

THE CLIMATE PROTECTION GAP AND PRIVATE SELF-INSURANCE

PAULINE CHIKHANI

ABSTRACT. This paper studies how climate tail risk affects household wealth accumulation when the primary saving asset is itself exposed to destruction. I embed an empirically estimated census-tract-level process for disaster-induced depreciation into a heterogeneous-agent model with incomplete markets. Because wealth simultaneously serves as self-insurance and the object at risk, disaster exposure weakens precautionary saving incentives and reshapes the wealth distribution. Quantitatively, aggregate effects are modest: empirically calibrated disaster risk reduces aggregate capital by less than 1 percent. Distributional effects, however, are substantial. Households exposed to greater climate risk accumulate less wealth, are more likely to remain hand-to-mouth, and experience greater financial fragility. The share of hand-to-mouth households increases by roughly 30 percentage points, while wealth inequality rises markedly. These findings suggest that climate tail risk operates primarily through a distributional channel rather than through large aggregate contractions.

JEL: Q54, D91, E21, D52, D31.

Keywords: Climate risk, capital destruction, precautionary saving, self-insurance, incomplete markets, wealth inequality.

This version: May 2026.

UNIVERSITY OF LAUSANNE (HEC), DEPARTMENT OF ECONOMICS, LAUSANNE, CH-1015, SWITZERLAND

E-mail address: `pauline.chikhani@unil.ch`.

I am grateful to Aurélien Eyquem, Luca Gemmi, Sébastien Houde, Fidel Petros, Jean-Paul Renne, Adrien Tschopp, and Gauthier Vermandel for helpful discussions and comments. I also thank seminar and conference participants at the University of Lausanne, the CMAP-CREST QSEF Seminar, and the CESifo Venice Summer Institute workshop on Green Energy Transition and the Role of Energy in the Macroeconomy for valuable feedback. All remaining errors are my own.

I declare that I have received no external funding for this research and have no financial relationships or conflicts of interest to disclose.

1. INTRODUCTION

Global temperatures have continued to rise in recent years: 2025 ranked as the third warmest year on record, following 2024 and 2023 ([Berkeley Earth, 2026](#)). Yet beyond the rise in global average temperatures lies substantial spatial heterogeneity. Some regions already experience temperature increases of up to 4°C relative to their 1951–1980 baseline, and exposure to climate-related damages varies markedly not only across countries but also within them. While coordinated global mitigation would be the first-best response, international cooperation remains limited, in part because the largest emitters are not necessarily those bearing the greatest immediate damages ([IPCC, 6th Assessment report](#)). In this environment, adaptation and risk management increasingly fall on national governments, local communities, and private agents.

In principle, insurance and financial markets could allow these risks to be shared more broadly. In practice, however, a substantial fraction of disaster losses remains uninsured. The reinsurance industry refers to the share of economic losses not covered by insurance as the “climate protection gap” ([Swiss Re Institute, 2025](#)). In the United States, this gap is estimated at roughly 43% and has widened over recent decades. Consequently, a large share of climate-related losses ultimately falls on household balance sheets rather than being transferred to insurance markets. Precautionary saving, therefore, represents the most basic form of climate-risk management and the first line of defense against the climate protection gap. Understanding how households self-insure is a necessary step before studying more sophisticated forms of risk sharing through financial markets, insurance contracts, or public intervention.

This paper studies the first layer of the climate protection gap: self-insurance through wealth accumulation when wealth itself is exposed to disaster risk. I develop a heterogeneous-agent model with incomplete markets and borrowing constraints in which households face location-specific destruction shocks that directly reduce accumulated wealth. The key departure from standard precautionary-saving environments is that the asset used for self-insurance is also the asset exposed to destruction. Wealth, therefore, plays a dual role: it provides a buffer against future losses, but it also determines the extent to which it can be destroyed. This interaction creates a trade-off between insurance and exposure, so that greater disaster risk need not generate uniformly higher precautionary saving.

The model introduces spatial heterogeneity through climate-risk zones that differ in the distribution and persistence of destruction shocks. Households face common aggregate

prices and identical current economic conditions, but differ in their exposure to climate risk. Climate risk affects household decisions not only through realized losses but also through expectations about future destruction. This structure makes it possible to isolate how differences in the distribution of future disaster risk affect saving behavior and wealth accumulation.

To discipline these destruction processes, I construct a synthetic census-tract-level panel of disaster-induced depreciation rates using disaggregated U.S. county-level natural-disaster loss data for 2002–2022. County-level aggregation substantially smooths localized destruction, so I allocate observed county-year losses across census tracts using tract-level exposure shares to recover the distribution of local tail risk. I then summarize the resulting spatial heterogeneity into two climate-risk zones, each characterized by a distinct distribution of destruction shocks. These estimated processes discipline the quantitative model and allow me to study how localized climate tail risk reshapes household saving behavior and the distribution of wealth.

The model delivers three main results. First, climate risk has a non-monotonic effect on precautionary saving incentives because wealth simultaneously provides insurance and determines exposure to destruction. At low wealth levels, the exposure effect dominates: households in high-risk locations save less because future assets are themselves vulnerable to loss. As wealth increases, the insurance value of additional assets becomes more important and precautionary motives re-emerge. Consequently, saving responses to climate risk are highly nonlinear across the wealth distribution. Quantitatively, households in high-risk zones save less than those in safer areas at low wealth levels, but can save up to 4 times as much at intermediate wealth levels.

Second, the aggregate effects of empirically calibrated climate tail risk are modest, but the distributional effects are substantial. Introducing rare but severe depreciation shocks reduces steady-state aggregate capital by approximately 0.9%. However, the mass of households subject to the borrowing constraint rises sharply, from roughly 5% to about 35% of the population. Climate tail risk, therefore, operates primarily through redistribution and compression of the lower tail of the wealth distribution rather than through large aggregate contractions.

Third, geographic aggregation provides substantial implicit risk pooling. When losses are measured at the county level, the mean annual disaster-induced depreciation rate over 2002–2022 is only about 0.005 percent per county-year. This apparent smoothness masks

the localized nature of destruction: in affected tracts, households can lose a substantial fraction of their assets, and the probability of such losses varies sharply across locations. Aggregate data, therefore, understate the intensity of risk faced by exposed households and the behavioral responses it generates.

Taken together, the results suggest that policy can play a central role in reducing households' exposure to climate-related disaster risk. When households rely mainly on resources that are themselves exposed to local destruction, self-insurance remains limited: the same shock that creates the need for liquidity can also reduce households' ability to draw on accumulated wealth. A natural policy priority is therefore not only to expand the set of instruments that transfer local disaster risk away from household balance sheets, but also to reduce the barriers that prevent households from using existing ones. These barriers may include high premia, limited attention, low financial literacy, lack of trust, or imperfect information about coverage. Policies that improve affordability, transparency, and the take-up of insurance and related risk-pooling mechanisms can therefore have effects similar to those of introducing safer claims. This is especially relevant for households with limited wealth, for whom insufficient protection against climate risk can make losses harder to absorb and reinforce inequality across locations and income groups.

2. RELATED LITERATURE AND CONTRIBUTIONS

Empirical Evidence on Disaster Exposure and Household Adjustment. A large empirical literature documents substantial heterogeneity in household outcomes following natural disasters, even when aggregate economic effects remain relatively modest. At the household level, disasters disrupt local labor markets through business closures, infrastructure damage, residential destruction, and job separations. Using matched employer-employee data for Hurricanes Katrina and Rita, [Groen et al. \(2020\)](#) document sizable short-run earnings losses among directly exposed households, particularly for workers experiencing severe residential damage. [Gallagher and Hartley \(2017\)](#) similarly find persistent deterioration in household financial conditions following flood exposure, with larger effects among financially constrained households.

At the aggregate level, however, many disasters appear to generate comparatively limited long-run macroeconomic effects. [Boustan et al. \(2020\)](#) show that migration, reconstruction, and local adaptation often mitigate aggregate consequences of U.S. natural

disasters. Nevertheless, severe disasters can produce persistent reductions in local economic activity and productive capacity (Hsiang and Jina, 2014; Strobl, 2011). Moreover, incomplete insurance coverage and imperfect public compensation imply that households remain partially exposed to uninsured destruction risk (Gallagher, 2014; Bakkensen and Barrage, 2022).

This paper reconciles these two empirical facts: disasters often generate limited aggregate effects while simultaneously producing persistent financial vulnerability among exposed households. I show that localized destruction risk can weaken precautionary saving and increase borrowing constraints even when aggregate fluctuations remain modest.

Climate Risk in Heterogeneous-Agent and Spatial Models. A growing literature studies climate risk in heterogeneous-agent and spatial macroeconomic environments. Fried (2022) embeds storm risk and federal disaster aid in a regional Aiyagari model, emphasizing adaptation incentives and redistribution across risk zones. Fried (2026) extends this framework to rising temperature exposure and heterogeneous income dynamics. In spatial equilibrium settings, Bilal and Rossi-Hansberg (2023) analyze local capital destruction, migration, and spatial reallocation under climate change, while Indaco and Ortega (2024) study local population adjustment to climate exposure.

These papers emphasize migration, spatial reallocation, adaptation, and public insurance as the primary adjustment margins to climate risk. By contrast, I focus on household balance-sheet adjustment in incomplete markets and on the residual exposure households face when insurance and public compensation remain incomplete. The analysis shows that climate risk can materially alter wealth accumulation and self-insurance behavior even in the absence of migration or large spatial reallocations.

Rare Disasters and Stochastic Depreciation. The paper also relates to the rare-disaster literature pioneered by Barro (2009), Gourio (2012), and Gabaix (2012), where low-probability aggregate shocks shape macroeconomic dynamics and asset prices. Extensions incorporate heterogeneous beliefs and production heterogeneity to explain risk premia and valuation dynamics (Chen et al., 2012; Gomes et al., 2019). In these frameworks, disasters are typically aggregate events affecting the economy as a whole.

Unlike this literature, I study localized destruction shocks that leave aggregate productivity largely intact while substantially altering the cross-sectional distribution of wealth. The contribution is to show that disaster risk may operate primarily through distributional channels even when aggregate macroeconomic effects remain limited. In

this environment, the relevant amplification mechanism is not a collapse in aggregate productivity, but the concentration of destruction risk among households whose precautionary buffers remain imperfectly protected.

Relatedly, [Ambler and Paquet \(1994\)](#) introduce stochastic depreciation in an RBC model. I extend this approach by embedding empirically disciplined stochastic depreciation into an incomplete-markets economy with spatially heterogeneous exposure. This allows stochastic depreciation to affect not only aggregate capital dynamics, but also borrowing constraints, precautionary saving, and wealth dispersion.

Stochastic Returns and Wealth Inequality. Finally, the paper relates to the literature on stochastic returns and wealth inequality in incomplete-markets economies. A large body of work shows that multiplicative return risk and heterogeneous investment opportunities generate persistent wealth dispersion and fat-tailed wealth distributions ([Angeletos, 2007](#); [Benhabib et al., 2011, 2015, 2019](#); [Moll et al., 2022](#)). In these frameworks, inequality dynamics are primarily driven by persistent heterogeneity in returns, portfolio opportunities, or investment technologies.

I use this framework to study spatially heterogeneous disaster-induced destruction risk rather than persistent return heterogeneity. In standard stochastic-return models, multiplicative shocks primarily amplify upper-tail wealth concentration. In my model, climate risk instead reshapes the lower tail of the wealth distribution by weakening vulnerable households' ability to accumulate precautionary buffers.

Taken together, these literatures have emphasized aggregate disaster dynamics, spatial reallocation, public insurance, or persistent return heterogeneity as the primary channels through which climate risk affects the economy. This paper instead studies a setting in which the main saving instrument simultaneously serves as the primary source of self-insurance and as the object exposed to destruction risk. Embedding this mechanism in an incomplete-markets framework with spatially heterogeneous exposure generates large distributional effects even when aggregate fluctuations remain limited.

3. MODEL

The central mechanism studied in this paper operates by exposing the household's primary savings instrument to uninsured disaster-induced capital destruction. This assumption is empirically relevant because a large share of household wealth is held in real estate ([Jones and Neelakantan, 2023](#)), a locally exposed asset whose value can be

directly affected by natural-disaster risk. Because physical capital is both the vehicle for wealth accumulation and the object at risk, disaster exposure is not independent of household balance sheets — a feature that distinguishes this environment from standard incomplete-markets models with exogenous income risk. This mechanism is closely related to the interpretation in [Hsiang and Jina \(2014\)](#), in which repeated disaster exposure acts similarly to an increase in effective capital depreciation. In the present framework, however, this depreciation risk is heterogeneous across households and only partially insurable, generating endogenous distributional consequences through wealth accumulation dynamics.

Natural disasters also disrupt labor income through business closures, infrastructure damage, and job separations (e.g., [Groen et al., 2020](#); [Strobl, 2011](#)). This channel is incorporated via a state-dependent productivity shifter driven by the same Markov process as capital destruction and calibrated using existing empirical estimates of disaster-induced earnings losses. The inclusion of labor-income risk maintains the standard precautionary-saving environment familiar from the [Aiyagari \(1994\)](#) literature and provides a benchmark against which the additional effects of climate-related wealth destruction can be isolated. The labor-income channel primarily amplifies the distributional consequences of wealth destruction rather than constituting a separate source of climate-risk heterogeneity.

The productivity states should therefore be interpreted as idiosyncratic labor-efficiency units capturing these labor-income disruptions rather than differences in local wages. Households across all climate zones face common equilibrium factor prices and differ only in their productivity realizations and exposure to destruction risk.

3.1. Households. The household sector consists of a unit mass of infinitely-lived households, indexed by $i \in [0, 1]$.¹

Households have identical CRRA preferences over stochastic consumption streams $\{c_t^i\}_{t \geq 0}$:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{(c_t^i)^{1-\gamma}}{1-\gamma}, \quad (1)$$

¹In the data, I observe one representative agent per census tract across a finite number of tracts. In the model, I assume a continuum of tracts, playing the role of a standard continuum of households in heterogeneous-agent models.

where $\beta = \frac{1}{1+\rho}$ denotes the discount factor and $\gamma \geq 1$ is the coefficient of relative risk aversion.

Households face idiosyncratic exposure to natural disasters. Disaster realizations are persistent and heterogeneous across households, capturing differences in local exposure and vulnerability.

Let $\mathcal{J} = \{1, \dots, J\}$ denote the set of disaster regimes. The idiosyncratic disaster state variable x_t^i evolves according to a finite-state Markov chain and takes values in

$$x_t^i \in \left\{ e_1^{(J)}, \dots, e_J^{(J)} \right\}, \quad (2)$$

where $e_j^{(J)}$ denotes the j th column of the identity matrix $I_J \in \mathbb{R}^{J \times J}$, so that $x_t^i = e_j^{(J)}$ places household i in disaster regime j at time t . This selection-vector representation provides a compact characterization of state-dependent depreciation and labor-income shocks.

The dynamics of x_t^i are governed by the transition probability matrix $\Pi = (\pi_{jj'})_{j,j' \in \mathcal{J}}$, where

$$\mathbb{P} \left(x_{t+1}^i = e_{j'}^{(J)} \mid x_t^i = e_j^{(J)} \right) = \pi_{jj'}. \quad (3)$$

Each row of Π sums to one.

The disaster state jointly determines effective capital destruction and labor productivity. Capital holdings are subject to idiosyncratic depreciation shocks such that

$$\delta_t^i = \delta + s_t^i, \quad (4)$$

where $\delta > 0$ denotes standard physical depreciation and

$$s_t^i = \mu^\top x_t^i, \quad (5)$$

with

$$\mu = (\mu_1, \dots, \mu_J)^\top \quad (6)$$

collecting the state-dependent disaster-induced depreciation components. Hence,

$$x_t^i = e_j^{(J)} \quad \Leftrightarrow \quad s_t^i = \mu_j.$$

Similarly, let

$$\ell = (\ell_1, \dots, \ell_J)^\top \quad (7)$$

denote the vector of state-dependent labor productivity levels. Household labor income is given by

$$y_t^i = w_t \ell^\top x_t^i, \quad (8)$$

where w_t denotes the equilibrium wage rate. Note that the vector ℓ is normalized such that its stationary mean equals one,

$$\bar{\pi}^\top \ell = 1, \quad (9)$$

where $\bar{\pi}$ denotes the stationary distribution of x_t^i . This ensures that aggregate labor income equals w_t in equilibrium and that the productivity shifter introduces cross-sectional dispersion without affecting aggregate quantities or equilibrium prices.²

This structure allows moderate disaster regimes to primarily affect labor income while severe regimes jointly disrupt earnings and household wealth.

