)	**0.177**	**63**	**$1,842**	**1.99**	**$207**
11 Squared Coeffs (incl. Joint Probability^2)	0.177	64	$1,840	1.88	$206

Notes: Robustness of the paper's baseline reclassification-risk model across key design choices. All rows use the SSP2-4.5 projection on $N = 3,529,254$ NFIP policies. Panel A compares the three Lasso models against OLS variants with and without censoring; Panels B–D each vary one dimension of Model 3 (censoring level, interaction terms, squared terms) holding the rest at baseline. Idiosyncratic Risk is the standard deviation of the within-policy premium residual; Mean Insurance Value is the average CARA welfare gain across policies from an $s = 0.5$ subsidy at risk aversion $a = 5 \times 10^{-4}$. Baseline rows are bolded.

Table A25. Projected 2070 Hazards by Climate Scenario

	Current	SSP2-4.5	SSP5-8.5
Panel A: Hurricane Hazards (coastal policies)			
100-yr Rainfall Depth (mm)	176	688	831
100-yr Storm Tide Water Level (mm)	2,473	3,503	3,663
Joint Probability ($\times 100$)	0.32	4.82	6.17
Panel B: Fluvial and Pluvial Hazards (all policies)			
Max River Discharge (m^3/s)	622	839	791
Min River Discharge (m^3/s)	36.91	147.08	137.17
Mean River Discharge (m^3/s)	134	193	178
Total Runoff (mm/day)	1.85	2.31	2.27
Surface Runoff (mm/day)	0.80	1.83	1.80
Subsurface Runoff (mm/day)	0.67	2.26	2.50
Rainfall Flux (mm/day)	4.46	6.04	6.25
Snowfall Flux (mm/day)	0.788	1.790	1.589

Notes: Mean flood-hazard values across NFIP policies, by warming scenario; Panel A's hurricane hazards are restricted to the $\approx$2.6 million coastal policies. "Current" = 2023 hazard inputs (Gori et al. 2021 for hurricanes, ISIMIP2b 1975–2005 baseline for inland); each 2070 column scales those current values by the corresponding model's projected future/historical hazard ratio, averaged over 5–8 climate models per hazard. The SSP2-4.5 hurricane column further rescales the Gori RCP8.5 ratios using Xi–Lin (2023) estimates of how hurricane hazards differ between SSP2-4.5 and SSP5-8.5 (Gori et al. 2021 only ran RCP8.5); the inland SSP2-4.5 column uses ISIMIP2b RCP6.0 as the CMIP5 analog. Hurricane hazards are 100-year return-period intensities; inland hazards are annual maxima.

affect the cost of elevating, so the implied national elasticity may differ from the Florida estimate. To speak to this, we (a) regress the FL-estimated treatment effects (at the NFIP community level) on 18 observables, and (b) reweight the implied elasticity using national values of those observables. Table A26 reports the former. Three covariates have multivariate slopes significant at $p < .10$: tract median household income, claims per policy, and non-SFHA share. Table A27 shows how Florida compares to the rest of the country on those observables. Evaluating the fitted regression at the rest-of-US covariate means yields an extrapolated E_m of approximately 0.075 (compared to the Florida estimate of 0.34), with the three significant covariates accounting for most of the gap. Substituting this extrapolated E_m into the optimal subsidy gives $s^* = 65\%$, versus the Florida baseline of 52%.

B.20 Robustness to Long-Run Behavioral Responses

The mitigation and location elasticities in Sections 7 and 8 are estimated from the short-run RR2.0 price shock. In a fully dynamic model, households who draw bad realizations of the future climate state could undertake additional self-protection in response. To study the welfare implications, we

Table A26. Heterogeneity in the Mitigation Elasticity

Covariate	MV-binary predicted E_m		
	E_m at lo (SE)	E_m at hi (SE)	Δ (SE)
Tract median household income	0.1425 (0.0436)	0.0156 (0.0411)	-0.1268^{*} (0.0589)
Building age (community mean)	0.0627 (0.0402)	0.1489 (0.0505)	0.0862 (0.0673)
Primary-residence share	0.0671 (0.0463)	0.1063 (0.0550)	0.0392 (0.0814)
Claims per policy (community mean)	-0.0406 (0.0603)	0.1240 (0.0311)	0.1646^{**} (0.0592)
Tract residential-mobility rate	0.1419 (0.0495)	0.0511 (0.0408)	-0.0908 (0.0654)
Tract median age	0.1024 (0.0699)	0.0844 (0.0352)	-0.0179 (0.0800)
Tract median tenure (yrs)	0.0981 (0.0359)	0.0631 (0.0584)	-0.0350 (0.0681)
NFIP policy tenure (yrs)	0.1296 (0.0437)	0.0525 (0.0397)	-0.0771 (0.0558)
Basement/crawl/subgrade share	0.1010 (0.0359)	0.0524 (0.0574)	-0.0485 (0.0668)
Single-story share	0.1120 (0.0506)	0.0673 (0.0507)	-0.0447 (0.0806)
Manufactured-home share	0.0924 (0.0361)	0.0770 (0.0523)	-0.0154 (0.0609)
Single-family detached share	0.0976 (0.0359)	0.0462 (0.0864)	-0.0514 (0.0989)
Post-2000 housing share	0.0572 (0.0476)	0.1171 (0.0481)	0.0599 (0.0730)
V-zone share	-0.0036 (0.0836)	0.1204 (0.0401)	0.1239 (0.1034)
Non-SFHA (X-zone) share	0.0397 (0.0457)	0.1805 (0.0597)	0.1407^{+} (0.0854)
CRS class (10 = no discount)	0.0934 (0.0329)	0.0319 (0.0588)	-0.0615 (0.0640)
Median home value	0.0821 (0.0698)	0.0910 (0.0447)	0.0089 (0.0952)
Construction wage (NAICS 236)	0.1478 (0.0522)	0.0505 (0.0447)	-0.0974 (0.0747)