Each household supplies one unit of labor inelastically and accumulates physical capital, which is rented competitively to firms at the gross return r_t^K . Realized returns and labor income differ across households because disaster exposure generates household-specific depreciation and productivity shocks.

Formally, household i 's decisions are subject to the period- t budget constraint

$$c_t^i + a_{t+1}^i = (1 + r_t^K - (\delta + s_t^i)) a_t^i + y_t^i, \quad (10)$$

together with the borrowing constraint

$$a_{t+1}^i \geq -\phi. \quad (11)$$

The household chooses next-period assets a' to maximize expected lifetime utility, taking factor prices as given. The recursive problem is characterized by the value function $V(a, x)$ satisfying

$$V(a, x) = \max_{a'} \frac{c^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}[V(a', x') \mid x] \quad (12)$$

²The formal normalization of ℓ is detailed in Appendix A.5.

subject to

$$c + a' = (1 + r^K - (\delta + s(x))) a + y(x), \quad (13)$$

$$a' \geq -\phi, \quad (14)$$

where

$$s(x) = \mu^\top x, \quad y(x) = w\ell^\top x, \quad (15)$$

and expectations are taken with respect to the Markov transition matrix Π .³ The primes denote next-period variables.

3.2. Spatial Heterogeneity in Climate Risk. I consider a specification in which disaster risk varies persistently across space due to geographic differences in exposure. I partition the economy into a finite set of climate zones $\mathcal{Z} = \{1, \dots, Z\}$, where each climate zone corresponds to a distinct long-run exposure to natural disasters. Climate zone z hosts a continuum of households of measure λ_z , with $\sum_{z=1}^Z \lambda_z = 1$, and each household resides permanently in exactly one climate zone.⁴ The climate zone assignment, therefore, captures persistent geographic exposure.

Climate zones are ordered such that the distribution of disaster regimes becomes more severe with z , implying larger expected depreciation losses in more exposed regions. Formally

$$\mathbb{E}[s_t^i \mid z = Z] > \dots > \mathbb{E}[s_t^i \mid z = 1].$$

This ordering reflects more severe average damage in more exposed regions, while remaining agnostic to differences in higher-order moments across zones.

Within each climate zone $z \in \mathcal{Z}$, households face short-run disaster events indexed by $u \in \mathcal{U} = \{1, \dots, U\}$. Conditional on climate zone z , households transition across these event realizations according to a climate zone-specific U -state Markov chain with transition matrix $\Pi^z = (\pi_{uu'}^z)_{u, u' \in \mathcal{U}}$.

³Formally, when $x = e_j^{(J)}$,

$$\mathbb{E}[V(a', x') \mid x] = \sum_{j'=1}^J \pi_{jj'} V(a', e_{j'}^{(J)}).$$

⁴One could allow climate exposure itself to evolve over time, so that locations gradually transition across climate zones as risk intensifies or spreads. In that case, the pair (z, u) would simply follow a single J -state Markov process, as presented in the general framework above. I abstract from such climate-zone transitions in the baseline specification.

Let

$$\mu^z = (\mu_1^z, \dots, \mu_U^z)^\top$$

denote the vector of disaster-induced depreciation realizations associated with climate zone z . The vectors μ^z govern the severity of disaster events, while the transition matrices Π^z govern their persistence.

The full vector of possible depreciation realizations is obtained by concatenating the zone-specific vectors:

$$\mu = ((\mu^1)^\top, \dots, (\mu^Z)^\top)^\top \in \mathbb{R}^J, \quad J = Z \times U.$$

An analogous construction applies to the vector of state-dependent labor productivity levels ℓ .⁵

The pair $(z, u) \in \mathcal{Z} \times \mathcal{U}$ partitions the general J -state process (represented by the selection vector x_t^i) into economically interpretable climate-zone and event blocks, with each pair corresponding to a unique index $j \in \mathcal{J} = \{1, \dots, J\}$.

As in the general formulation, the climate-risk state is encoded by the selection vector

$$x_t^i \in \{e_1^{(J)}, \dots, e_J^{(J)}\},$$

which can be written as the Kronecker product

$$x_t^i = \zeta_t^i \otimes v_t^i,$$


where

$$\zeta_t^i \in \{e_1^{(Z)}, \dots, e_Z^{(Z)}\}$$

denotes the climate-zone selection vector and

$$v_t^i \in \{e_1^{(U)}, \dots, e_U^{(U)}\}$$

⁵The normalization of ℓ is applied economy-wide, weighting each zone's stationary distribution by its population share λ_z , as detailed in Appendix A.5.



Notes: Organization of the $J = 6$ regimes used in the model. Each regime is indexed by $x \in \{1, \dots, 6\}$ and corresponds to a pair (z, u) , where $z \in \{1, 2\}$ denotes the *climate type* (with $z = 1$ representing a moderate climate and $z = 2$ an intensive climate), and $u \in \{1, 2, 3\}$ denotes the disaster-event realization within each climate type. These realizations generate different levels of economic damages, with $u = 1$ corresponding to low depreciation losses s_1 , $u = 2$ to medium losses s_2 , and $u = 3$ to high losses s_3 . Climate zones differ in the distribution of these realizations, implying larger expected damages in the intensive climate ($z = 2$) than in the moderate climate ($z = 1$). Formally, $\mathbb{E}[s_t^i | z = 2] > \mathbb{E}[s_t^i | z = 1]$.

FIGURE 3.1. Regimes Organization

the within-zone event realization vector.⁶ The depreciation and productivity realizations are still

$$s_t^i = \mu^\top x_t^i,$$

$$y_t^i = w_t \ell^\top x_t^i,$$

respectively, so that each pair (z, u) corresponds to $s_t^i = \mu_u^z$, and $y_t^i = w_t \ell_u^z$.

In the following analysis, I consider the case $Z = 2$ climate zones and $U = 3$ disaster-event realizations per zone. This implies a total of $J = 2 \times 3 = 6$ distinct climate-risk states. The six states correspond to all combinations $(z, u) \in \{1, 2\} \times \{1, 2, 3\}$ and are indexed by $x \in \{1, \dots, 6\}$. Figure 3.1 illustrates the organization of these regimes.

Heterogeneity therefore arises along two dimensions: across zones, due to differences in long-run exposure (persistent climate risk), and within zones, due to stochastic regime realizations (transitory disaster risk). Importantly, the household's recursive problem is invariant to whether climate risk is represented directly as a J -state Markov chain or decomposed into Z zones of U regimes. Conditional on zone z , households transition across the event states $u \in \mathcal{U}$ according to Π^z , with each event state associated with a depreciation realization μ_u^z and productivity level ℓ_u^z . Thus, the zone-regime structure

⁶For example, when $Z = 2$ and $U = 3$, a household in zone 1 and regime 2 is represented by $\zeta_t^i = (1, 0)^\top$, $v_t^i = (0, 1, 0)^\top$, and $x_t^i = (0, 1, 0, 0, 0, 0)^\top$. The individual zone and regime components can be recovered from x_t^i using the following equivalence

$$\begin{aligned} (x_t^i = e_j^{(J)}) &\Leftrightarrow (s_t^i = \mu_j) \\ &\Leftrightarrow \begin{cases} \zeta_t^i = (I_Z \otimes \mathbf{1}_{\{1 \times U\}}) x_t^i, \\ v_t^i = (\mathbf{1}_{\{1 \times Z\}} \otimes I_U) x_t^i. \end{cases} \end{aligned}$$

imposes economically interpretable spatial structure on the underlying J -state process without altering the household's recursive problem.

3.3. The Representative Firm. The production sector consists of a representative firm operating a constant returns to scale production function, combining capital K_t and labor L_t to produce final output $Y_t = F(K_t, L_t)$. I assume a standard Cobb-Douglas production function,

$$F(K_t, L_t) = AK_t^\alpha L_t^{1-\alpha},$$

with an exogenous technological factor $A > 0$, and $\alpha \in (0, 1)$.

The firm hires labor and rents capital in competitive factor markets at prices w and r^K , respectively. Importantly, the model features a single representative firm operating at the national level rather than separate production sectors for each climate zone. Households from all climate zones supply labor and savings to a common economy and therefore face the same equilibrium wage and interest rate. Consequently, climate zones interact through their joint contribution to aggregate capital accumulation and the determination of equilibrium factor prices rather than through local production differences.

The aggregate capital stock is determined by the economy-wide asset holdings of households across all climate zones:

$$K_t = \int_0^1 a_t^i di = \sum_{z \in \mathcal{Z}} \lambda_z \sum_{u \in \mathcal{U}} \int a_t m_t^z(da, u).$$

Given prices, the firm chooses inputs to maximize its profit:

$$\max_{K_t, L_t} Y_t - r_t^K K_t - w_t L_t, \tag{16}$$

$$\text{s.t. } Y_t = AK_t^\alpha L_t^{1-\alpha}. \tag{17}$$

The first-order conditions give the firm's demand for capital and labor:⁷

$$r_t^K = \alpha AK_t^{\alpha-1}, \tag{18}$$

$$w_t = (1 - \alpha)AK_t^\alpha. \tag{19}$$

3.4. The Recursive Competitive Equilibrium. Households are permanently assigned to a climate zone $z \in \mathcal{Z}$, where each has a total population share $\{\lambda_z\}_{z=1}^Z$, satisfying $\sum_{z=1}^Z \lambda_z = 1$.

⁷Given the normalization $L_t = 1$.

Conditional on climate zone z , households transition across within-zone disaster-event realizations $u \in \mathcal{U} = \{1, \dots, U\}$ according to the zone-specific transition matrix $\Pi^z = (\pi_{uu'}^z)_{u, u' \in \mathcal{U}}$.

Although households differ in their exposure to climate risk across zones, capital and labor markets remain fully integrated. Households therefore face common equilibrium factor prices (r^K, w) , and climate heterogeneity affects equilibrium only through the distribution of household states and asset holdings across zones.

A *stationary recursive competitive equilibrium* consists of climate zone-specific policy functions $\{g^z\}_{z \in \mathcal{Z}}$, climate zone-specific stationary distributions $\{m^z\}_{z \in \mathcal{Z}}$, an aggregate capital stock K , and prices (r^K, w) such that:

- (i) **Households.** Given prices (r^K, w) , for each zone $z \in \mathcal{Z}$ the policy function $a' = g^z(a, u)$ solves the recursive problem (12), where expectations are taken with respect to the zone-specific transition matrix Π^z .⁸
- (ii) **Firm.** Given prices, the representative firm behaves optimally, and factor prices satisfy (18)–(19).
- (iii) **Market clearing.** Given the distribution $m^z(a, u)$, markets clear:

$$K = \sum_{z \in \mathcal{Z}} \lambda_z \sum_{u \in \mathcal{U}} \int a m^z(da, u), \quad L = \sum_{z=1}^Z \lambda_z = 1.$$

- (iv) **Stationarity.** For each zone $z \in \mathcal{Z}$, given the saving policy function $g^z(a, u)$ from (i) and transition matrix Π^z , $m^z(a, u)$ is the stationary distribution.

$$m^z(a', u') = \sum_{u \in \mathcal{U}} \int_{a: a' = g^z(a, u)} m^z(da, u) \pi_{uu'}^z.$$

4. DATA

This section constructs empirical measures of disaster-induced capital depreciation for model calibration. A central challenge is that natural-disaster damages are highly localized: county-level aggregates mechanically smooth destruction across households and may therefore substantially understate effective exposure to capital risk (Strobl, 2011; Fried, 2022). To address this issue, I combine county-level disaster damages with

⁸A full description of the solution method used to numerically solve the model can be found in Appendix A.

spatial exposure measures from FEMA’s National Risk Index (NRI) to construct tract-level depreciation rates. The resulting tract-year panel is then used to estimate the stochastic depreciation processes s governing household capital risk in the model.

4.1. A County-Level Depreciation Rate. The empirical counterpart to the model’s depreciation shock s is the county-level loss rate

$$d_t^i \equiv \frac{\text{NDL}_t^i}{K_{t-1}^i}, \quad (20)$$

where NDL_t^i denotes direct economic losses from natural hazards in county i and year t , and K_{t-1}^i denotes the corresponding beginning-of-period county capital stock. Disaster losses are obtained from SHEL DUS ([Center for Emergency Management and Homeland Security, 2023](#)), which reports county-level property and crop damages for all U.S. natural hazard events from 1960–2022. Following the asset boundary of the BEA fixed-assets accounts, NDL_t^i is restricted to property losses, as cultivated assets are not capitalized in the national accounts.⁹ County capital stocks are constructed by allocating national BEA fixed assets according to county GDP shares ([U.S. Bureau of Economic Analysis, 2025](#)); these county-level estimates form the basis for the tract-level capital stock measures constructed in Section 4.2.¹⁰

The analysis focuses on the 2002–2022 period. Although SHEL DUS records extend further back in time, the dataset’s geographic and hazard-level detail improves substantially beginning in the early 2000s. In addition, county-level GDP data from the BEA, used to construct county capital stocks, are only consistently available starting in 2001. Because the depreciation measure uses beginning-of-period capital stocks, the resulting dataset forms a balanced panel of 3,103 counties over 2002–2022.^{11,12}

I next characterize the spatial distribution of county-level disaster risk by contrasting the frequency of natural disasters with their severity. Figure 4.1 maps two complementary dimensions of exposure over 2002–2022. Panel A reports the percentile rank of

⁹Appendix Figure B.2 reports estimated depreciation rates incorporating crop losses.

¹⁰This procedure implicitly assumes a common capital-output ratio across counties. Appendix B.1 discusses the resulting measurement error and evaluates its relationship with disaster exposure.

¹¹Appendix B.2 describes the geographic harmonization procedures between datasets. Appendix B.3 additionally reconciles SHEL DUS with FEMA Presidential Disaster Declaration records to mitigate documented undercoverage of major disaster events in the early part of the sample.

¹²While the data allow for substantially finer disaggregation – including by hazard type and by major versus minor disaster events – I focus here on aggregate natural-disaster exposure. SHEL DUS distinguishes detailed hazard categories, which I aggregate into broader groups for expositional simplicity; the corresponding mapping is reported in Appendix Table B.1. Appendix B.4 reports additional evidence on heterogeneity in disaster frequency and severity across hazard categories.

cumulative hazard records per county, capturing the frequency of disaster occurrence. Panel B reports the percentile rank of the mean depreciation rate conditional on positive losses, capturing the severity of destruction when disasters occur.




FIGURE 4.1. Frequency and Severity of Natural Disaster Exposure across U.S. Counties

Notes: This figure displays the county-level percentile distribution of natural disaster exposure and associated capital depreciation rates across the United States. Panel A maps the percentile rank of the cumulative number of disaster records per county over 2002–2022. Panel B maps the percentile rank of the mean county-level capital depreciation rate conditional on strictly positive disaster losses over the same period. Percentile ranks are computed within each panel across all U.S. counties. Darker shades indicate higher percentile ranks. Grey counties have no strictly positive loss observation and are excluded from the Panel B distribution. Figure based on SHEL DUS disaster records and the author’s depreciation-rate estimates.

Disaster exposure displays substantial geographic heterogeneity along both dimensions. The spatial distribution of disaster frequency differs markedly from that of conditional depreciation rates: counties with frequent disasters do not necessarily face high capital losses per event, whereas some counties exposed to fewer events systematically face higher conditional depreciation. The figure, therefore, highlights that disaster frequency is not a sufficient statistic for disaster-induced capital risk.

On the quantitative side, Table 4.1 reports the distribution of county-level depreciation rates conditional on positive losses. The distribution is strongly right-skewed: the mean exceeds the 90th percentile, indicating that aggregate averages are driven by a small number of catastrophic realizations rather than by typical events.

At the same time, the table reveals an important limitation of county-level measurement. Even in the upper tail of the distribution, implied depreciation rates remain

quantitatively small. This is a direct consequence of geographic aggregation. By construction, d_t^i spreads total county losses over the entire county capital stock, including households and firms that were never directly affected by the disaster. A disaster that destroys a substantial share of capital in a localized area may therefore appear as only a negligible depreciation rate once losses are averaged across the entire county economy.

TABLE 4.1. Distribution of County-Level Capital Depreciation Rates

Share NDL > 0	Mean (%)	Median (%)	P90 (%)	P95 (%)	P99 (%)
73	0.136	0.003	0.080	0.225	1.894

Notes: This table reports summary statistics for county-level capital depreciation rates induced by natural disasters. The depreciation rate is defined as total reported disaster damages in a county-year divided by the county’s capital stock in the previous year. The first column reports the share of county-year observations with positive reported natural disaster losses (NDL); The mean, the median, and percentiles (P90, P95, and P99) are computed over that subsample. All values are expressed in percentage points.