Notes: Heterogeneity in the mitigation elasticity E_m across 18 covariates, from a single multivariate regression on the Florida elevation-certificate sample. The treatment is the post-2021 V/X-shaded reclassification; the outcome is whether the building elevates at least one foot above the base flood elevation (or three feet above grade outside the SFHA). *E_m at lo / E_m at hi* are predicted values with covariate k forced below / above its national median, holding the other 17 fixed; $\Delta = E_m$@hi$-E_m$@lo is the regression slope on the interaction — the heterogeneity test for that covariate.
Significance: $+\, p < 0.10$, $*\, p < 0.05$, $**\ p < 0.01$, $***\, p < 0.001$.

Table A27. Florida vs. Rest-of-U.S. Community Characteristics

Covariate	Unit	FL mean (μ_{FL})	Rest of U.S. mean (μ_{RoC})
Tract median household income	per $10k	$93,929	$95,280
Building age (community mean)	per 10 yrs	35.1	40.8
Primary-residence share	per 10 pp	0.793	0.763
Claims per policy (community mean)	per 0.1 claims	0.1435	0.0244
Tract residential-mobility rate	per 10 pp	0.126	0.118
Tract median age	per 10 yrs	49.9	42.9
Tract median tenure (yrs)	per 5 yrs	9.6	11.0
NFIP policy tenure (yrs)	per 5 yrs	12.0	11.6
Basement/crawl/subgrade share	per 10 pp	0.025	0.162
Single-story share	per 10 pp	0.558	0.552
Manufactured-home share	per 1 pp	0.027	0.044
Single-family detached share	per 10 pp	0.524	0.649
Post-2000 housing share	per 5 pp	0.280	0.280
V-zone share	per 10 pp	0.043	0.020
Non-SFHA (X-zone) share	per 10 pp	0.368	0.488
CRS class (10 = no discount)	per class step	5.44	7.66
Median home value	per $10k	$334,880	$317,219
Construction wage (NAICS 236)	per $100/wk	$1,524	$1,460

Notes: Policy-count-weighted community statistics for the 18 covariates used in the mitigation-elasticity extrapolation, in Florida vs. the rest of the U.S. Sample: 451 Florida communities (0.9M policies) vs. 15,505 elsewhere (2.2M policies). Sources: NFIP FOIA, ACS county housing data, BLS QCEW (construction wages), and the FEMA CRS Community Status Book.

re-solve for the optimal subsidy under a long-run extrapolation: we scale the mitigation and location elasticities up by multipliers κ_{Em} and κ_{Ex}, on the assumption that households respond to the 2070 climate-driven price changes with the same elasticity they exhibit toward the short-run RR2.0 price changes we measure.

Table A28 reports the result. Under this dynamic extrapolation the elasticities scale up by roughly 40–70% relative to the short-run baseline, and the optimal subsidy falls by 4–5 percentage points, from 52% to 47% under SSP2-4.5 and to 48% under SSP5-8.5. The static framework therefore overstates the welfare case for a subsidy only modestly. This rescaling implicitly allows households who experience large risk shocks to be those most likely to relocate or mitigate, thereby reducing the value of a subsidy that dulls those incentives. Third-order feedback effects — subsidies dulling the impact of price changes on mitigation and relocation, which in turn add a small amount of reclassification risk — are quantitatively negligible.

B.21 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 benefits to assessed damage. We do this separately by damage level to understand if underinsurance covaries with flood severity. The results are in Figure 14 below.