Survey evidence corroborates this aggregation mechanism. While county-level data imply near-universal exposure, with 73% of county-years recording positive losses, the Federal Reserve’s Survey of Household Economics and Decisionmaking (SHED; [Board of Governors of the Federal Reserve System, 2024](#)) indicates that only 10–20% of households report any financial impact from natural disasters in the preceding year, and roughly 2% report substantial impacts.¹³ The gap between these two figures reflects the smoothing induced by county-level aggregation: disasters are geographically concentrated within counties, and averaging losses over the entire county capital stock mechanically understates the depreciation experienced by directly affected households.

Taken together, the evidence points to a framework in which disaster risk is characterized by both spatially heterogeneous frequency and severity rather than a uniform aggregate process – a structure the model is designed to capture. County-level depreciation rates nevertheless provide a downward-biased proxy for household-level exposure to disaster-induced capital risk. The following section addresses this measurement challenge by constructing census-tract-level depreciation measures that preserve national coverage while substantially reducing aggregation bias.

4.2. Construction of Census-Tract-Level Depreciation Rates. Census tracts are relatively permanent subdivisions of counties containing, on average, roughly 4,000 residents. To construct tract-level depreciation rates, I combine county-level disaster losses

¹³The complete table is reported in Appendix Table [B.5](#).

with tract-level information from FEMA’s National Risk Index (NRI), a composite measure of local natural-disaster risk reported for all U.S. census tracts ([Federal Emergency Management Agency, 2025](#)).¹⁴

The NRI combines three dimensions of local disaster risk. Expected annual losses (EAL) capture baseline hazard exposure, while the Social Vulnerability Index (SVI) and Community Resilience Index (CRI) characterize the susceptibility of local communities to disaster damages and their capacity to recover from them. FEMA further classifies census tracts into five ordered risk categories, ranging from “Very Low” to “Very High” risk, using a k-means clustering procedure that groups locations with similar underlying risk profiles. These measures are used to distribute county-level disaster exposure across space and to construct tract-level effective capital stocks.¹⁵

Table 4.2 summarizes the main assumptions and data sources underlying the construction of tract-level depreciation rates.

TABLE 4.2. Construction of Tract-Level Depreciation Rates

Stage	Object	Rule	Definition / Source
<i>Exposure Allocation</i>			
1	Exposure shares	Proportional to within-county expected annual losses	$w_{i j} = f(EAL^i)$ FEMA NRI
2	Event allocation	Multinomial allocation of county disaster events	$Records_t^i \sim \mathcal{M}(Records_t^j, w_{i j})$ SHELDUS
<i>Capital Construction</i>			
3	Physical capital	Proportional to tract building-value shares	$k_t^i = f(K_t^j, BV_i^{2020})$ BEA + FEMA NRI
4	Effective capital	Adjusted for vulnerability and resilience	$\tilde{k}_t^i = f(k_t^i, SVI^i, CRI^i)$ FEMA NRI
<i>Loss Allocation</i>			
5	Tract losses	Aggregation-consistent spatial allocation	$ndl_t^i = f(NDL_t^j, w_{i j}, \tilde{k}_t^i)$ SHELDUS
6	Depreciation rate	Allocated losses relative to effective capital	$d_t^i = ndl_t^i / \tilde{k}_t^i$

Notes: The table summarizes the main steps used to construct tract-level depreciation rates from county-level disaster losses and FEMA tract-level risk measures. Full derivations and implementation details are reported in [Appendix B.6](#).

¹⁴Appendix B.6 provides full construction details and derivations.

¹⁵See [Federal Emergency Management Agency \(2025, Section 3.2\)](#), for details on the construction of NRI ratings.

1) *Exposure shares*. Within each county, disaster events are allocated across census tracts according to tract-level expected annual loss (EAL) shares,

$$w_{i|j} = \frac{EAL_i}{\sum_{k \in j} EAL_k}, \quad \sum_{i \in j} w_{i|j} = 1,$$

where $w_{i|j}$ denotes the exposure share of tract i within county j .

Figure 4.2 illustrates how the allocation procedure maps tract-level hit probabilities into FEMA risk categories. Tracts with low hit probabilities are predominantly classified as “Very Low” or “Relatively Low” risk, whereas tracts with higher hit probabilities are increasingly concentrated in “Relatively High” and “Very High” risk categories. This monotonic relationship reflects the role of expected annual losses in distributing county-level disaster intensity across space: tracts with greater baseline exposure are more likely to experience disaster events.




FIGURE 4.2. Census Tract Hit Probability by FEMA National Risk Index Rating

Notes: The figure reports the composition of census tracts by FEMA National Risk Index (NRI) rating across the distribution of tract-level hit probabilities. Hit probabilities are constructed by allocating county-level disaster events to tracts based on their expected annual loss (EAL) shares. For each value on the horizontal axis, the stacked areas report the share of tracts in each NRI rating category.

2) *Event allocations.* Conditional on the realized number of county-level disaster events, tract-level event counts follow a multinomial allocation with probabilities $\{w_{i|j}\}_{i \in j}$.¹⁶ This construction preserves observed county-level disaster counts while allowing disaster exposure to vary systematically across census tracts within a county: tracts with higher exposure shares $w_{i|j}$ are more likely to experience disasters than lower-exposure tracts. In a typical draw, roughly 16% of the U.S. population resides in census tracts classified as hit within a given year, closely matching the household exposure rates reported in the SHED survey evidence.¹⁷

3) *Physical capital.* Tract-level physical capital is constructed by allocating county capital according to 2020 tract building replacement-value (BV) shares derived from FEMA’s NRI:

$$\kappa_{i|j} = \frac{BV_i^{2020}}{\sum_{k \in j} BV_k^{2020}}, \quad \sum_{i \in j} \kappa_{i|j} = 1,$$

where $\kappa_{i|j}$ denotes tract i ’s share of county j ’s capital stock. These shares are held fixed over time, implying

$$k_t^i = \kappa_{i|j} K_t^j. \quad (21)$$

4) *Effective capital.* Physical capital alone is unlikely to provide a sufficient measure of disaster exposure at the local level. Two census tracts with identical physical capital stocks may differ substantially in the extent to which that capital is protected against, exposed to, or able to recover from natural disasters. Local socioeconomic vulnerability, infrastructure quality, insurance coverage, and institutional capacity all affect the amount of capital effectively at risk during disaster events.

To capture these differences, I construct a measure of tract-level effective capital that scales physical capital using FEMA measures of social vulnerability (SVI) and community resilience (CRI):

$$\tilde{k}_t^i = k_t^i \omega_i^{SVI} \omega_i^{CRI},$$

¹⁶This follows from the conditional multinomial property: if tract-level event arrivals are independent Poisson processes, then conditioning on the observed county-level number of events implies a multinomial allocation across census tracts. The formal description in Appendix B.6.3, Eq.(31).

¹⁷Conditional on county-level disaster counts, the tract allocation is stochastic. In practice, however, the resulting exposure distribution is highly stable because the panel contains many disaster events observed over many years; Appendix B.6.3 provides additional details and robustness checks.

where the vulnerability and resilience wedges, ω_i^{SVI} and ω_i^{CRI} respectively, are exponential transformations of the normalized SVI and CRI indices.¹⁸ By construction, effective capital is weakly lower than physical capital and declines more strongly in socially vulnerable or low-resilience areas. The adjustment, therefore, concentrates disaster exposure toward locations where capital is both more fragile and less recoverable, allowing tract-level depreciation rates to better approximate localized disaster risk.




FIGURE 4.3. Distribution of Tract-Level Capital

Notes: The figure reports the distribution of tract-level physical capital k_t^i and effective capital $\tilde{k}_t^i$ over 2002–2022. Effective capital scales physical capital using FEMA measures of social vulnerability and community resilience, implying lower effective capital in more vulnerable and less resilient census tracts. The horizontal axis reports log capital, and the vertical axis the corresponding kernel density.

5) *Loss allocation.* Observed county-level disaster losses are allocated across hit census tracts to construct tract-level realized losses ndl_t^i for tract i in year t . The aggregation-consistent procedure is designed to generate localized severe destruction together with smaller losses in surrounding affected areas. In county-years with multiple hit tracts, the procedure first identifies an epicenter tract based on disaster frequency and FEMA-based historical expected depreciation intensity. Losses are then concentrated toward the

¹⁸Because higher values of the Community Resilience Index correspond to more resilient communities, I use the resilience deficit $cd^i = 1 - CRI^i/100$ so that both vulnerability and low resilience reduce effective capital through the same monotonic transformation. Appendix B.6.4 provides additional details.

epicenter, while remaining damage is distributed across other affected tracts according to their historically expected depreciation rates. All tract-level allocations are capped at effective capital and remain fully consistent with observed county-level disaster losses.¹⁹

6) *Depreciation rate.* The tract-level allocation procedure generates a micro-level disaster process characterized by localized severe destruction together with smaller losses in surrounding affected areas. Relative to county-level measures, the construction substantially reduces the mechanical smoothing induced by geographic aggregation while remaining fully consistent with observed county-level disaster losses.

The resulting tract-level depreciation rate is defined as

$$d_t^i = \frac{ndl_t^i}{\tilde{k}_t^i},$$

where ndl_t^i denotes realized disaster losses allocated to tract i in year t and $\tilde{k}_t^i$ denotes tract-level effective capital.

The tract-level allocation substantially alters the distribution of measured disaster exposure relative to county-level averages. As shown in Figure 4.4, tract-level depreciation rates are considerably more dispersed and exhibit a substantially thicker upper tail, reflecting the localized concentration of disaster-induced capital destruction within counties.

This concentration is also visible in Table 4.3. Although only 16% of tract-year observations record positive disaster losses, the upper tail of the conditional depreciation distribution is economically large: the 99th percentile exceeds 13%. Relative to county-level measures, the tract-level construction yields substantially higher realized depreciation rates by concentrating aggregate losses in localized areas rather than mechanically averaging them across entire counties. The resulting process, therefore, captures rare but severe local capital destruction largely absent from aggregate county-level data.

The resulting dataset is a balanced tract-year panel of 84,093 tracts over 2002–2022 combining national capital stocks, county-level disaster losses, and tract-level exposure and vulnerability measures. It provides the empirical foundation for estimating heterogeneous stochastic depreciation processes across climate zones in Section 5.

¹⁹Appendix B.6.5 provides the full allocation algorithm and parameterization.




FIGURE 4.4. Distribution of Log Depreciation Rates: County vs. Census Tract

Notes: The figure compares the distribution of log depreciation rates at the county and census-tract levels for observations with positive losses over 2002–2022. County-level depreciation rates are defined as total reported disaster losses divided by county capital stock. Census-tract depreciation rates are constructed by allocating county-level losses across hit tracts in proportion to their historical disaster exposure, then dividing the allocated losses by tract-level effective capital.

TABLE 4.3. Distribution of Tract-Level Capital Depreciation Rates

Share $ndl > 0$	Mean (%)	Median (%)	P90 (%)	P95 (%)	P99 (%)
16	0.593	0.000	0.092	0.613	13.477

Notes: This table reports summary statistics for tract-level capital depreciation rates induced by natural disasters. The depreciation rate is defined as total reported disaster damages in a tract-year divided by the tract’s effective capital stock in the previous year. The first column reports the share of tract-year observations with positive reported natural disaster losses (ndl); The mean, median, and percentiles (P90, P95, and P99) are computed over that subsample. All values are expressed in percentage points.

5. ESTIMATION OF THE DISASTER-INDUCED DEPRECIATION PROCESS

This section estimates the stochastic process governing disaster-induced capital depreciation across census tracts. The objective is to recover a parsimonious regime-switching process that can be embedded directly into the quantitative heterogeneous-agent model developed in the following section.




FIGURE 5.1. Climate Zones

To capture persistent geographic differences in disaster exposure, I estimate the depreciation process separately across climate zones. Throughout the exposition, the climate-zone index $z \in \mathcal{Z}$ is omitted whenever no confusion arises.

5.1. Climate Zones. Natural-disaster exposure is highly heterogeneous across locations, requiring a distinction between regions facing moderate and intensive climate risk. I therefore partition census tracts into two climate zones using FEMA’s National Risk Index (NRI) classification ([Federal Emergency Management Agency, 2025](#)). Census tracts classified as *Relatively High* or *Very High* risk are assigned to the intensive zone, while all remaining tracts belong to the moderate zone.

The resulting classification isolates regions historically characterized by elevated disaster exposure and larger realized depreciation losses. Figure 5.1 displays the geographic distribution of the two zones. Intensive tracts, which represent 16% of census tracts in the sample, are concentrated along coastal and wildfire-prone regions, broadly consistent with the spatial distribution of realized depreciation rates shown in Panel B of Figure B.4.

5.2. Regime-Switching Depreciation Process. Disaster-induced depreciation is modeled as a latent regime-switching process combining infrequent zero-depreciation realizations with persistent high-damage states. Within each climate zone, census tracts are

assumed to be independent and identically distributed. Let $i = 1, \dots, N$ index census tracts and $t = 1, \dots, T$ index years. The observed depreciation rate in tract i and year t is denoted by $d_t^i \geq 0$.

Conditional on the latent regime, positive depreciation rates follow an logit-normal distribution with regime-specific destruction intensity. Specifically,

$$\log \left(\frac{d_t^i}{1 - d_t^i} \right) \mid s_t^i = u \sim \mathcal{N}(\mu_u, \sigma^2), \quad (22)$$

where μ_u denotes the regime-specific destruction intensity and σ is a common within-regime dispersion parameter.

Because many tract-year observations exhibit zero depreciation, the estimation incorporates a zero-mass component alongside the continuous distribution of positive losses. Conditional on regime u , depreciation equals zero with probability $1 - p_u$ and is strictly positive with probability p_u .

Persistence in disaster exposure is captured by a latent first-order Markov chain that governs transitions across regimes. The transition dynamics are summarized by the matrix

$$\Pi = (\pi_{uu'})_{u,u'=1}^U.$$

The model is estimated using maximum likelihood with a Hamilton filter (Hamilton, 1994). Full derivations of the likelihood, filtering recursion, and panel representation are provided in Appendix C.

5.3. Structural Restrictions. The estimated process is disciplined to maintain a transparent mapping between the empirical specification and the discrete depreciation states used in the quantitative model. In the structural environment, households face a finite set of depreciation shocks governed by a Markov transition matrix. I therefore impose several economically motivated restrictions on the empirical specification.

The first restriction separates normal periods from realized disaster states. I set $U = 3$ and impose

$$p = (0, 1, 1)^\top,$$

so that the first regime corresponds to a no-disaster state with degenerate zero depreciation.

The second restriction disciplines the estimation to capture rare but economically meaningful tail events. In particular, the third regime is calibrated to target the upper tail of the empirical depreciation distribution and isolate severe disaster realizations.

The third restriction limits within-regime dispersion in order to preserve sharp differences across latent states. I therefore impose a common within-regime dispersion parameter σ and penalize excessively large values in the likelihood function. This restriction ensures that heterogeneity is captured primarily through differences in regime means and transition probabilities rather than excessive smoothing within states.

The resulting estimates recover a small number of economically interpretable disaster regimes that can be directly embedded into the heterogeneous-agent model. Estimated parameters are reported in Table 5.1.

Zone	s_{low}	s_{medium}	s_{high}
Moderate	0.0000	0.0002	6.4473
Intensive	0.0000	0.0006	23.7150

<i>Common Parameters</i>			
p	0	1	1
σ		1	

Moderate Zone				
	s_{low}	s_{medium}	s_{high}	$\bar{\pi}$
s_{low}	0.893	0.101	0.006	0.852
s_{medium}	0.634	0.333	0.033	0.137
s_{high}	0.375	0.485	0.139	0.011

Intensive Zone				
	s_{low}	s_{medium}	s_{high}	$\bar{\pi}$
s_{low}	0.830	0.147	0.023	0.539
s_{medium}	0.170	0.767	0.064	0.415
s_{high}	0.461	0.384	0.155	0.046

TABLE 5.1. Estimated Idiosyncratic Shock States and Transition Matrices

Notes: The top panel reports the estimated discrete depreciation shocks (in percent) for each regime of the Markov-switching model. The idiosyncratic shock process follows a finite-state Markov chain and is estimated by maximum likelihood using the Hamilton filter in a panel setting. Parameters p and σ are common across zones. The bottom panels report the estimated transition matrices; rows sum to one and the final column shows the implied stationary distribution. The Aiyagari model is calibrated using the estimated depreciation shocks s and the corresponding transition matrices.

5.4. Estimation Results. Estimated disaster dynamics differ sharply across climate zones. Tracts located in intensive regions face both larger conditional destruction shocks and more persistent exposure to severe disaster realizations than tracts in moderate regions. Table 5.1 shows that the intensive zone is characterized by both higher regime-specific depreciation intensities and greater persistence of high-damage states.




FIGURE 5.2. Estimated and Empirical Depreciation Distributions

The estimated upper-tail depreciation shocks remain within the range considered in recent climate-macro calibrations. For example, [Bilal and Rossi-Hansberg \(2023\)](#) consider effective climate-related depreciation rates reaching roughly 17% in highly exposed regions. Although the underlying objects differ conceptually, the comparison suggests that the estimated local destruction intensities are economically plausible.

The estimated regime-switching process reproduces the main empirical features of the depreciation distribution across climate zones. In particular, the estimation captures both the large mass of zero-depreciation observations and the substantially thicker right tail observed in intensive regions. Figure 5.2 compares the estimated mixture distribution with the empirical density for the aggregate sample and separately by climate zone.

The objective of the estimation is to identify economically meaningful disaster regimes rather than perfectly match every local fluctuation of the empirical density. Consistent with this objective, the estimated distribution exhibits a distinct modal structure associated with the discrete latent states while abstracting from minor irregularities in the empirical kernel density.

The estimated latent states map naturally into the discrete depreciation shocks used in the structural economy. Figure 5.3 illustrates how the estimated regime probabilities and depreciation intensities translate into the discrete shock structure embedded in the quantitative model.




FIGURE 5.3. Discrete Regime Representation

6. QUANTITATIVE RESULTS

6.1. Calibration. I embed the results of my regime-switching process into my depreciation-shock Aiyagari model. Table 6.1 reports the parameters that are externally set rather than estimated. Preferences are standard and imply moderate risk aversion. Technology parameters follow conventional macroeconomic values. The productivity process is a symmetric three-state Markov chain, normalized to ensure the stationary mean of effective labor equals one.

Parameter	Description	Value
γ	Relative risk aversion	2
ρ	Discount rate	0.05
δ	Depreciation rate	0.10
α	Capital share in production	0.33
Z	Number of climate zones	2
U	Number of productivity states	3
<i>Productivity States (normalized)</i>		
ℓ_1	High productivity	1.02
ℓ_2	Middle productivity	0.94
ℓ_3	Low productivity	0.85

TABLE 6.1. Model Calibration (Non-Estimated Parameters)

Notes: The productivity process is symmetric and calibrated to generate a stationary mean of effective labor equal to one. Destruction risk differs across climate zones through the distribution of idiosyncratic depreciation shocks.

The idiosyncratic labor-income process is calibrated externally rather than estimated within the model. A growing empirical literature documents substantial earnings losses and persistent labor-market disruptions following natural disasters (Pleninger, 2022; Groen et al., 2020). Studying Hurricanes Katrina and Rita, Groen et al. (2020, Table 2) estimate that workers exposed to minor residential or workplace damage experienced average first-year earnings losses of 6%-8%, while those exposed to major damage experienced losses of 15%-17%. Motivated by these estimates, I calibrate a symmetric three-state productivity process with labor-productivity states

$$\ell = [1.02, 0.94, 0.85],$$

normalized to generate a stationary mean of effective labor equal to one.²⁰

Because the model features a common national labor market with a single equilibrium wage, these productivity states should be interpreted as idiosyncratic labor-efficiency units that mimic the earnings losses documented in the empirical literature rather than changes in local wages. This approach allows the model to capture disaster-induced labor-income disruptions while maintaining a tractable benchmark environment with integrated factor markets.

The calibration intentionally adopts conservative productivity losses relative to the upper end of the empirical estimates. The objective of the paper is not to model post-disaster labor-market dynamics, but rather to isolate the role of disaster-induced capital destruction within a heterogeneous-agent framework. For this reason, the labor-income process is held fixed across climate zones. All cross-zone heterogeneity enters through the estimated depreciation process and its transition matrix.

To provide economic intuition for the results that follow, the joint productivity–destruction states can be interpreted as representing different disaster environments. The high state corresponds to periods without significant disruption, in which households experience normal labor income and no capital destruction. The middle state captures moderate events that primarily affect labor income and economic activity while generating limited destruction of accumulated wealth. The low state represents severe disasters that combine substantial earnings losses with large depreciation shocks. By holding the labor-income process fixed across climate zones, the quantitative analysis

²⁰See Appendix A.5 for the normalization procedure.

isolates how differences in the distribution and persistence of destruction risk affect household behavior and wealth accumulation.

6.2. Quantitative Results. Before introducing disaster risk, the model displays the familiar precautionary saving behavior of incomplete-markets economies. Wealth provides self-insurance against future income fluctuations, so households facing greater uncertainty accumulate larger buffers and dissave less aggressively in response to adverse shocks. Accordingly, the dashed lines in Figure 6.1 show that households in Zone 2 save more than those in Zone 1 despite identical current productivity realizations, reflecting their greater exposure to future risk.

The intuition follows directly from the standard buffer-stock saving motive. In a conventional Aiyagari environment, wealth is a safe asset that provides insurance against future income fluctuations. Additional savings, therefore, relax future borrowing constraints and reduce the marginal utility cost of adverse realizations. Because wealth itself is not exposed to destruction, the precautionary motive operates monotonically: more risk increases the value of holding assets. The same mechanism also explains why dissaving becomes less pronounced in bad states. Households facing greater future uncertainty preserve a larger precautionary buffer rather than fully smoothing contemporaneous consumption.

6.2.1. Climate Risk Weakens Self-Insurance Incentives. Introducing disaster-induced depreciation risk fundamentally changes the precautionary-saving mechanism because the saving asset itself becomes exposed to shocks. Wealth now plays a dual role: it remains the primary instrument of self-insurance, yet it becomes vulnerable to destruction. This interaction generates highly nonlinear saving responses across wealth levels. The solid lines in Figure 6.1 show that, once disaster risk is introduced, households in the high-risk zone no longer display uniformly stronger precautionary saving behavior. Instead, saving policies become sharply non-monotone. At low and intermediate wealth levels, households in Zone 2 dissave aggressively even in relatively favorable productivity states, while at higher wealth levels, precautionary accumulation re-emerges.

The key mechanism is that the insurance value of wealth now depends on its exposure to future destruction. When wealth is low, an additional unit of savings provides only limited insurance because a sufficiently adverse disaster realization may destroy a substantial share of those assets in the following period. In that region, the expected




FIGURE 6.1. State-Contingent Saving Policies Across Climate Zones

Notes: The figure compares saving policies with and without natural-disaster risk across climate zones and shock states. Dashed lines correspond to the model without disaster risk, while solid lines correspond to the model with disaster-induced depreciation risk. Colors distinguish the moderate climate zone (Zone 1) from the intensive climate zone (Zone 2). The left panels report the low-shock state, while the right panels report the high-shock state. Productivity risk is common across zones, with $\ell = [1.02, 0.94, 0.84]$, whereas depreciation risk varies by climate zone and is substantially higher in the intensive zone. Overall, disaster risk generates strong asymmetry in saving behavior across both wealth levels and climate zones.

marginal utility gain from future self-insurance can fall below the marginal utility of current consumption. Households, therefore, prefer to consume today rather than accumulate vulnerable wealth. Disaster risk weakens the standard precautionary motive precisely where borrowing constraints are most binding.

At higher wealth levels, however, the logic reverses. Wealthier households can accumulate sufficiently large buffers such that the insurance value of additional assets dominates expected destruction losses. In that region, households once again engage in precautionary accumulation. The result is the U-shaped saving profile observed for Zone 2 in Figure 6.1: households first reduce savings sharply as disaster exposure dominates, before precautionary motives gradually recover at higher asset levels. The nonlinearity becomes particularly pronounced in the medium- and high-shock states, where dissaving increases strongly with wealth before flattening out.

Figure 6.2 highlights the equilibrium implications of the distorted saving policies. In the intensive-risk zone, households consume more at low wealth levels despite facing

greater future risk. This pattern reflects the lower marginal value of wealth accumulation when future assets remain exposed to destruction. As households place less value on building precautionary buffers, they reallocate resources away from future wealth accumulation and toward current consumption. Over time, these decisions reshape the stationary distribution, generating a substantially larger mass of households near the borrowing constraint. Climate risk, therefore, affects not only how households allocate resources between current consumption and future wealth accumulation, but also the long-run distribution of wealth.

6.2.2. Future Destruction Risk Distorts Behavior. The previous figures compare saving behavior across climate zones and disaster states. In the medium- and high-shock states, households differ both in their realized depreciation shock and in their exposure to future disaster risk. To isolate the role of future destruction risk, I therefore focus on the low-shock state, where households in both zones face the same contemporaneous realization. Conditional on this common state, differences in behavior arise exclusively from cross-zone differences in the distribution of future depreciation risk.

Figure 6.3 directly compares households across climate zones in this common state. Current productivity, income, and realized depreciation are identical across zones. Nevertheless, households in the intensive-risk zone consume more and save less than households in the moderate-risk zone, despite facing the same contemporaneous conditions.

, therefore, The figure therefore shows that climate risk affects behavior even in the absence of realized losses. Anticipated future destruction reduces the value of accumulating precautionary wealth, leading households to substitute away from self-insurance and toward current consumption. Climate risk thus operates not only through realized capital losses but also through the expectation of future losses, which weakens precautionary saving before disasters occur.

Table 6.2 quantifies these differences at selected wealth levels. At low wealth levels, households in the intensive-risk zone save substantially less than otherwise identical households in the moderate-risk zone. For example, at $a = 5$, households in Zone 1 save 0.011 units while households in Zone 2 optimally dissave 0.015 units. The difference reverses at higher wealth levels, where precautionary accumulation gradually re-emerges. At $a = 25$ and $a = 50$, households in the intensive-risk zone save roughly three times more than households in the moderate-risk zone. These patterns mirror the non-monotonic saving schedules observed in Figure 6.3 and highlight the central role of




FIGURE 6.2. Stationary Policies and Wealth Distributions by Climate Zone

Notes: The figure reports consumption $c(a)$, net saving $a'(a)$, and stationary wealth density $g(a)$ for the moderate (Zone 1) and intensive (Zone 2) climate zones. The productivity shock is common across zones and equal to $\ell = \{1.02, 0.94, 0.85\}$. Differences across zones arise solely from cross-zone differences in the distribution of idiosyncratic depreciation risk. Low, Med, and High denote the household state and the corresponding additional shock to the disaster-induced depreciation rate. The Low state corresponds to no additional shock (0-shock), while High corresponds to an additional depreciation rate of 6.4% and 23.7% in Zones 1 and 2, respectively.

wealth in determining whether disaster exposure or self-insurance incentives dominate household behavior.

6.2.3. Distributional Effects Dominate Aggregate Effects. The distributional consequences documented above do not translate into equally large aggregate effects. Instead, Figure 6.4 compares the benchmark economy with only idiosyncratic labor-productivity risk




FIGURE 6.3. Identical Realization, Different Risk Distribution: Low-State Comparison

Notes: The figure compares Zone 1 and Zone 2 in the state with no contemporaneous depreciation shock (0-shock) that corresponds to the *Low* state. Productivity risk is identical in realization across zones. Since current income and realized destruction are the same, differences in consumption, saving, and the stationary wealth distribution are driven exclusively by cross-zone differences in the distribution of future idiosyncratic depreciation risk. The intensive climate zone exhibits weaker precautionary saving at low wealth levels and a larger mass of households near the borrowing constraint.

Wealth (a)	Moderate Zone (Z1)	Intensive Zone (Z2)	Relative Saving (Z2/Z1)
1	0.009	0.000	0.0
5	0.011	-0.015	-1.4
10	0.014	0.054	3.8
25	0.025	0.091	3.6
50	0.044	0.129	2.9

TABLE 6.2. Net Saving by Climate Zone at Selected Wealth Levels

Notes: The table reports net saving $a'(a)$ at selected wealth levels in Zone 1 and Zone 2, conditional on regime s_1 . The first two columns report optimal savings at wealth level a in each zone. The third column reports the savings ratio Z2/Z1.

(No ND) to the economy that additionally incorporates idiosyncratic disaster-induced depreciation risk (ND). Despite the substantial changes in household behavior documented above, the model generates a striking disconnect between household-level responses and aggregate outcomes: aggregate capital declines by less than one percent after introducing disaster risk.

The reason is that disaster risk affects households asymmetrically across the wealth distribution. Poor households reduce precautionary saving because accumulated wealth provides little effective insurance against future destruction. By contrast, wealthier




FIGURE 6.4. Aggregate Capital and Distributional Effects Across Model Variants

Notes: The figure compares the stationary equilibria with and without natural-disaster risk. The model labeled “No ND” includes only idiosyncratic productivity risk, $\ell = [1.02, 0.94, 0.84]$, while the model labeled “ND” additionally incorporates idiosyncratic depreciation risk associated with natural disasters. The left panel reports aggregate capital. The percentage difference in aggregate capital between the two equilibria is indicated above the bars. The middle panel reports the share of households at the borrowing constraint, and the right panel reports the Gini coefficient of the stationary wealth distribution. The Gini coefficient is computed from the invariant wealth distribution $g(a)$ using the discrete approximation $G = \frac{1}{N^2 \bar{a}} \sum_{i=1}^N (2i - N - 1) a_{(i)}$, where $a_{(i)}$ denotes ordered wealth levels and $\bar{a}$ their mean.

households continue to accumulate assets because sufficiently large buffers retain substantial insurance value. The decline in savings among poorer households is therefore partially offset by continued accumulation among wealthier households.

Although aggregate effects remain modest, the distributional consequences are substantial. The share of hand-to-mouth households rises sharply, and wealth inequality increases markedly. Climate risk thus reshapes who holds wealth rather than dramatically reducing the total amount of wealth held in the economy. The primary effect of disaster risk is not aggregate capital destruction, but a distortion of self-insurance incentives that shifts households toward the lower tail of the wealth distribution.

6.2.4. *Wealth-Dependent Exposure as the Mechanism.* To assess whether wealth-dependent exposure is central to the results, I consider a concave destruction technology,

$$\tilde{d}(a) = sa^\tau,$$

under which marginal exposure to destruction declines with increasing wealth. Therefore, the budget constraint of household i is

$$c_t^i + a_{t+1}^i = (1 + r_t^K - \delta) a_t^i + w_t - \tilde{d}(a_t^i).$$

Concavity captures the idea that wealthier households may be better able to protect their assets against climate risk. Higher-wealth agents can invest in adaptation measures (e.g., reinforced structures, elevated housing, protective infrastructure), choose less exposed locations, or access formal insurance markets. As a result, additional wealth does not increase effective exposure to destruction in proportion.

This specification contrasts with the baseline linear case, where exposure scales proportionally with wealth. While the empirical section 4.2 shows that effective capital is lower in structurally vulnerable locations, the model's concave formulation, in contrast, captures household-level adaptive capacity: higher wealth reduces marginal exposure by enabling investment in protection, insurance, or safer assets.

Figure 6.5 illustrates how saving behavior changes as destruction risk becomes less sensitive to wealth. For moderate reductions in τ , precautionary-saving incentives are progressively restored. Households in the intensive-risk zone become less willing to dis-save at low and intermediate wealth levels because additional assets no longer generate proportional increases in disaster exposure.

For sufficiently low values of τ , however, the mechanism changes qualitatively. When $\tau = 0.10$ or $\tau = 0.05$, disaster exposure becomes highly concentrated among households with very little wealth.²¹ As wealth increases, marginal exposure falls rapidly, so additional assets provide insurance with relatively little additional disaster risk. Households, therefore, face a strong incentive to accumulate enough wealth to escape the highly vulnerable region near the borrowing constraint. Once this threshold is crossed, precautionary motives strengthen sharply. The effect is particularly visible in the medium- and high-shock states, where saving policies become steeply upward sloping beyond approximately one unit of wealth.

The implications for the stationary wealth distribution are shown in Figure 6.6. Moving from the linear benchmark toward moderate concavity substantially reduces the mass

²¹For $\tilde{d}(a) = sa^\tau$, marginal exposure is $\partial\tilde{d}(a)/\partial a = s\tau a^{\tau-1}$. When $\tau < 1$, marginal exposure diverges as $a \rightarrow 0$ and declines monotonically with wealth. The concavity specification, therefore, concentrates disaster exposure among households with very low asset holdings.