Figure 14. Insurance Coverage Relative to Flood Damage

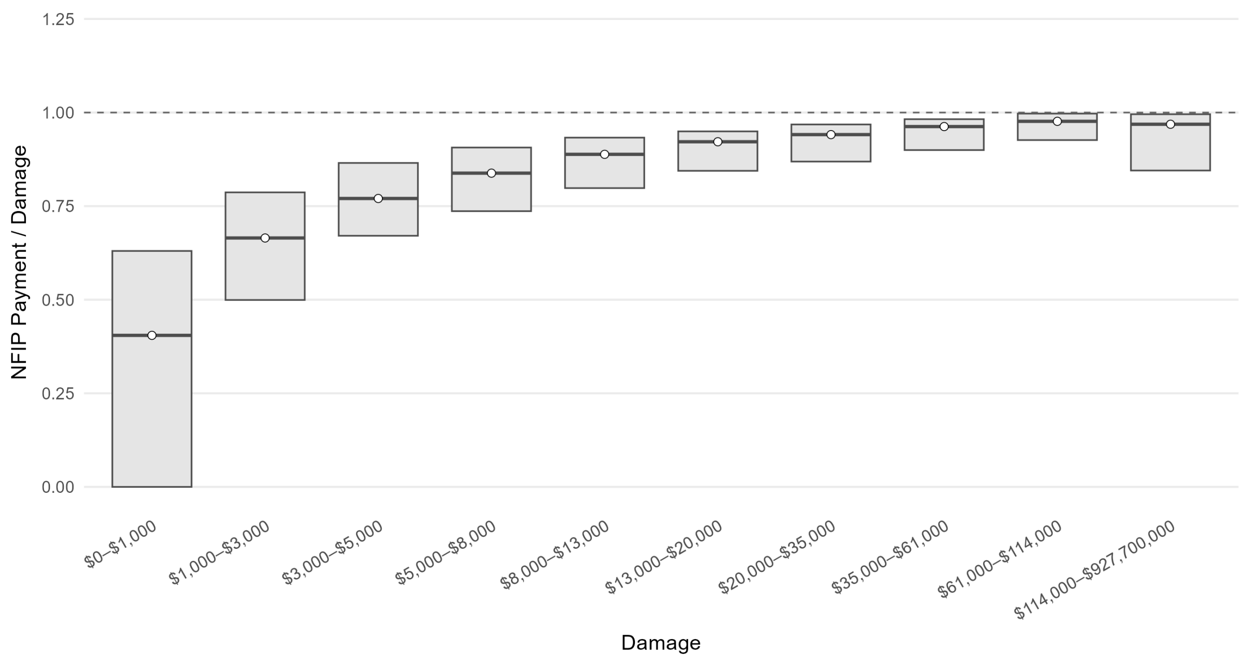

Notes: Each box summarizes the distribution of the claim-level insurance payment per dollar of damage. Boxes show the inter-quartile range; the bold horizontal segment marks the median. The dashed horizontal line at 1.0 indicates full coverage.

Figure 14 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.22 Behavioral Misperception Wedges (Internality Corrections)

Here we derive a quantitative measure, in the spirit of Section 9, of the welfare effects of two behavioral internalities discussed in Section 3.1: household under-perception of flood risk and household over-perception of FEMA disaster aid. Both biases lower demand-relevant willingness-to-pay for NFIP coverage relative to its welfare-relevant unbiased value.

Setup. Consider a Bernoulli flood-loss risk: household i experiences gross loss L_f with probability p_i and zero otherwise. FEMA disaster aid covers fraction f of the gross loss when the household is uninsured; the residual $(1-f)L_f$ is the loss against which NFIP insures.[44] Under CARA utility with coefficient γ, a small-risk approximation implies that the certainty-equivalent value of NFIP coverage is

$$V_i^*(p_i, f) \;\approx\; p_i(1-f)L_f \;+\; \frac{\gamma}{2}\, p_i(1-p_i)(1-f)^2 L_f^2. \tag{8}$$

We take $L_f = \$29{,}267$ and $\bar{P} = \$1{,}739$, which implies a PDD-support flood probability $p = \sigma_{\text{PDD}}\bar{P}/L_f \approx 0.055$. The spillover regression of Section 4 identifies the per-insured per-year FEMA-aid offset as $|\beta| = p\, f\, L_f$, so $f = |\beta|/(\sigma_{\text{PDD}}\bar{P}) = 216/(0.93 \cdot 1739) \approx 0.133$, implying that FEMA covers approximately \$3,900 per affected uninsured household. A premium subsidy that reduces log prices by $\Delta \ln \tilde{P}$ induces additional coverage of $\Delta q \simeq \eta^q\, \Delta \ln \tilde{P}$. For any wedge Δ_i^B between welfare-relevant and decision-relevant value, the welfare gain per house-year is

$$\Delta W^B \;\approx\; \overline{\Delta^B} \;\cdot\; \eta^q \;\cdot\; \Delta \ln \tilde{P}, \tag{9}$$

where $\overline{\Delta^B}$ is the wedge averaged across households marginal to the subsidy.

B.22.1 Misperception of Flood Risk

Suppose households under-perceive the flood probability, evaluating insurance as if the probability were $\bar{p}_i = k_R p_i$ with $k_R < 1$. Define the wedge

$$\begin{aligned} \Delta_i^{B,R} &\equiv V_i^*(p_i, f) - V_i^*(\bar{p}_i, f) \\ &\approx (1-k_R)(1-f)P_i \;+\; \tfrac{\gamma}{2}(1-f)^2 L_f^2\big[p_i(1-p_i) - k_R p_i(1-k_R p_i)\big], \end{aligned} \tag{10}$$

where we have used actuarial fairness $p_i L_f = P_i$ for the linear term.

Calibration. Following Bakkensen and Barrage (2022), we set $k_R = 0.57$ (a 43% under-perception), and assume our headline coefficient of absolute risk aversion $\gamma = 5 \times 10^{-4}$. Eq. (10) yields $\Delta^{B,R} \approx \$4{,}461$ per newly insured household-year, and the implied internality benefit per house-year for a 1% log price cut is approximately \$14.38.

B.22.2 Misperception of FEMA Generosity

Suppose households over-perceive FEMA generosity, evaluating the share of loss FEMA will cover as $\tilde{f} = k_F f$ with $k_F > 1$.[45] Holding p_i at its true value, the wedge is

$$\begin{aligned} \Delta_i^{B,F} &\equiv V_i^*(p_i, f) - V_i^*(p_i, \tilde{f}) \\ &\approx (k_F - 1)\, f\, P_i \;+\; \tfrac{\gamma}{2}\, p_i(1-p_i)L_f^2\big[(1-f)^2 - (1-k_F f)^2\big]. \end{aligned} \tag{11}$$

[44] We treat NFIP as fully covering the residual loss conditional on take-up, consistent with the empirical verification of (effectively) complete coverage on the intensive margin in Appendix B.21.

[45] Section 3.1 cites Kousky and Shabman (2012) for the qualitative claim that households mistakenly believe FEMA is more generous than it actually is. Kousky et al. (2018) document an aid-to-insurance crowd-out ratio of \$1.3–\$1.7 per realized FEMA dollar, broadly consistent with k_F in this range.

Both terms are positive when $k_F > 1$: over-perception of FEMA inflates the perceived no-NFIP residual loss reduction, depressing decision-relevant NFIP value below welfare-relevant value.

Calibration. Bakkensen and Barrage (2022) directly elicited household expectations of government assistance as a percentage of flood damages in their Rhode Island survey ($n = 187$), reporting a sample mean of approximately 12% — close to the empirically calibrated $f = 13.3\%$ implied by our spillover regression. Their finding is consistent with k_F near unity in that sample, though small sample size and wide standard errors leave room for moderate over-perception. The qualitative literature (Kousky and Shabman, 2012) and revealed-preference evidence on disaster-aid crowd-out (Kousky et al., 2018) both indicate over-perception as the typical direction. We adopt $k_F = 1.5$ as a conservative calibration consistent with this qualitative consensus. At $k_F = 1.5$ and our headline $\gamma = 5 \times 10^{-4}$, Eq. (11) yields $\Delta^{B,F} \approx \$1{,}362$ per newly insured household-year, and the implied internality benefit per house-year for a 1% log price cut is approximately $4.39.

B.22.3 Joint Wedge and Robustness of the Optimal Subsidy

When both biases are operative, the welfare-relevant minus decision-relevant value is

$$\Delta_i^{B,J} \equiv V_i^*(p_i, f) - V_i^*(\bar{p}_i, \tilde{f}). \tag{12}$$

At the central calibrations of Sections B.22.1 and B.22.2, the joint wedge is approximately $5,257 per newly insured household-year, comprising the $4,461 risk wedge plus the $1,362 FEMA wedge less a $566 negative cross-term.

Table A29 reports the pooled optimal subsidy s^* under each combination of misperception assumptions, across the same five-point CARA risk-aversion grid γ used in Table 5. The no-misperception benchmark itself rises with γ (from 49% at $\gamma = 5 \times 10^{-5}$ to 59% at $\gamma = 5 \times 10^{-3}$), reflecting that the reclassification-risk welfare value scales with risk aversion. The risk-misperception channel is the dominant driver: at the central calibration $k_R = 0.57$ it raises s^* to 89%–100% depending on γ, whereas FEMA misperception alone ($k_F = 1.5$) has a more modest effect (58%–100%); the joint cell at $(k_R = 0.57, k_F = 1.5)$ delivers $s^* \in [94\%, 100\%]$. Halving each wedge's departure from the no-misperception point ($k_R = 0.785$, $k_F = 1.25$) attenuates these effects — to 68%–100% for risk misperception, 54%–100% for FEMA misperception, and 72%–100% jointly — while leaving the qualitative ranking intact.

Table A29. Robustness of optimal subsidy s^* to behavioral misperception wedges

Misperception scenario	CARA risk-aversion coefficient γ				
	5×10^{-5}	1×10^{-4}	2×10^{-4}	5×10^{-4}	5×10^{-3}
No misperception	49%	50%	51%	52%	59%
Half risk misperception ($k_R = 0.785$)	68%	76%	90%	100%	100%
Half FEMA misperception ($k_F = 1.25$)	54%	57%	63%	78%	100%
Half both misperceptions ($k_R = 0.785$, $k_F = 1.25$)	72%	82%	100%	100%	100%
Full risk misperception ($k_R = 0.57$)	89%	100%	100%	100%	100%
Full FEMA misperception ($k_F = 1.5$)	58%	64%	74%	100%	100%
Full both misperceptions ($k_R = 0.57$, $k_F = 1.5$)	94%	100%	100%	100%	100%

Notes: Each cell reports the pooled optimal subsidy s^* under the joint CARA-Bernoulli wedge of Eq. (12). All cells hold $\eta = 0.32$, $\lambda_\beta = 1.0$, and $f = |\beta|/(\sigma_{\text{PDD}}\bar{P}) = 0.133$. Misperception calibrations: $k_R = 0.57$ from Bakkensen and Barrage (2022); $k_F = 1.5$ as a central calibration informed by the direct elicitation in Bakkensen and Barrage (2022) (sample mean $\approx 12\%$ expected coverage, $n = 187$) and the revealed-preference crowd-out evidence in Kousky et al. (2018), with sensitivity bracket $[1.0, 1.5]$. The "Full" rows apply these central calibrations; the "Half" rows halve each wedge's departure from the no-misperception benchmark, i.e. $k_R = 0.785$ and $k_F = 1.25$.