FIGURE 6.5. Saving Policies under Alternative Degrees of Wealth Exposure

Notes: The figure reports optimal net saving policies $a'(a)$ for alternative values of the wealth-exposure parameter τ . The benchmark specification $\tau = 1$ corresponds to a linear destruction technology in which disaster-induced depreciation is proportional to wealth. Lower values of τ make destruction risk increasingly concave, reducing the marginal exposure of additional wealth to future losses. Columns correspond to the low-, medium-, and high-shock states, while rows correspond to the moderate (Zone 1) and intensive (Zone 2) climate zones. As τ decreases, precautionary-saving incentives are progressively restored, particularly in the intensive-risk zone, where households become less willing to dissave at low and intermediate wealth levels.

of households near zero wealth, consistent with the restoration of precautionary-saving incentives. For $\tau = 0.10$ and $\tau = 0.05$, the distribution changes more dramatically. The cumulative distribution becomes noticeably flatter, indicating that a substantial fraction of households hold much larger wealth.

Taken together, the saving policies and wealth distributions suggest the emergence of a threshold effect. Households that remain near the borrowing constraint continue to hold little wealth, while households that accumulate enough assets to move beyond the highly exposed region have a strong incentive to continue building precautionary buffers. Rather than simply reducing the mass of poor households, extreme concavity generates




FIGURE 6.6. Saving Policies under Alternative Degrees of Wealth Exposure

Notes: The figure reports the cumulative distribution function of stationary wealth for alternative values of the wealth-exposure parameter τ . The benchmark case $\tau = 1$ corresponds to a linear destruction technology, while lower values of τ imply increasingly concave exposure to disaster-induced depreciation. As the marginal exposure of wealth declines, households accumulate larger precautionary buffers and the stationary distribution shifts toward higher wealth levels. The concentration of households near the borrowing constraint falls substantially, indicating that wealth-dependent exposure is a key determinant of the distributional effects of disaster risk.

a more polarized wealth distribution, with households increasingly concentrated either near the lower tail or far into the upper tail.

These results reinforce the paper's central mechanism. The distortion to precautionary saving arises because wealth simultaneously provides insurance and increases exposure to destruction. Once the link between wealth accumulation and marginal exposure is weakened, precautionary-saving incentives re-emerge, and the distributional consequences of disaster risk are substantially altered.

6.2.5. *Summary.* Four results emerge.

- (1) Climate risk weakens the standard precautionary-saving motive. Because wealth is both the primary self-insurance device and the object exposed to destruction, greater risk can reduce rather than increase desired savings.

- (2) Anticipated future destruction risk affects behavior even in the absence of realized losses. Households facing identical current conditions but different future risk distributions consume more and save less when future wealth accumulation provides less effective insurance.
- (3) The consequences are primarily distributional rather than aggregate. Aggregate capital declines by less than one percent, yet the share of hand-to-mouth households rises sharply and wealth inequality increases substantially.
- (4) Wealth-dependent exposure is the central mechanism. Reducing the marginal exposure of wealth to destruction restores precautionary-saving incentives and dramatically alters the stationary wealth distribution.

Overall, climate risk operates less through aggregate capital destruction than through a distortion of household self-insurance incentives. By reducing the effectiveness of wealth accumulation as insurance, disaster risk reshapes the long-run distribution of wealth while leaving aggregate capital largely unchanged.

7. DISCUSSION

The analysis in this paper studies a deliberately restrictive environment in which households insure themselves against disaster risk only through precautionary saving in an asset that is itself exposed to destruction. This restriction isolates the paper’s central mechanism: when the primary savings instrument is also the object at risk, wealth accumulation provides insurance while simultaneously increasing exposure. At the same time, the framework is sufficiently rich to accommodate additional layers of climate-risk management. Because climate risk is modeled as a spatially heterogeneous destruction process, the same environment can be used to study how households respond when they gain access to safe assets, diversified financial claims, private insurance, public transfers, or adaptation technologies. The baseline model should therefore be interpreted not as a complete description of all available risk-sharing mechanisms, but as a benchmark from which the contribution of each additional instrument can be evaluated.

Financial Structure, the Climate Protection Gap, and Risk Sharing. The baseline economy intentionally abstracts from financial diversification by assuming that all wealth is exposed to disaster risk and that no risk-free asset exists. This assumption should not be interpreted as a claim that safe assets are absent in developed economies. Rather, it captures, in reduced form, the effective exposure faced by households whose

wealth is concentrated in locally exposed assets and who do not fully insure climate-related losses. The key issue is therefore not the formal existence of safer assets or insurance contracts, but whether households actually use them to protect their balance sheets against local destruction risk.

This financial structure is consistent with the composition of household portfolios in the United States. Housing is a central component of household wealth and is directly exposed to local natural-disaster risk. Households in the middle of the wealth distribution hold most of their wealth in real estate, while wealthier households invest more heavily in stocks and private business equity (Jones and Neelakantan, 2023). More precisely, real estate is the dominant asset between the 25th and 99th percentiles of the net-worth distribution, whereas business equity and stocks become quantitatively more important only in the upper part of the distribution and outweigh real estate only at the very top. For households below the 25th percentile, net worth is sufficiently low that portfolio shares are difficult to interpret, which is itself consistent with limited scope for meaningful diversification. The baseline model, therefore, captures an important feature of household balance sheets: for much of the population, wealth accumulation is closely tied to locally exposed physical assets.

This interpretation connects directly to the climate protection gap. Even with insurance and financial instruments in place, a substantial share of disaster losses remains uninsured. The gap may reflect missing markets, but it may also reflect limited take-up of existing instruments because of high premia, borrowing constraints, limited attention, low financial literacy, lack of trust, or imperfect information about coverage. In this sense, the baseline economy represents the lowest layer of the climate protection gap: the case in which households rely primarily on exposed wealth for self-insurance.

Adding financial assets or insurance markets is then a natural counterfactual use of the framework. A risk-free asset, diversified claims, or insurance contracts would introduce successive layers of risk sharing on top of the baseline economy. These instruments would allow households to separate precautionary saving from local disaster exposure and would enable quantification of how much of the distributional cost of climate risk can be reduced when households have access to safer claims. The destructible asset would remain central, however, because it represents housing and other locally exposed assets that provide housing services, collateral value, and local consumption benefits.

This extension is especially useful because access to risk-sharing instruments is likely to be unequal across the wealth distribution. The same empirical evidence that motivates the baseline also suggests that safer or more diversified claims are more relevant for households higher in the wealth distribution, who participate more in broader capital markets. A richer version of the model could therefore distinguish between the existence of risk-sharing instruments and their effective use. It could examine whether additional assets close the climate protection gap broadly or primarily benefit households already able to access financial markets.

Finally, the framework provides a disciplined benchmark for the role of public policy. If private insurance and diversified assets are unavailable, too costly, or insufficiently used by exposed households, governments remain the residual insurers of large disaster losses. The model could therefore evaluate whether policies that improve affordability, transparency, and take-up of private insurance reduce reliance on ex post public transfers. The objective would not be to replace public intervention entirely, but to distinguish risks that can be transferred through private mechanisms from those that remain systemic, non-diversifiable, or distributionally costly and therefore require public support.

Climate Change and Time-Varying Risk. Although the quantitative analysis focuses on current spatial heterogeneity in climate risk, the framework is sufficiently rich to capture future changes in its distribution. The model is formulated as a general J -state Markov process for climate exposure. In the baseline quantitative exercise, I use a restricted version of this process in which climate zones are independent risk environments, and the mass of households in each zone is fixed. This restriction isolates the role of current spatial heterogeneity. In the general case, however, the same Markov structure can allow transitions across climate-risk regimes. Climate change can then be represented as a shift in transition probabilities, with households or locations initially associated with moderate risk becoming more likely to move toward intensive-risk states. The framework can therefore describe not only heterogeneous climate exposure today, but also the gradual movement of the economy toward more exposed risk regimes over time.

Adaptation can be incorporated within the same structure as the counterpart to climate-driven risk intensification. While climate change can be represented as increasing the probability of moving toward more exposed risk regimes, adaptation reduces the effective risk associated with a given location. Investment in resilient housing, protective

infrastructure, or adaptation technologies can lower either the probability or the severity of destructive shocks, allowing households in physically risky areas to face a less risky effective distribution. Adaptation, therefore, need not change the household's location; it changes the climate-risk process relevant for household decisions.

Taken together, these applications show that the model provides a flexible laboratory for studying the limits of private climate-risk management. The baseline economy identifies the consequences of relying on exposed wealth as the only form of self-insurance. Richer versions of the same framework can then introduce safe assets, portfolio diversification, insurance markets, adaptation, and public transfers one at a time. This structure makes it possible to quantify which parts of the climate protection gap can be closed through private financial mechanisms and which parts remain as residual exposure requiring public policy. The contribution of the framework is therefore not only to characterize the distributional effects of current heterogeneous climate risk, but also to provide a tractable environment for evaluating alternative mechanisms of risk sharing as climate risk evolves.

8. CONCLUSION

This paper examines how spatially heterogeneous climate tail risk reshapes household wealth accumulation when capital serves both as the vehicle for saving and as the object at risk. By constructing a census-tract-level distribution of disaster-induced depreciation from observed loss data and embedding the estimated destruction process into a heterogeneous-agent model with wealth-proportional exposure, I isolate a mechanism that differs fundamentally from standard precautionary-saving environments.

When disaster losses scale with assets, additional wealth raises not only intertemporal resources but also exposure to catastrophic loss. As a result, the marginal return to saving is endogenously reduced in high-risk environments. This mechanism generates sharply nonlinear saving behavior: at low wealth levels, households in intensive climate zones save less and cluster at the borrowing constraint; at intermediate levels, precautionary motives strengthen but remain distorted by exposure. Importantly, these differences arise even when current income and realized shocks are identical, reflecting differences in the distribution of future tail risk.

Quantitatively, the aggregate capital stock responds only modestly to empirically calibrated disaster risk. However, distributional adjustments are large. Climate tail risk

primarily compresses the wealth distribution toward the lower tail, increasing the mass of hand-to-mouth households and raising inequality. Aggregate smoothness, therefore, masks substantial micro-level vulnerability.

The results also highlight that the economic consequences of climate risk depend not only on the frequency of disasters but also on the financial structure through which households manage exposure. When the primary saving asset is itself vulnerable to destruction, self-insurance becomes inherently limited. Policies and financial innovations that decouple wealth accumulation from physical exposure may therefore generate substantial welfare and distributional gains even when aggregate output effects remain modest.

While the analysis focuses on existing spatial heterogeneity in disaster exposure, the results suggest that any future increase in climate-related tail risk would likely operate through similar balance-sheet channels. The mechanism identified here implies that a higher destruction risk may reshape wealth distribution and increase financial fragility, even when aggregate effects remain modest.

An important direction for future research is to enrich the model's financial structure. The present framework represents a lower bound on available insurance opportunities: households can only self-insure by accumulating an asset that is itself exposed to destruction. In practice, households also hold financial assets, purchase insurance, and benefit from public support. Introducing safe assets, insurance markets, and financial intermediaries would enable the study of how climate risk is shared across households, firms, and regions, and quantify the extent to which private markets can attenuate climate-induced wealth compression.

More broadly, such extensions would provide a quantitative framework for studying the climate protection gap. A central policy question is how much disaster risk can be absorbed through private precautionary saving, financial diversification, insurance contracts, and adaptation investments, and how much ultimately remains for the public sector. As climate change increases the frequency and severity of extreme events, risks that are currently diversifiable may become increasingly systemic, expanding the role of government intervention. Quantifying this transition would help identify both the possibilities and the limits of private climate-risk management and provide a clearer benchmark for climate policy in a warming world.

REFERENCES

- Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., and Moll, B. (2022). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The review of economic studies*, 89(1):45–86.
- Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–684.
- Ambler, S. and Paquet, A. (1994). Stochastic Depreciation and the Business Cycle. *International Economic Review*, 35:101–116.
- Angeletos, G.-M. (2007). Uninsured Idiosyncratic Investment Risk and Aggregate Saving. *Review of Economic dynamics*, 10(1):1–30.
- 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.
- Barro, R. J. (2009). Rare Disasters, Asset Prices, and Welfare Costs. *American Economic Review*, 99(1):243–264.
- Benhabib, J., Bisin, A., and Luo, M. (2019). Wealth Distribution and Social Mobility in the US: A Quantitative Approach. *American Economic Review*, 109(5):1623–1647.
- Benhabib, J., Bisin, A., and Zhu, S. (2011). The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents. *Econometrica*, 79(1):123–157.
- Benhabib, J., Bisin, A., and Zhu, S. (2015). The Wealth Distribution in Bewley Economies with Capital Income Risk. *Journal of Economic Theory*, 159:489–515.
- Berkeley Earth (2026). Global Temperature Report for 2025. <https://berkeleyearth.org/global-temperature-report-for-2025/>. Accessed February 17, 2026.
- Bilal, A. and Rossi-Hansberg, E. (2023). Anticipating Climate Change across the United States. Technical report, National Bureau of Economic Research.
- Board of Governors of the Federal Reserve System (2024). Survey of Household Economics and Decisionmaking. Technical report, Federal Reserve System. Microdata from the 2021–2024 survey waves.
- 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.
- Center for Emergency Management and Homeland Security (2023). Spatial hazard events and losses database for the united states (sheldus[®]).
- Chen, H., Joslin, S., and Tran, N.-K. (2012). Rare Disasters and Risk Sharing with Heterogeneous Beliefs. *The Review of Financial Studies*, 25(7):2189–2224.
- Federal Emergency Management Agency (2025). National Risk Index Data. Technical report, U.S. Department of Homeland Security. December 2025 release, Version 1.20.
- Fried, S. (2022). Seawalls and Stilts: A Quantitative Macro Study of Climate Adaptation. *The Review of Economic Studies*, 89(6):3303–3344.
- Fried, S. (2026). A Macro Study of the Unequal Effects of Climate Change. Forthcoming.

- Gabaix, X. (2012). Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance. *The Quarterly journal of economics*, 127(2):645–700.
- 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. Technical report, Working Paper.
- Gallagher, J. and Hartley, D. (2017). Household Finance after a Natural Disaster: The Case of Hurricane Katrina. *American Economic Journal: Economic Policy*, 9(3):199–228.
- Gomes, J. F., Grotteria, M., and Wachter, J. A. (2019). Cyclical Dispersion in Expected Defaults. *The Review of Financial Studies*, 32(4):1275–1308.
- Gourio, F. (2012). Disaster Risk and Business Cycles. *American Economic Review*, 102(6):2734–2766.
- Groen, J. A., Kutzbach, M. J., and Polivka, A. E. (2020). Storms and Jobs: The Effect of Hurricanes on Individuals’ Employment and Earnings over the Long Term. *Journal of Labor Economics*, 38(3):653–685.
- Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.
- Hsiang, S. M. and Jina, A. S. (2014). The Causal Effect of Environmental Catastrophe on Long-run Economic Growth: Evidence from 6,700 Cyclones. Technical report, National Bureau of Economic Research.
- Indaco, A. and Ortega, F. (2024). Adapting to Climate Risk? Local Population Dynamics in the United States. *Economics of Disasters and Climate Change*, 8(1):61–106.
- Integrated Research on Disaster Risk (2014). *Peril Classification and Hazard Glossary*. Integrated Research on Disaster Risk, Beijing. IRDR DATA Publication No. 1.
- Intergovernmental Panel on Climate Change (IPCC) (2023). *Climate Change 2022 – Impacts, Adaptation and Vulnerability: Working Group II Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.
- Jones, J. B. and Neelakantan, U. (2023). Portfolios Across the U.S. Wealth Distribution. Economic Brief 23-39, Federal Reserve Bank of Richmond. Accessed May 30, 2026.
- Moll, B., Rachel, L., and Restrepo, P. (2022). Uneven Growth: Automation’s Impact on Income and Wealth Inequality. *Econometrica*, 90(6):2645–2683.
- National Weather Service (2021). Nwsi 10-1605: Storm data preparation. Technical Report NWSI 10-1605, U.S. Dept. of Commerce, NOAA, National Weather Service. Accessed 2025-10-16.
- Pleninger, R. (2022). Impact of Natural Disasters on the Income Distribution. *World Development*, 157:105936.
- Soloveichi, R. (2021). Tracking Cultivated Assets in Measures of Capital. Technical report, US Department of Commerce, Bureau of Economic Analysis.

- Strobl, E. (2011). The Economic Growth Impact of Hurricanes: Evidence from US Coastal Counties. *Review of Economics and Statistics*, 93(2):575–589.
- Swiss Re Institute (2025). Natural catastrophes in 2024: Climate risk and the protection gap. Technical Report 1/2025, Swiss Re Institute, Zurich. Available at <https://sigma.swissre.com/system/files/sigma>
- U.S. Bureau of Economic Analysis (2025). Table 1.1. current-cost net stock of fixed assets and consumer durable goods. Fixed Assets Accounts (iTables). Accessed 2025-02-12.
- U.S. Census Bureau (2020). TIGER/Line Shapefiles: Census Tracts. 2020 Census tract boundaries.