B.23 Income Distribution of Net Gains from a Subsidy

Figure 15 bins NFIP policyholders by tract median income, and within each bin reports the average net gain or loss from a simulated 1% premium subsidy funded by a lump-sum tax on all insured households. The qualitative point—a relatively weak positive association between income and net gains, which could be reversed under progressive financing—is preserved.

Figure 15. Net Gains/Losses by Income Bin, 1% Subsidy Funded by Lump-Sum Tax

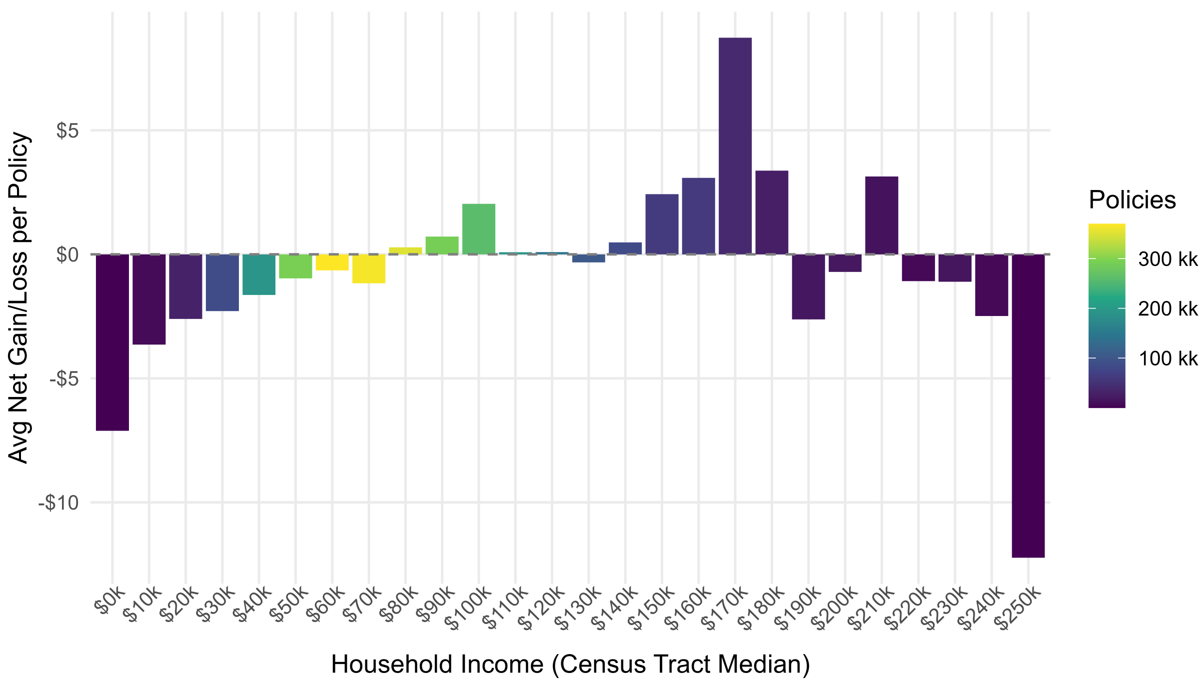

Notes: Bar chart of average net gain/loss by income bin (2023 dollars) from a simulated 1% subsidy on NFIP premiums funded by lump-sum tax on all insured households. The color scale represents the number of policies (sample size) per bin. The simulation assigns tract-median income from the ACS to each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

B.24 Private Flood Market Crowdout

Our subsidy prescription applies to all flood insurance, public or private. The historical NFIP-only subsidy regime reflects its status as the sole flood-insurance provider for most of its history. This is not the policy we propose. Nonetheless, an NFIP-only subsidy could in principle crowd out a nascent private market: a household induced into the NFIP rather than into private coverage forfeits the difference in coverage limits, which carries a welfare cost. This appendix quantifies that channel in two steps. First, we estimate the crowd-out empirically from state-level variation as RR2.0 has rolled out. Second, we compute the welfare loss from each crowded-out household forfeiting above-cap private coverage and propagate it into the optimal-subsidy formula.

Crowd-out estimation. We use state-level variation in Risk Rating 2.0 (RR2.0) premium shocks. Private flood policy counts come from NAIC data-call annual filings (2018–2024); NFIP counts come from OpenFEMA. We regress the state-level change in log private First-Dollar Standalone policies (2022–23 average minus 2018–20 average) on the state's policies-in-force-weighted RR2.0 relative gap, weighted by pre-period NFIP policies.

Table A30 reports multiple specifications. The full-sample yields a noisy null. Various sample restrictions (controlling for hurricane-hit states and dropping Hawaii which was a sharp outlier in

RR2.0 price changes) and additional controls (realized damage and region fixed effects) generates, at most a 20% crowd out. We use this as a conservative upper bound on crowdout, but note that generally the effect is indistinguishable from zero.

Coverage-limit welfare wedge. The primary welfare loss when a household is diverted from private to public flood insurance by a (hypothetical) NFIP-only subsidy is that NFIP has a sharp coverage limit. The NFIP caps coverage at $250,000 per building and $100,000 per contents, while private policies routinely write above these limits. A household crowded out of a private contract therefore loses protection for flood damages that exceed this cap. Under CARA preferences with absolute risk aversion γ, the certainty-equivalent value of insurance against this above-cap loss is

$$WTP^{\text{cap}} \;\approx\; pL \;+\; \tfrac{\gamma}{2}\, p(1-p)\, L^2,$$

where p is the per-policy-year probability of exceeding the cap and L the conditional mean shortfall. We estimate p and L from the NFIP claims file, evaluating the building and contents wedges separately and summing. At $\gamma = 5 \times 10^{-4}$ this gives a per-policy-year wedge of approximately $\$147$ under current flood risk. Under a 2070 projection in which p and L scale jointly to match the cross-year decomposition of log(pL) variance observed in the claims data, the wedge rises to approximately $\$1{,}022$.

Effect on the optimal subsidy. Because each marginal displaced household forfeits a fixed WTP^{cap} regardless of the subsidy level, the crowd-out welfare cost enters the optimal-subsidy formula linearly in s. This shifts the linear-in-s FEMA-savings coefficient down by $\eta\,\theta\, WTP^{\text{cap}}$. Re-solving for s^* as we vary the degree of crowd-out and the insurable risk above the coverage cap yields Table A31. Under current flood risk and the central crowd-out estimate ($\theta = 0.11$, $WTP^{\text{cap}} = \$147$) the optimal subsidy falls by less than 1 percentage point from its baseline of 52%. Under the most adversarial calibration considered ($\theta = 0.20$, $WTP^{\text{cap}} = \$1{,}022$) it falls by 6 percentage points to 46%. The corresponding row of Panel D of Table 6 combines this calibration with the other adversarial assumptions explored in Section 9.

B.25 FEMA-IHP Approval Selection on Pre-Flood NFIP Coverage

The paper's spillover estimates are identified from county-disasters in which FEMA's Individuals and Households Program (FEMA-IHP) was authorized by presidential declaration. A concern is that the IHP-declaration decision itself selects on insurance coverage — counties with more NFIP take-up could be less likely to be approved, biasing the sample. This appendix tests for that selection directly.

Data construction. Approved county-disasters are the paper's 886-row main subsample. From FEMA's preliminary damage assessment records, which document the property-damage estimates

Table A28. Optimal Subsidy under Long-Run Behavioral Responses

Scenario	κ_{Em}/κ_{Ex}	s^* (FOC only)
Paper baseline (RR2.0 short-run, SSP2-4.5 anchor)	1.00 / 1.00	52%
Long-run extrapolation, SSP2-4.5	1.43 / 1.69	47%
Long-run extrapolation, SSP5-8.5	1.41 / 1.72	48%

Notes: Optimal subsidy s^* recomputed with the mitigation and location elasticities scaled by long-run multipliers (κ_{Em}, κ_{Ex}), which assume households respond to the 2070 climate-driven price changes with the same elasticity as they do to the RR2.0 price changes.

Table A30. Private Flood Insurance Take-Up: State-Level Evidence

	All 51	**Drop Gulf 5**	**Drop Gulf 5 + HI**
	(1)	**(2)**	**(3)**
Outcome: Change in Private Flood Insurance Policies, 2022-23 − 2018-20			
Bivariate	-0.107	0.113	-0.015
	(0.134)	(0.091)	(0.165)
+ log$(1+$ damage$_{2020\text{-}22})$	0.089	0.126	-0.024
	(0.149)	(0.098)	(0.201)
+ Region FE + log$(1+$ damage$)$	0.074	0.196^{*}	0.113
	(0.149)	(0.085)	(0.182)
Num. Obs.	51	46	45

Notes: State-level regression of private flood insurance take-up on NFIP price changes due to RR2.0. The regressor is the policy-weighted mean across a state's counties of the post-RR2.0 premium relative to the pre-RR2.0 premium; the outcome is the change from 2018-20 to 2022-23 in (the log of) private flood-only policies in force. The second row controls for realized flood damage from 2020-22; the third adds regional fixed effects. The second column drops hurricane-exposed states (LA, FL, TX, MS, AL); the third drops Hawaii, which had substantially larger RR2.0 price increases than any other state. Regressions are weighted by pre-period NFIP policies. Significance: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

that accompany each IHP request, we identify requests where IHP assistance was requested but no declaration followed, restricting to 2010–2022 county-disasters with reported property damage above zero. This yields 36 denied county-disasters.