APPENDIX A. MODEL

A.1. Continuous-Time Representation. To solve the model numerically, I recast the household problem in continuous time following the finite-difference approach of [Achdou et al. \(2022\)](#). The continuous-time representation provides a convenient characterization of the stationary equilibrium through a coupled system of Hamilton–Jacobi–Bellman (HJB) and Kolmogorov Forward (KFE) equations.

I first present the general formulation based on the J -state disaster process $x \in \{e_1^{(J)}, \dots, e_J^{(J)}\}$. I then specialize the framework to the baseline specification with independent climate zones weighted by population shares λ_z .

Let Q denote the continuous-time generator associated with the Markov process for x . When the transition process is specified in discrete time through the transition matrix Π , I define

$$Q = \frac{\Pi - I_J}{\Delta t},$$

where Δt denotes the discretization interval.

Given disaster state x , the depreciation and labor-productivity realizations are

$$s(x) = \mu^\top x, \quad \ell(x) = \ell^\top x.$$

The household's continuous-time asset accumulation equation is

$$\dot{a} = w\ell(x) + (r^K - \delta - s(x))a - c,$$

subject to the borrowing constraint

$$a \geq -\phi.$$

A.2. Stationary HJB Equation. Let $V(a, x)$ denote the stationary value function associated with disaster state x . The stationary HJB equation is

$$\rho V(a, x) = \max_c \left\{ \frac{c^{1-\gamma}}{1-\gamma} + V_a(a, x) [w\ell(x) + (r^K - \delta - s(x))a - c] + \sum_{x'} Q(x, x') V(a, x') \right\}.$$

The first-order condition implies

$$c(a, x) = (V_a(a, x))^{-1/\gamma}.$$

A.3. Stationary Distribution. Given optimal household policies, the stationary distribution $g(a, x)$ solves the Kolmogorov Forward Equation associated with the optimal savings dynamics and generator matrix Q :

$$0 = -\frac{\partial}{\partial a} [\dot{a}(a, x)g(a, x)] + \sum_{x'} Q(x', x)g(a, x').$$

Aggregate capital is obtained by integrating over household assets and disaster states:

$$K = \sum_x \int a g(a, x) da.$$

A.4. Block Decomposition with Climate Zones. In the baseline specification, the general J -state process is decomposed into climate zones and within-zone disaster-event realizations. Specifically, each state x corresponds to a pair (z, u) , where

$$z \in \mathcal{Z} = \{1, \dots, Z\}, \quad u \in \mathcal{U} = \{1, \dots, U\}.$$

Conditional on climate zone z , households transition across disaster-event realizations according to the zone-specific transition matrix

$$\Pi^z = (\pi_{uu'}^z)_{u, u' \in \mathcal{U}}.$$

The associated continuous-time generator matrix is

$$Q^z = \frac{\Pi^z - I_U}{\Delta t}.$$

The disaster and labor-productivity realizations are given by

$$s_t^i = \mu_u^z, \quad y_t^i = w_t \ell_u^z.$$

Let $g^z(a, u)$ denote the stationary distribution within climate zone z . Aggregate capital is then obtained by combining the independent zone-specific distributions using population weights λ_z :

$$K = \sum_{z \in \mathcal{Z}} \lambda_z \sum_{u \in \mathcal{U}} \int a g^z(a, u) da,$$

where

$$\sum_{z \in \mathcal{Z}} \lambda_z = 1.$$

Thus, the baseline spatial specification can be interpreted as a block decomposition of the general J -state process, where each climate zone defines an independent transition block weighted by its population share.

A.5. Normalization of Labor Productivity. To preserve a non-degenerate stationary wealth distribution, I introduce a small state-dependent labor productivity component linked to the same disaster process governing capital destruction.

Let

$$\ell = (\ell_1, \dots, \ell_J)^\top$$

denote the vector of labor-productivity realizations associated with the general J -state process. Let $\bar{\pi}$ denote the stationary distribution satisfying

$$\bar{\pi}^\top Q = 0, \quad \sum_{j=1}^J \bar{\pi}_j = 1.$$

The labor-productivity vector is normalized so that its stationary mean equals one:

$$\bar{\pi}^\top \ell = 1.$$

This normalization ensures that aggregate labor income equals the equilibrium wage rate w , so that the labor-productivity component generates cross-sectional heterogeneity without affecting aggregate quantities or the firm's problem.

In the baseline spatial specification, the stationary distribution is obtained by aggregating the zone-specific stationary distributions using population weights. Let $\bar{\pi}^z$ denote the stationary distribution associated with transition matrix Π^z . The normalization condition becomes

$$\sum_{z \in \mathcal{Z}} \lambda_z (\bar{\pi}^z)^\top \ell = 1.$$

In the quantitative implementation, a subset of productivity realizations is calibrated using existing empirical estimates of disaster-induced earnings losses, while one realization is determined residually to satisfy the normalization condition.

Let $r \in \mathcal{U}$ denote the residual event state. Then

$$\sum_{z \in \mathcal{Z}} \lambda_z \sum_{u \in \mathcal{U}} \bar{\pi}_u^z \ell_u = 1.$$

Solving for the residual productivity realization gives

$$\ell_r = \frac{1 - \sum_{z \in \mathcal{Z}} \lambda_z \sum_{u \in \mathcal{U} \setminus \{r\}} \bar{\pi}_u^z \ell_u}{\sum_{z \in \mathcal{Z}} \lambda_z \bar{\pi}_r^z}.$$

In the baseline calibration, the low-damage state is chosen as the residual state, so that $r = 1$.

A.6. Numerical Solution. The stationary equilibrium is computed using an implicit finite-difference scheme following [Achdou et al. \(2022\)](#). For a given interest rate r^K , I solve the stationary HJB equation using upwind discretization of the asset drift. The stationary distribution is then recovered from the associated Kolmogorov Forward Equation.

The equilibrium interest rate is updated iteratively until aggregate household asset supply equals firm capital demand:

$$K^S(r^K) = K^D(r^K).$$

The numerical implementation therefore combines:

- (1) solution of the HJB equation for optimal household policies,
- (2) computation of the stationary distribution through the KFE,
- (3) aggregation across climate zones using population weights λ_z ,
- (4) market-clearing iteration on the equilibrium interest rate.

APPENDIX B. DATA

This appendix describes the construction of the disaster-induced depreciation rates used to estimate the stochastic depreciation processes in the model. The procedure combines county-level disaster losses from SHEL DUS, county capital stocks constructed from BEA fixed assets, and tract-level exposure measures from FEMA’s National Risk Index (NRI). The final output is a balanced tract-year panel of localized depreciation shocks for 2002–2022.

B.1. County-Level Disaster Losses and Capital Stocks. The county-level disaster-induced depreciation rate is defined as

$$d_t^j \equiv \frac{\text{NDL}_t^j}{K_{t-1}^j}, \quad (23)$$

where NDL_t^j denotes direct disaster losses in county j and year t , and K_{t-1}^j is the county capital stock at the beginning of the period.

Capital Stocks. The Bureau of Economic Analysis (BEA) reports annual national estimates of the current-cost net stock of fixed assets and consumer durable goods, denoted K_t^N .²² Because comparable capital stock measures are not available at the county level, I construct county capital stocks by allocating aggregate national capital proportionally to county GDP shares.

This allocation implicitly assumes that counties share a common capital-output ratio. Under a Cobb-Douglas production technology with common factor shares, output shares coincide with capital shares, implying

$$K_t^i = \frac{\text{GDP}_t^i}{\text{GDP}_t^N} K_t^N, \quad (24)$$

where GDP_t^i and GDP_t^N denote county and national GDP, respectively. The procedure preserves aggregation consistency by construction:

$$\sum_i K_t^i = K_t^N.$$

This approach abstracts from cross-county differences in capital intensity. Consequently, capital stocks may be understated in relatively capital-intensive counties and overstated in less capital-intensive counties.²³

To assess the potential magnitude of this measurement issue, I construct a proxy for county capital intensity using BEA county-industry accounts. The proxy is defined as the share of county GDP originating from sectors typically characterized by high physical-capital intensity: mining and oil extraction, utilities, construction, manufacturing, and transportation and warehousing. These sectors rely disproportionately on structures, machinery, industrial equipment, energy infrastructure, and transportation networks relative to service-oriented industries.

I then compare this proxy with FEMA National Risk Index measures to evaluate whether the resulting measurement error is systematically related to disaster exposure. The Spearman rank correlation between county capital intensity and the FEMA National Risk Index composite score is 0.074, while the correlation with expected annual losses is 0.064. These correlations remain small, suggesting that potential measurement error in county capital stocks does not mechanically align with geographic disaster exposure.

²²Table 1.1, Current-Cost Net Stock of Fixed Assets and Consumer Durable Goods, (U.S. Bureau of Economic Analysis, 2025).

²³This measurement issue mechanically biases estimated depreciation rates upward in low-capital-intensity counties and downward in high-capital-intensity counties.

Natural-Disaster Losses (SHELDUS). Economic losses are drawn from the Spatial Hazard Events and Losses Database for the United States (SHELDUS), developed by [CEMHS](#). SHELDUS reports all county-level events from 1960–2022 that meet at least one of the following criteria: (i) reported property or crop damage of at least \$1, (ii) at least one injury, or (iii) at least one fatality. For each event, SHELDUS reports direct economic losses disaggregated into property and crop components. Following [National Weather Service \(2021\)](#), property losses encompass damages to private structures, vehicles, and public infrastructure/facilities, whereas crop losses refer to damages to agricultural production intended for harvest or food.

When an event is reported as affecting multiple counties or multiple hazard types, SHELDUS allocates the reported totals equally across all listed counties and across all associated hazard categories. This equal-splitting rule applies to property losses, crop losses, injuries, and fatalities. When source documents provide more detailed geographic information, SHELDUS uses that data instead of equal-splitting. For events reported by National Weather Service (NWS) forecast zones, SHELDUS first maps the forecast zones to counties; when a zone spans multiple counties, SHELDUS again distributes losses evenly across those counties.

Because cultivated assets are not capitalized in the official BEA fixed-assets accounts, the baseline definition of the disaster-induced depreciation rate excludes crop losses from NDL_t^i and uses only property losses.²⁴

The analysis window is restricted to 2002–2022 to align SHELDUS losses with county-level macroeconomic controls. Dollar values are in current dollars.

SHELDUS distinguishes 19 hazard types. For parsimony, I aggregate them into eight broader categories following [IRDR \(2014\)](#). Appendix Table [B.1](#) reports the mapping.

B.2. Geographic Harmonization. U.S. counties are identified by unique Federal Information Processing Standards (FIPS) codes. The Bureau of Economic Analysis (BEA), however, aggregates certain counties under modified or composite FIPS codes for its local-area statistics. To ensure consistency across datasets, I apply the same aggregation structure to the SHELDUS loss data.

For counties that experienced boundary changes—such as splits, mergers, or redefinitions—over the sample period, I adopt a harmonized coding scheme that treats each county as if it had maintained a consistent boundary (merged or separate) throughout. This approach ensures temporal consistency and avoids making assumptions about the allocation of GDP or disaster damages across newly formed or reconfigured counties.

A complete list of FIPS code adjustments is provided in Appendix Table [B.2](#).

BEA groupings — grouped by <i>to</i> FIPS		
From	To	Aggregate Name
[15005, 15009]	15901	Kalawao + Maui

²⁴[Soloveichi \(2021\)](#) provides estimates of cultivated assets up to 2019. Appendix Figure [B.2](#) reports estimated depreciation rates incorporating crop losses.

[51003, 51540]	51901	Albemarle + Charlottesville
[51005, 51580]	51903	Alleghany + Covington
[51015, 51790, 51820]	51907	Augusta + Staunton + Waynesboro
[51031, 51680]	51911	Campbell + Lynchburg
[51035, 51640]	51913	Carroll + Galax
[51053, 51570, 51730]	51918	Dinwiddie + Colonial Heights + Petersburg
[51059, 51600, 51610]	51919	Fairfax + Fairfax City + Falls Church
[51069, 51840]	51921	Frederick + Winchester
[51081, 51595]	51923	Greensville + Emporia
[51089, 51690]	51929	Henry + Martinsville
[51095, 51830]	51931	James City + Williamsburg
[51121, 51750]	51933	Montgomery + Radford
[51143, 51590]	51939	Pittsylvania + Danville
[51149, 51670]	51941	Prince George + Hopewell
[51153, 51683, 51685]	51942	Prince William + Manassas + Manassas Park
[51161, 51775]	51944	Roanoke + Salem
[51163, 51530, 51678]	51945	Rockbridge + Buena Vista + Lexington
[51165, 51660]	51947	Rockingham + Harrisonburg
[51175, 51620]	51949	Southampton + Franklin City
[51177, 51630]	51951	Spotsylvania + Fredericksburg
[51191, 51520]	51953	Washington + Bristol
[51195, 51720]	51955	Wise + Norton
[51199, 51735]	51958	York + Poquoson

County split (AK changes) — grouped by *from* FIPS

From	To	Aggregate Name
02232	[02105, 02230]	*Skagway–Hoonah–Angoon + *Skagway–Yakutat + Skagway + Hoonah–Angoon
02280	[02195, 02275]	*Wrangell–Petersburg + Wrangell + Petersburg
02261	[02063, 02066]	*Valdez–Cordova + Chugach + Copper River
02201	[02130, 02198]	*Prince of Wales–Outer Ketchikan + Prince of Wales–Hyder + *Ketchikan + Ketchikan Gateway

County merged — grouped by *to* FIPS

From	To	Old vs. New County Name
[46113]	46102	Shannon to Oglala Lakota
[02270]	02158	Wade Hampton to Kusilvak

[51515]	51019	Bedford City to Bedford
[51560]	51005	Clifton Forge to Alleghany

TABLE B.2. FIPS Remapping

Notes: This table summarizes the remapping of county-level Federal Information Processing Standards (FIPS) codes as reflected in the SHELDUS database. The top panel reports county aggregations used by the Bureau of Economic Analysis (BEA), grouped by the resulting (“to”) FIPS code and associated aggregated name. The second panel lists instances where historical counties have been split into multiple entities, with entries grouped by the original (“from”) FIPS codes. The third panel reports cases where counties were administratively dissolved and merged into already existing counties. In each case, the “from” FIPS code corresponds to a former independent county that ceased to exist, and the “to” FIPS code identifies the absorbing county. These changes reflect administrative reclassifications. Counties marked with an asterisk (*) no longer exist in their original form but are included for consistency with historical data. These administrative changes have been implemented in the construction of the dataset and are consistent with the modifications adopted by SHELDUS and BEA datasets.

Connecticut requires a separate geographic harmonization because the most recent FEMA National Risk Index (NRI) adopts the revised post-2023 FIPS convention for the state. To ensure consistency across datasets, all Connecticut counties are aggregated into a single state-level unit (FIPS 09000). This harmonization is applied throughout the sample period, including years prior to 2023.

B.3. Major vs. Minor Events. I partition the loss data into two groups: (i) *major events* and (ii) *minor events*. SHELDUS identifies major events as those associated with a “large-scale, high-impact” disaster and relies on two external sources: (i) NCEI’s Billion-Dollar Weather and Climate Disasters (BDD) database and (ii) FEMA’s Presidential Disaster Declarations (PDD).²⁵

While all NCEI billion-dollar events are contained within the aggregate SHELDUS dataset, documented discrepancies exist between SHELDUS and the PDD records. In particular, [Fried \(2022\)](#) and [Gallagher \(2014\)](#) report omitted events in SHELDUS relative to PDD. To mitigate this undercoverage, I reconcile the two “major” sources as follows. For each county–year–hazard cell flagged as major by either source, I define the major-event loss as the maximum of the two reported amounts:

$$\text{NDL}_{i,t,h}^{\text{major}} = \max\{\text{NDL}_{i,t,h}^{\text{NCEI}}, \text{NDL}_{i,t,h}^{\text{PDD}}\}.$$

This procedure preserves the full set of NCEI major events while correcting for documented undercoverage of PDD events in SHELDUS.^{26,27}

²⁵SHELDUS also includes the Global Disaster Identifier Number (GLIDE). GLIDE-coded events are largely subsumed by the NCEI major-event series in this setting.

²⁶In the assembled dataset, this reconciliation increases the number of county–year–hazard major-event observations by roughly 20% relative to relying on NCEI alone.

²⁷The discrepancy between SHELDUS and PDD records declines over time as SHELDUS coverage improves, implying that the supplementation primarily affects the early 2000s (see also [Gallagher, 2023](#)).

Category	Hazards
Geophysical & Mass Movement	Earthquake
	Volcano
	Tsunami/Seiche
	Avalanche
	Landslide
Flooding & Coastal Hazards	Flooding
	Coastal
Drought/Heat	Drought
	Heat
Wildfire	Wildfire
Winter Weather	Winter Weather
Hurricane/Tropical Storm	Hurricane
	Tropical Storm
Tornado	Tornado
Severe Weather	Severe Storm/Thunder Storm
	Lightning
	Hail
	Fog
	Wind

TABLE B.1. Hazard Categories

Notes: This table reports the hazard aggregation used throughout the paper. The hazards listed in the second column correspond to the original hazard classifications reported in the Spatial Hazard Events and Losses Database for the United States (SHELDUS; CEMHS, 2023). The broader hazard categories reported in the first column are constructed by the author following the classification framework of the Integrated Research on Disaster Risk (IRDR, 2014).

Minor events are defined residually as all SHELDUS events not flagged as major by either NCEI or PDD. Finally, I adjust county-year-hazard loss aggregates to incorporate the supplemented PDD-based major events.

The category “Drought and Heat” primarily affects crops, which are excluded from the baseline definition of economic losses because cultivated assets are not capitalized in BEA fixed-asset accounts. Appendix Figure B.2 reports depreciation rates for 2002–2019 period, when both crop losses and cultivated assets are included. The most notable difference appears in the distributions for “Drought and Heat”, for both major and minor events. In practice, drought conditions often increase the likelihood of severe wildfires or floods, and major disasters typically co-occur with drought episodes. Because SHELDUS records all hazard types contributing to a multi-hazard event, droughts frequently appear as major events even though, on their own, they do not damage physical capital to the same extent as hurricanes, tropical storms, or wildfires.




FIGURE B.1. Disasters' Depreciation Rate Distribution

Notes: The figure displays the separate distributions of the logarithm of the empirical depreciation rate, $\log(d_t^i)$, for each hazard type h , distinguishing major (red) and minor (blue) disasters, across U.S. counties i , and years t from 2001 to 2022. For each hazard h , the left panel reports the Dirac mass at zero, corresponding to county-year observations with no recorded capital losses. The right panel plots the kernel density of $\log(d_t^i)$ conditional on $d_t^i > 0$, with major and minor events shown separately. A disaster is classified as major if it appears in either NCEI's Billion-Dollar Weather and Climate Disasters (BDD) database or FEMA's Presidential Disaster Declarations (PDD). Densities are weighted by the relative frequency of each disaster type (major/minor) among positive-loss county-year observations. The continuous component integrates to the empirical share of strictly positive-loss observations within each hazard and disaster type; together with the Dirac mass at zero, the total probability mass for each distribution sums to one.




FIGURE B.2. Disasters' Depreciation Rate Distribution with Crop Damages

Notes: The figure displays the distribution of the logarithm of the depreciation rate, $\log(\delta)$, across U.S. counties by hazard type over the period 2001-2019. Damages included in the depreciation rate are property and crop damages. The capital is corrected accordingly, taking into account an estimate of cultivated assets from Soloveichi (2021). For each hazard, the left panel shows the Dirac mass at zero—i.e., the share of county-year observations with no capital loss—while the right panel reports kernel density estimates conditional on $\delta > 0$. Red and blue denote major and minor disaster designations, respectively. For each hazard, the densities integrate to one after accounting for the zero Dirac mass, so that the total mass (Dirac plus continuous density) equals unity.

B.4. Heterogeneity in Disasters Data. The county-level dataset is a balanced panel covering 3,103 U.S. counties from 2002–2022. It combines information on natural hazard events from the SHELDUS database with county-level capital stock estimates constructed by allocating BEA national fixed assets according to local GDP shares. Together, these components yield the natural-disaster-induced depreciation rate d , defined in Eq. (20).²⁸ The dataset captures both the frequency and severity of natural disasters across eight hazard categories: “Hurricane and Tropical Storm”, “Wildfire”, “Drought and Heat”, “Tornado”, “Flooding and Coastal Hazards”, “Severe Weather”, “Winter Weather”, and “Geophysical and Mass Movement” events.

A central descriptive feature of the data is the distinction between disaster frequency and disaster severity. As shown in Table B.3, “Severe Weather” accounts for a substantial share of all recorded events but is associated with a relatively modest average depreciation rate. By contrast, “Hurricanes and Tropical Storms” and “Drought and Heat” occur less frequently at the national level yet display considerably higher mean depreciation conditional on positive losses. These categories also exhibit a higher fraction of events classified as major disasters.²⁹ Wildfires similarly generate elevated conditional losses in the counties where they occur, despite a smaller national count of events.³⁰

The aggregate distribution of depreciation rates d_t^i is depicted in Figure B.3, with each hazard’s contribution represented by a color. The overall distribution is composed of two parts: a Dirac mass at zero, corresponding to county–year–hazard observations with no recorded capital loss, and a continuous component over strictly positive depreciation rates. Approximately 80% of observations are zero, highlighting the sparsity of damaging events at the county–year–hazard level. Conditional on $d_t^i > 0$, the distribution of $\log(d_t^i)$ is moderately right-skewed.

Hazard composition varies systematically across the distribution. Lower percentiles are dominated by “Severe Weather” and “Flooding and Coastal Hazards”. At the other extreme, Table B.4 shows that “Flooding and Coastal Hazards” together with “Hurricane and Tropical Storm” account for about 50% of events in the top 1% of the depreciation distribution. Overall, “Hurricane and Tropical Storm” accounts for less than 1% of all natural disasters in the sample. Nevertheless, its share rises monotonically toward the upper tail of the depreciation distribution of d , reaching 23%, whereas most other hazards

²⁸In the dataset, five observations (out of 521,304) exhibit a natural-disaster-induced depreciation rate exceeding one, which is not economically plausible. I correct these cases by capping their values at 0.99. These outliers likely stem from two sources: (i) the assumption that capital stock is proportional to county-level GDP, and (ii) the allocation method used by SHELDUS, which distributes economic losses equally across all affected counties in multi-county events.

²⁹Major events are defined as disasters that either exceed USD 1 billion in direct economic damages, following the Billion-Dollar Disaster (BDD) classification, or trigger a federal emergency response under a Presidential Disaster Declaration (PDD). Detailed definitions are provided in Appendix B.3.

³⁰The category “Drought and Heat” primarily affects crops, which are excluded from the baseline definition of economic losses because cultivated assets are not capitalized in BEA fixed-asset accounts. In practice, drought conditions often increase the likelihood of severe wildfires or floods, and major disasters typically co-occur with drought episodes. Because SHELDUS records all hazard types contributing to a multi-hazard event, droughts frequently appear as major events even though, on their own, they do not damage physical capital to the same extent as hurricanes, tropical storms, or wildfires.

Hazard Type	Records	Share Major (%)	Mean d (%)	Mean $d > 0$ (%)	Std. Dev. $d > 0$
Hurricane and Tropical Storm	3679	80	0.032	0.89	0.056
Wildfire	2933	68	0.005	0.25	0.031
Drought and Heat	6424	64	0.002	0.26	0.009
Tornado	16352	42	0.011	0.07	0.010
Flooding and Coastal Hazards	49249	35	0.024	0.08	0.011
Severe Weather	384980	22	0.012	0.02	0.006
Winter Weather	20082	16	0.007	0.06	0.004
Geophysical and Mass Movement	2719	11	0.001	0.05	0.003

TABLE B.3. Statistics of SHELDUS Natural Hazards

Notes: The table reports summary statistics by hazard type for a balanced panel of 3,103 U.S. counties observed annually from 2002 to 2022. *Records* denote the total number of hazard occurrences across all counties over the sample period. *Share Major (%)* corresponds to the fraction of records classified as major disasters under FEMA and PDD designations. The variable d (%) measures the depreciation rate of the county’s capital stock attributable to natural disasters. *Mean d* is computed over all county-year observations, including those without disaster (i.e., $d \geq 0$). *Mean $d > 0$* and *Std. Dev. $d > 0$* restrict the sample to county-years in which a natural disaster occurred and destroyed at least \$1 of capital.

either decline or remain roughly constant. Flood-related events, by contrast, are present throughout the distribution, with substantial representation in both extreme percentiles.

Hazard Type	Total Share	Share events in the top <i>percentile</i> :		
		5%	2.5%	1%
Hurricane and Tropical Storm	0.8	11	16	23
Wildfire	0.6	4	4	4
Drought and Heat	1.3	2	3	6
Tornado	3.4	17	16	16
Flooding and Coastal Hazards	10.1	32	31	27
Severe Weather	79.1	22	17	10
Winter Weather	4.1	11	11	12
Geophysical and Mass Movement	0.6	1	1	1

TABLE B.4. Shares in the Upper Tail of the Depreciation Distribution

Notes: The table summarizes the composition of hazard events in the upper tail of the aggregate distribution of county-level depreciation rates. The sample is a balanced panel of 3,103 U.S. counties observed annually from 2002 to 2022. “Total Share” reports the fraction of all records attributable to each hazard type. “Share of events in the top” 5%, 2.5%, and 1% report, respectively, the fraction of events of each hazard type among all county-year observations whose depreciation rate d_t^i lies in the top 5%, 2.5%, and 1% of the aggregate distribution of d_t^i . These statistics characterize which hazards account for the most severe realized capital losses.

To illustrate this spatial heterogeneity, Figure B.4 maps county-level exposure for 2002–2022. For each hazard category, the figure reports the within-hazard percentile rank of the cumulative number of events recorded in each county. Counties are shaded by




FIGURE B.3. Disasters' Depreciation Rate Distribution

Notes: The figure displays the distribution of the logarithm of the empirical disaster-induced depreciation rate, $\log(d_t^i)$, across all hazard types h , U.S. counties i , and years t from 2002 to 2022. The left panel reports the Dirac mass at zero, representing the share of county-year-hazard observations with no recorded capital losses. The right panel plots the kernel density of $\log(d_t^i)$ conditional on $d_t^i > 0$, with colors distinguishing each hazard type. Each density is weighted by its corresponding hazard's frequency among positive-loss observations. The continuous component, therefore, reflects the aggregate empirical distribution conditional on $d_t^i > 0$, and integrates to the corresponding share of observations; together with the Dirac mass at zero, the full distribution has total probability mass equal to one. Dashed and dotted vertical lines mark the sample mean and one standard deviation above the mean (+1SD), respectively.

their percentile rank within each hazard, which facilitates comparison of spatial patterns across hazard types while abstracting from differences in aggregate frequency.

The maps reveal apparent geographic clustering in the exposure to specific hazard types. Hurricanes and tropical storms are concentrated along the Atlantic and Gulf Coasts, with intensity increasing in the Southeastern U.S., consistent with prevailing storm tracks and landfall patterns. Wildfires are concentrated in the Western U.S., particularly in California, Oregon, and parts of the Mountain West, reflecting the intersection of dry climate, vegetation, and topography. Droughts and heatwaves are most prevalent in the Southwest and central Great Plains.

Flooding and coastal hazards exhibit a pattern similar to the map of hurricanes and tropical storms, but also extend inland along major river systems, including the Mississippi Basin.

Tornadoes and severe weather events exhibit broad spatial coverage, spanning much of the Midwest, South, and Eastern U.S., consistent with meteorological data on storm activity zones such as “Tornado Alley”. Winter weather event records are most concentrated




FIGURE B.4. Spatial Distribution of Natural Hazard Records across U.S. Counties

Notes: The figure displays the spatial distribution of disaster incidence by hazard types across U.S. counties from 2002 to 2022. For each hazard, shading reflects the percentile rank of the total number of recorded events per county. Darker shades indicate counties with higher relative frequency within each hazard category.

in the Upper Midwest and Northeast, in line with patterns of seasonal snowstorms and cold extremes. Finally, “Geophysical and Mass Movement” hazards (e.g., earthquakes, landslides) are highly localized, with records primarily in a small number of western counties.

Overall, the figure highlights substantial heterogeneity in geographic exposure across hazard categories. Natural disasters differ not only in their spatial concentration but also in the magnitude of the losses they generate. Hurricanes and tropical storms are associated with substantially higher depreciation rates than other hazards, with conditional average losses more than three times those associated with droughts, heatwaves, or wildfires. By contrast, severe and winter weather events occur more frequently but generate comparatively small depreciation rates conditional on their occurrence.

These patterns indicate substantial heterogeneity in disaster risk. Both the frequency and the severity of realized depreciation shocks vary across hazards and across space, suggesting that households in different locations face systematically different distributions of disaster-induced capital losses.

TABLE B.5. Financial Impact Across SHED Survey Waves

	Wave 2024	Wave 2023	Wave 2022	Wave 2021
<i>Panel A. Number of respondents</i>				
No	12,295	11,400	11,667	11,874
Yes	2,524	2,137	1,257	1,531
Slightly	1,567	1,346		
Moderately	715	579		
Substantially	242	212		
<i>Panel B. Share of respondents in percent</i>				
No	79	81	89	87
Yes	21	19	11	13
Slightly	13	12		
Moderately	6	5		
Substantially	2	2		

Notes: Each column corresponds to a wave of the Federal Reserve’s Survey of Household Economics and Decisionmaking (SHED). Respondents indicate whether, in the past year, they experienced a financial impact from natural disasters (“Yes” or “No”). In Waves 2024 and 2023, respondents who answered “Yes” were further asked to classify the severity of the impact as slight, moderate, or substantial. Panel A reports respondent counts; Panel B reports corresponding shares. Empty cells indicate that follow-up severity questions were not asked in earlier waves. The exact questions correspond to *ND0* in waves 2024 and 2023, and to *ND1* (*a - d*) in waves 2022 and 2021.

B.5. Household Exposure Evidence from SHED.

B.6. Construction of Census-Tract Depreciation Rates.

B.6.1. Census-Tracts and FEMA National Risk Index Measures.

Census-Tracts. Census tracts are a relatively permanent subdivision of counties and constitute the next administrative level below counties in the United States. Each county is partitioned into mutually exclusive census tracts, which are designed to support the analysis of small-area demographic and economic outcomes. On average, census tracts contain approximately 4,000 residents and are intended to remain stable over time, with boundary changes occurring only infrequently in response to major population shifts identified in the decennial Census. I therefore adopt the 2020 Census tract delineation from the U.S. Census Bureau and treat tract boundaries as time-invariant over the sample period.³¹

In the resulting data, the 84,093 census tracts indeed contain, on average, roughly 4,000 residents, consistent with their intended design. Counties vary substantially in the number of tracts they contain, reflecting heterogeneity in population size and density: the average county contains about 27 census tracts, with counts ranging from a single tract in sparsely populated counties to nearly 2,500 tracts in the largest urban counties. This institutional structure permits disaster damages to be highly localized within counties while remaining consistent with county-level aggregation in the observed loss data.

National Risk Index. I use FEMA’s National Risk Index (NRI) at the census-tract level to construct a disaggregated dataset capturing both baseline hazard exposure and socioeconomic vulnerability.

The NRI combines three conceptually distinct components to measure baseline climate risk for a given spatial unit, ranging from broad administrative regions (e.g., states or counties) to fine geographic subdivisions such as census tracts. Expected Annual Losses (EAL) quantify the long-run average economic losses from natural hazards and capture baseline hazard exposure. The Social Vulnerability Index (SVI) measures the susceptibility of local populations to harm based on demographic and socioeconomic characteristics, reflecting differential sensitivity and adaptive capacity. The Community Resilience Index (CRI) captures a community’s ability to respond to and recover from disasters, incorporating factors such as housing conditions, infrastructure, and institutional capacity. The NRI composite score aggregates these components into a single measure of baseline risk. Following FEMA’s definition, risk is constructed as expected annual losses scaled by a community-specific risk factor that depends on social vulnerability and resilience,

$$\text{NRI}_i = \text{EAL}_i \times f\left(\frac{\text{SVI}_i}{\text{CRI}_i}\right), \quad (25)$$

where $f(\cdot)$ denotes FEMA’s transformation of social vulnerability relative to community resilience into a community-level risk factor capturing the potential consequences of disaster exposure.³²

³¹Census tract shapefiles (U.S. Census Bureau, 2020) are obtained using the `tigris` package in R.

³²See Federal Emergency Management Agency (2025) for details on the normalization and aggregation of each component.

Combining these measures enables county-level expected losses to be allocated across census tracts in a manner that reflects within-county heterogeneity in both hazard exposure and vulnerability, providing a coherent foundation for disaggregating disaster impacts.

B.6.2. Census Tract Hit Probability. I construct tract-level hit probabilities by combining county-level disaster-event frequencies with measures of relative exposure across census tracts within a county.