Summary statistics. Panel B of Table A32 compares pre-flood NFIP coverage between approved and denied county-disasters. Approved counties have substantially higher pre-flood coverage: a mean of 4.8% versus 2.3% in denied counties. This is the opposite of a selection story in which assistance flows to less-insured places; if anything, approvals concentrate in higher-coverage counties. The same panel compares property damage: approved county-disasters have a mean damage of $82.8 million versus $6.2 million for denied, consistent with the IHP-declaration decision being primarily driven by event severity.

Linear-probability regression. Panel A reports a damage-weighted linear-probability regression of approval on pre-flood NFIP coverage, varying the fixed-effect structure. Across all five specifications - the coefficient on NFIP coverage is statistically insignificant and quantitatively tiny. Approval is essentially driven by damage, not by coverage.

Selection within the damage band that drives the headline estimate. As additional evidence that the IHP-approval margin is immaterial to our findings, Table A33 re-runs the paper's primary IV specification on three damage subsamples. Events with at least $100 million in property damage account for approximately 98% of total damage in the paper's sample and drive the damage-weighted IV estimate (-241.6 on 509 observations, similar to the headline of -215.7). All 36 denied county-disasters fall below the $100 million cut, so selection on the declaration decision cannot operate within the damage band that identifies the headline result. Put another way, no disaster with more than $100 million dollars of damage has ever been denied IHP, and those are the disasters that constitute 98% of our sample.

B.26 NFIP Claim Payouts Outside FEMA-Approved Disasters

We compute the share of NFIP claim payouts that fall outside any FEMA Individual-Assistance declaration window. Table A34 categorizes every NFIP claim by the disaster-declaration status of its county-disaster. NFIP paid out roughly $2.5 billion across 166,000 claims in flood events that never received a presidential disaster declaration, plus a further $26 million in counties where an assistance request was denied. In these cases FEMA household aid was unavailable, and insurance was the only tool for households to recover losses. This makes NFIP subsidies even more attractive to the planner.

Table A31. Optimal Subsidy s^* with Private-Market Crowd-Out

Scenario	Crowdout	Value of Insurance Above NFIP Cap	s^*
Paper baseline (no crowd-out)	0.00	$0	52%
Today, central	0.11	$147	52%
Today, upper	0.20	$147	52%
2070 worst-case, central	0.11	$1,022	49%
2070 worst-case, upper	0.20	$1,022	46%

Notes: Optimal subsidy s^* recomputed under five scenarios that vary the degree of crowd-out and the welfare loss from risk left uninsured above the NFIP coverage cap, measured either today (rows 2–3) or in 2070 (rows 4–5).

Table A32. FEMA-IHP Aid Approval and Pre-Flood NFIP Coverage

Panel A: *Linear-Probability Regression of FEMA-IHP Approval on Pre-Flood NFIP Coverage*

	Dependent variable: FEMA-IHP declared (= 1 if approved)				
	(1) no FE	(2) State FE	(3) Year FE	(4) State + Year FE	(5) Disaster FE
NFIP Policies per 100 Houses	$-1.1e-05$	$-1.1e-05$	$-1.9e-05$	$8.9e-05$	$5.4e-07$
	(7e-05)	(3.3e-05)	(6.9e-05)	(9.1e-05)	(1.1e-06)
State FE		✓		✓	
Year FE			✓	✓	
Disaster FE					✓
R^2	0.023	0.146	0.051	0.202	0.582
Num. Obs.	922	917	922	917	893

Panel B: *NFIP Coverage and Property Damage, FEMA-IHP-Approved vs. FEMA-IHP-Denied Counties*

Group	**N**	**Mean cov. (%)**	**Median cov. (%)**	**SD cov. (%)**	**Mean damage ($M)**	**Median damage ($M)**
approved	886	4.783	1.648	9.09	82.81	1.05
denied	36	2.286	1.083	3.50	6.24	1.00

Notes: Panel A reports linear-probability regressions of FEMA's Individual-and-Households-Program (FEMA-IHP) approval decision on county-level NFIP coverage. All specifications include polynomial damage controls, are damage-weighted, and cluster standard errors at the county level, with column (5) comparing counties hit by the same declared disaster. Panel B compares pre-flood NFIP coverage and property damage between FEMA-IHP-approved county-disasters (the paper's 886-row main subsample) and FEMA-IHP-denied county-disasters (a governor's request appears in FEMA's preliminary-damage-assessment records but no declaration followed). Significance: $+\,p < 0.10, *\,p < 0.05, **\ p < 0.01, ***\,p < 0.001$.