Let $Records_t^j$ denote the number of natural disaster events observed in county j in year t . I model disaster occurrence at the county level as a Poisson process,

$$Records_t^j \sim \text{Poisson}(\lambda_j), \quad (26)$$

where λ_j is a county-specific intensity parameter capturing persistent differences in disaster exposure across counties. Under the assumption of a constant arrival rate over time, the maximum likelihood estimator of λ_j is given by the sample mean,

$$\lambda_j = \frac{1}{T} \sum_{t=1}^T Records_t^j, \quad (27)$$

with T denoting the length of the observation period. This parameter summarizes the average frequency of disasters in county j .

While disaster frequency varies across counties, exposure within counties is highly heterogeneous. To allocate county-level disaster activity across census tracts, I rely on tract-level expected annual losses (EAL) computed by FEMA's National Risk Index. For a census tract i located in county j , I define the exposure share

$$w_{i|j} = \frac{EAL_i}{\sum_{k \in j} EAL_k}, \quad (28)$$

which satisfies $\sum_{i \in j} w_{i|j} = 1$. These shares assign a larger fraction of county-level disaster intensity to tracts with higher baseline risk.

I assume that county-level disaster events are independently allocated across census tracts according to these exposure shares. By a standard Poisson thinning argument,

$$Records_t^i \sim \text{Poisson}(\lambda_i), \quad \text{where} \quad \lambda_i = w_{i|j} \lambda_j. \quad (29)$$

This construction preserves aggregation consistency, in the sense that tract-level intensities sum to the county-level intensity, $\sum_{i \in j} \lambda_i = \lambda_j$.

Finally, I define a census tract as being “hit” in a given year if it experiences at least one disaster event. Under the Poisson assumption, the unconditional probability that tract i is hit is given by

$$\mathbb{P}(\text{hit}^i) = 1 - \exp(-\lambda_i). \quad (30)$$

This probability increases with both county-level disaster frequency and tract-specific exposure, thereby spatially concentrating disaster impacts within counties.

B.6.3. *Hit Census Tracts.* I assume that each time a county j experiences a disaster event in year t , exactly one census tract within that county is hit. The affected tract is randomly drawn according to conditional tract-level exposure shares, $\omega_{i|j} = \mathbb{P}(\text{hit}^i \mid \text{hit}^j)$. As a result, counties that experience frequent disaster activity may see the same census tract hit multiple times within a year, while other events affect different tracts.

Formally, conditional on the realized number of disaster events in a county-year, Records_t^j , events are independently allocated across census tracts according to their exposure shares $\{\omega_{i|j}\}_{i \in j}$. Under this assumption, the vector of tract-level event counts satisfies

$$(\text{Records}_t^i)_{i \in j} \mid \text{Records}_t^j \sim \text{Multinomial}(\text{Records}_t^j, \{\omega_{i|j}\}_{i \in j}). \quad (31)$$

This conditional allocation preserves the observed county-year count of disaster events by construction and implies that expected tract-level event intensities are proportional to exposure shares.³³

This assumption captures the highly localized nature of natural disasters: each recorded disaster event is assumed to affect a single census tract, corresponding to an average of roughly 4,000 individuals. At the same time, repeated events within a county generate heterogeneity in exposure across tracts, consistent with the concentration of damages observed in household-level survey evidence.

In a single simulated draw, approximately 16% of the population resides in census tracts classified as hit in a given tract-year. This level of exposure is consistent with the population shares implied by the county-level depreciation distribution reported in Table 4.1. Repeating the exercise over 200 independent draws yields a remarkably stable pattern across the full panel. At the annual level, exposure exhibits slightly greater variation, with the share of the population residing in hit tracts ranging from 13% to 22% across years.³⁴

Observed county-year disaster losses NDL_t^j are then allocated across the set of census tracts classified as affected in county j during year t . The allocation procedure is designed to capture the highly localized nature of disaster damage while remaining parsimonious and aggregation-consistent.

B.6.4. *Capital Allocations and Effective Capital.* Census tract-level capital stocks are constructed by allocating county-level capital K_t^j according to fixed tract shares based on building replacement values. Specifically, each census tract's share of county capital $\kappa_{i|j}$ is defined as the ratio of its building replacement value BV_{2020}^i to the county's total building replacement value BV_{2020}^j , as measured in the 2020 Census.³⁵

$$\kappa_{i|j} \equiv \frac{BV_i^{2020}}{\sum_{k \in j} BV_k^{2020}}, \quad \sum_{i \in j} \kappa_{i|j} = 1. \quad (32)$$

³³Indeed, the multinomial distribution implies $\mathbb{E}(\text{Records}_t^i \mid \text{Records}_t^j) = \omega_{i|j} \text{Records}_t^j$.

³⁴Population is held fixed at 2020 Census levels for all census tracts.

³⁵See footnotes 10-11 of [Federal Emergency Management Agency \(2025\)](#) for the exact construction of replacement building value.

These shares $\kappa_{i|j}$ are held constant over time and applied to the total capital stock in county j at time t , such that

$$k_t^i = \kappa_{i|j} K_t^j. \quad (33)$$

For each census tract i in county j and year t , I define effective capital as

$$\tilde{k}_t^i = k_t^i \omega_i^{SVI} \omega_i^{CRI}, \quad (34)$$

where k_t^i denotes the tract-level physical capital stock. The wedges $\omega_i^{SVI} \in (0, 1]$ and $\omega_i^{CRI} \in (0, 1]$ capture systematic differences in the productivity and durability of capital across locations.

The underlying premise is that identical physical capital stocks may generate different effective services depending on local structural conditions. In areas with greater social vulnerability and/or lesser community resilience, capital is more exposed to damage, less protected by mitigation investments, and embedded in environments with weaker financial and institutional buffers. As a result, one unit of physical capital delivers fewer effective services and is more susceptible to persistent losses. Effective capital, therefore, scales physical capital by tract-specific characteristics that affect both exposure and recovery capacity. In practice, this adjustment concentrates a larger share of disaster-induced depreciation toward socially vulnerable and low-resilience areas, allowing the tract-level construction to better capture the localized nature of disaster losses.

Social vulnerability is measured using the FEMA Social Vulnerability Index (SVI), reported as a percentile rank in $[0, 100]$. Let $sx^i = \text{SVI}^i/100$. I map vulnerability into a multiplicative discount factor according to

$$\omega_i^{SVI} = \exp(-\theta_S sx^i), \quad (35)$$

where $\theta_S > 0$ governs the sensitivity of effective capital to social vulnerability. The exponential specification ensures $\omega_i^{SVI} \in (0, 1]$, preserves strict positivity, and implies convex losses in effective capital as vulnerability increases.

Community resilience is measured by the FEMA Community Resilience Index (CRI), also reported as a percentile rank. Let $cx^i = \text{CRI}^i/100$ and define the resilience deficit $cd^i = 1 - cx^i$. Its effect is given by

$$\omega_i^{CRI} = \exp(-\theta_C cd^i), \quad (36)$$

where $\theta_C > 0$ governs the sensitivity to low resilience.

I impose $\theta_C < \theta_S$. This restriction reflects the interpretation that vulnerability captures structural exposure and susceptibility to damage, the primary determinant of effective capital at risk, whereas resilience modifies mitigation capacity and recovery dynamics. In this sense, social vulnerability governs the first-order adjustment to capital services, while resilience acts as a secondary amplifier or dampener.

This formulation interprets SVI and CRI as affecting the effective availability and absorptive capacity of local capital services rather than mechanically altering the physical stock.

B.6.5. *Attribution of County-Level Losses to Census Tracts.* Let NDL_t^j denote total realized disaster losses in county j and year t . I allocate these losses across hit census tracts in two stages to generate a tract-level damage distribution characterized by rare but severe local destruction and smaller impacts in other affected areas.

Stage 1: Extreme Events and Single-Hit Counties. If aggregate county losses exceed total county effective capital,

$$NDL_t^j \geq \tilde{K}_t^j \equiv \sum_{i \in j} \tilde{k}_t^i,$$

all H hit tracts are assigned a depreciation rate of one.

If exactly one census tract is hit in county j during year t ($H = 1$), all realized losses are assigned to that tract, up to its effective capital:

$$ndl_t^i = \min\{NDL_t^j, \tilde{k}_t^i\}.$$

Stage 2: Multi-Tract Allocation. For county-years with $H > 1$ and $NDL_t^j < K_t^j$, I first identify an *epicenter* tract. The epicenter is defined as the tract with the largest number of disaster records. If multiple tracts share the same maximum number of hits, the tie is broken by selecting the tract with the highest expected historical depreciation rate, defined as FEMA's expected building loss divided by 2020 building value.

Conditional on the epicenter selection, losses are allocated in three steps. First, the epicenter receives a baseline allocation equal to

$$\min\{0.2\tilde{k}_t^i, NDL_t^j(1 - \varepsilon H)\},$$

where ε , set at 0.1%, is a small constant that guarantees strictly positive losses for all hit tracts.

Second, each of the remaining $H - 1$ tracts receives a minimal share equal to εNDL_t^j .

Third, the remaining losses are distributed as follows: a fixed fraction (30%) is assigned to the epicenter, and the residual share is allocated across the remaining tracts proportionally to their historical expected depreciation rates. All allocations are capped at tract-level effective capital. Any excess losses above $\tilde{k}_t^i$ are reallocated to the epicenter, subject again to its capital constraint.

Finally, the tract-level depreciation rate is defined as

$$d_t^i = \frac{ndl_t^i}{\tilde{k}_t^i}, \tag{37}$$

where ndl_t^i denotes the allocated loss in tract i during year t .

APPENDIX C. ESTIMATION PROCEDURE

This appendix presents the full estimation procedure for the latent regime-switching depreciation process. The main text presents the model in reduced form for clarity of exposition. Here, I derive the full likelihood, Hamilton filter, and panel representation used for estimation.

C.1. Within-Zone Disaster Regimes. The estimation is conducted separately within each climate zone $z \in \mathcal{Z}$. Conditional on climate zone z , let $d_t^i \geq 0$ denote the observed depreciation rate for tract i in year t .

Within each climate zone, households transition across latent disaster-event realizations

$$u_t^i \in \mathcal{U} = \{1, \dots, U\},$$

following a zone-specific Markov process consistent with the theoretical model.

As in the model section, the disaster state is represented by the selection vector

$$v_t^i \in \left\{ e_1^{(U)}, \dots, e_U^{(U)} \right\},$$

where

$$v_t^i = e_u^{(U)} \Leftrightarrow u_t^i = u.$$

Conditional on the latent disaster regime, positive depreciation rates satisfy

$$\log \left(\frac{d_t^i}{1 - d_t^i} \right) = (\mu^z)^\top v_t^i + (\sigma^z)^\top v_t^i \varepsilon_{i,t}, \quad \varepsilon_{i,t} \sim \mathcal{N}(0, 1), \quad (38)$$

where

$$\mu^z = (\mu_1^z, \dots, \mu_U^z)^\top$$

collects the regime-specific depreciation intensities and

$$\sigma^z = (\sigma_1^z, \dots, \sigma_U^z)^\top$$

the regime-specific standard deviations.

Equivalently,

$$\log \left(\frac{d_t^i}{1 - d_t^i} \right) \mid v_t^i \sim \mathcal{N} \left((\mu^z)^\top v_t^i, ((\sigma^z)^\top v_t^i)^2 \right). \quad (39)$$

C.2. Zero-Inflated Mixture. A large fraction of observations exhibits zero depreciation, corresponding to years without realized disaster damage. To account for this feature, I augment the model with a zero-inflated mixture structure.

Conditional on latent disaster regime $u_t^i = u$ within climate zone z , depreciation equals zero with probability $1 - p_u^z$ and is strictly positive with probability p_u^z . The conditional density is therefore

$$\begin{aligned} f(d_t^i \mid v_t^i, z) &= (1 - (p^z)^\top v_t^i) \mathbb{1}_{\{d_t^i=0\}} \\ &\quad + ((p^z)^\top v_t^i) \mathbb{1}_{\{d_t^i>0\}} \frac{1}{d_t^i(1 - d_t^i)} \phi \left(\log \left(\frac{d_t^i}{1 - d_t^i} \right); (\mu^z)^\top v_t^i, (\sigma^z)^\top v_t^i \right), \end{aligned} \quad (40)$$

where $\phi(\cdot; \mu, \sigma)$ denotes the Gaussian density with mean μ and standard deviation σ .

C.3. Latent Markov Dynamics. Within each climate zone $z \in \mathcal{Z}$, persistence in disaster exposure is governed by a latent first-order Markov chain over disaster-event realizations $u \in \mathcal{U}$. The transition dynamics are summarized by the climate zone-specific transition matrix

$$\Pi^z = (\pi_{uu'}^z)_{u,u'=1}^U,$$

where

$$\pi_{uu'}^z = \mathbb{P}(u_{t+1}^i = u' \mid u_t^i = u, z).$$

Equivalently, in terms of the selection-vector representation,

$$\pi_{uu'}^z = \mathbb{P}(v_{t+1}^i = e_{u'}^{(U)} \mid v_t^i = e_u^{(U)}, z).$$

C.4. Hamilton Filter. Following [Hamilton \(1994\)](#), define the filtered probabilities

$$\hat{v}_{i,t|t}^z = (\mathbb{P}(u_t^i = 1 \mid \underline{d}_{i,t}, z), \dots, \mathbb{P}(u_t^i = U \mid \underline{d}_{i,t}, z))', \quad (41)$$

where

$$\underline{d}_{i,t} = (d_t^i, d_{t-1}^i, \dots, d_1^i)'. \quad (42)$$

Let

$$\eta_{i,t}^z = (f(d_t^i \mid u_t^i = 1, z), \dots, f(d_t^i \mid u_t^i = U, z))' \quad (42)$$

denote the vector of regime-specific conditional densities.

The Hamilton filter recursion is

$$\hat{v}_{i,t|t}^z = \frac{\hat{v}_{i,t|t-1}^z \odot \eta_{i,t}^z}{(\mathbf{1}_U)' (\hat{v}_{i,t|t-1}^z \odot \eta_{i,t}^z)}, \quad (43)$$

with prediction step

$$\hat{v}_{i,t+1|t}^z = (\Pi^z)' \hat{v}_{i,t|t}^z. \quad (44)$$

Here, $\odot$ denotes the Hadamard product and $\mathbf{1}_U$ is a $U \times 1$ vector of ones.

C.5. Panel Representation. Within each climate zone, observations are assumed to be independent and identically distributed conditional on latent disaster states.

Define the stacked vectors

$$D_t^z = (d_{1,t}^z, \dots, d_{N,t}^z)',$$

and

$$\Upsilon_t^z = ((v_{1,t}^z)', \dots, (v_{N,t}^z)')'.$$

Conditional on latent disaster states, the joint density factorizes across tracts:

$$f(D_t^z \mid \Upsilon_t^z) = \prod_{i=1}^N f(d_t^i \mid v_t^i, z). \quad (45)$$

Equivalently,

$$\log \left(\frac{D_t^z}{1 - D_t^z} \right) \Bigg| \Upsilon_t^z \sim \mathcal{N}_N (\Upsilon_t^z \mu^z, \Sigma_t^z), \quad (46)$$

where

$$\Sigma_t^z = \text{diag} ((\Upsilon_t^z \sigma^z)^2).$$

Define

$$H_t^z = (\eta_{1,t}^z, \dots, \eta_{N,t}^z)'.$$

The stacked filtering recursion becomes

$$\hat{\Upsilon}_{t|t}^z = (\hat{\Upsilon}_{t|t-1}^z \odot H_t^z) \oslash \left[(\hat{\Upsilon}_{t|t-1}^z \odot H_t^z) \mathbf{1}_U \right], \quad (47)$$

with prediction equation

$$\hat{\Upsilon}_{t+1|t}^z = \hat{\Upsilon}_{t|t}^z (\Pi^z)'. \quad (48)$$

C.6. Likelihood and Numerical Estimation. The log-likelihood function for climate zone z is³⁶

$$\mathcal{L}^z(\theta^z) = \sum_t (\mathbf{1}_N)' \log \left[(\hat{\Upsilon}_{t|t-1}^z \odot H_t^z) \mathbf{1}_U \right]. \quad (49)$$

The parameter vector is

$$\theta^z = (\mu^z, \sigma^z, \{\pi_{uu'}^z\}_{u,u'=1}^{U-1}, \{p_u^z\}_{u=1}^U).$$

Parameters are estimated independently within each climate zone by numerical maximum likelihood using the **Nelder-Mead** and **nlminb** optimization routines.

³⁶Note that the log is taken element-wise.