Table A33. Spillover Estimate by Damage Subsample

Damage subsample	OLS	IV	First-stage F	Num. Obs.	Share of damage
All (paper Table 2 sample)	-172.8^{***}	-215.7^{**}	81.98	886	100%
	(36.5)	(72.8)			
Damage $\geq$ \$100 M	-233.2^{***}	-241.6^{***}	78.81	509	$\approx$ 98%
	(31.6)	(65.9)			
Damage $<$ \$100 M	-77.5^{***}	-1195.0	0.10	377	$\approx$ 2%
	(20.3)	(5766.7)			

Notes: The paper's main IV spillover specification, re-run on the full sample and on county-disasters above and below \$100M in flood property damage. Significance: $+\ p < 0.10$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

Table A34. NFIP Claims and Payouts by Disaster-Declaration Status

Group	N claims	Total NFIP paid (\$M)	Share of total \$	Mean per claim	Median per claim
NFIP Claims from Approved PDD	$604{,}075$	33,638	93.1%	\$55,685	\$23,810
NFIP Claims from Denied PDD	936	26.4	$<$ 0.1%	\$28,168	\$12,697
NFIP Claims with no PDD	$165{,}707$	2,462	6.8%	\$14,858	\$4,284
Total	$770{,}718$	36,127	100.0%	\$46,874	—

Notes: NFIP claims, categorized by disaster-declaration status: claims inside an approved FEMA Individual Assistance declaration window for their county ("Approved PDD"); claims inside a county-disaster window where IA was requested but not granted, per FEMA's preliminary-damage-assessment records ("Denied PDD"); and claims with no presidential disaster declaration ("No PDD").

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:*

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

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 left-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 left-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 define 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 side of the optimality equation in Propositions 2 and 1 is augmented by:*

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

Binary Coverage With binary coverage (either no insurance or full insurance), the mitigation and location internalities are zero: $\Delta_i^m = \Delta_i^x = 0$. We empirically verify that coverage is essentially complete in Appendix B.21.

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

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

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*

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

Proof. By definition,

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

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

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

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

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

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

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

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

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

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

The planner's per–state budget is $\sum_i T(\omega) = \text{Out}(\omega; s),$ 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 $\text{Out}(\omega; s) = \sum_i sq_iP_i(\omega) + \sum_i f(1-q_i)P_i(\omega)$, we have $\partial\text{Out}/\partial s = \sum_i q_iP_i(\omega)$.

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

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

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

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

For the second term:

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

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 and 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

$$\frac{d\mathcal{L}}{ds} = \underbrace{D}_{\text{Step 1}} + \sum_i \sum_{a\in\{q,m,x\}} \underbrace{\frac{\partial\mathcal{L}}{\partial a_i}}_{\text{FOC gap}} \cdot \frac{da_i}{ds}. \tag{A.2}$$

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:

$$\frac{\partial\mathcal{L}}{\partial a_i} = -\mathbb{E}_\omega\Big[\mu(\omega)\,\frac{\partial \mathsf{Out}}{\partial a_i}\Big]. \tag{A.3}$$

Compute the outlay derivatives (per state):

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

$$\frac{\partial \mathsf{Out}}{\partial x_i} = s\,q_i\,\frac{\partial P_i}{\partial x_i}(\omega) + f(1-q_i)\,\frac{\partial P_i}{\partial x_i}(\omega). \tag{A.4}$$

Express da_i/ds via semi–elasticities:

$$\frac{da_i}{ds} = \frac{\partial a_i}{\partial \ln \mathbb{E}_\omega[\tilde{P}_i]}\cdot\frac{d\ln \mathbb{E}_\omega[\tilde{P}_i]}{ds} = -\frac{1}{1-s}\,\mathcal{E}_i^a. \tag{A.5}$$

The behavioral contribution is

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

Insert (A.5):

$$B = \frac{1}{1-s}\,\mathbb{E}_\omega\Big[\mu(\omega)\Big(\sum_i ((s-f)\,P_i(\omega))\,\mathcal{E}_i^q + \sum_i (s\,q_i + f(1-q_i))\,\frac{\partial P_i(\omega)}{\partial m_i}\,\mathcal{E}_i^m + \sum_i (s\,q_i + f(1-q_i))\,\frac{\partial P_i(\omega)}{\partial x_i}\,\mathcal{E}_i^x\Big)\Big]. \tag{A.6}$$

Setting $\frac{d\mathcal{L}}{ds} = 0$ at the optimum yields $D + B = 0$. From (A.6), $B = \frac{1}{1-s}\tilde{K}$, where $\tilde{K} = \mathbb{E}_\omega\Big[\mu(\omega)\Big(\sum_i ((s-f)\,P_i(\omega))\,\mathcal{E}_i^q + \sum_i (s\,q_i + f(1-q_i))\,\frac{\partial P_i(\omega)}{\partial m_i}\,\mathcal{E}_i^m + \sum_i (s\,q_i + f(1-q_i))\,\frac{\partial P_i(\omega)}{\partial x_i}\,\mathcal{E}_i^x\Big)\Big]$. Thus, $D = -\frac{1}{1-s}\tilde{K}$. 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, $\text{Cov}_i\,(q_i^*, u_i'(\omega)) = 0$; 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.

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,

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

Rewriting this yields

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

Hence

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

so the "average selection" and "ω-specific selection" terms vanish.

By A2, for each ω,

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

Using the definition of covariance, A3 implies

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

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 $\text{Cov}_\omega(\mu(\omega), g(\omega)) = 0$. Hence the RHS becomes

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

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.[46]

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

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

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

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

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 $(1-$

[46] 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.

$\lambda_i)\mathbb{E}[L_i]$. 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

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

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

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

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 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

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

as required.