THE CLIMATE PROTECTION GAP AND PRIVATE SELF-INSURANCE

PAULINE CHIKHANI

Abstract. This paper studies how localized climate tail risk affects household wealth accumulation when the primary saving instrument is itself exposed to stochastic destruction. I develop a heterogeneous-agent model in which disaster risk takes the form of idiosyncratic, wealth-proportional depreciation shocks, so that wealth simultaneously serves as a buffer against shocks and the object exposed to them. This interaction distorts precautionary saving and generates highly nonlinear responses. At low wealth levels, households in high-risk locations optimally accumulate less and cluster at the borrowing constraint relative to otherwise identical households in moderate-risk areas. Quantitatively, the aggregate effects of empirically calibrated disaster risk are modest: aggregate capital declines by less than 1 percent. In contrast, distributional effects are large. Introducing disaster-induced depreciation increases the share of hand-to-mouth households by 25 percentage points and substantially compresses the lower tail of the wealth distribution. Climate tail risk, therefore, operates primarily through a distributional channel rather than through large aggregate contractions.

JEL: Q54, D31.

Keywords: climate risk, heterogeneous agents, insurance.

This version: February 2026.

University of Lausanne (HEC), Department of Economics, Lausanne, CH-1015, Switzerland

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

The author declares that she has received no external funding for this research and has no financial relationships or conflicts of interest to disclose.

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 participants at the University of Lausanne for valuable feedback. All remaining errors are my own.

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.

The reinsurance industry reports that a substantial share of losses remains uninsured and therefore falls on households (Swiss Re Institute, 2025). The fraction of disaster-related economic losses not covered by insurance is referred to as the "climate protection gap". In the United States, it is estimated at roughly 43%, and has widened over recent decades. The relevance of this climate gap may increase in the future, as climate projections point to increases in both the frequency and severity of extreme weather events. These developments raise a central question: how does disaster risk affect household wealth accumulation when the primary saving instrument is itself subject to destruction?

This paper studies how households self-insure against localized climate risk through precautionary saving. I develop a heterogeneous-agent model with incomplete markets and borrowing constraints in which all accumulated assets are exposed to stochastic capital-destruction shocks. Climate risk is modeled as location-specific depreciation shocks that directly reduce household wealth.

The key feature of the environment is that the risk of destruction scales with accumulated assets: wealth is both a buffer against losses and the object exposed to them. In this environment, saving is shaped by two opposing forces. On the one hand, households accumulate assets to self-insure against future destruction shocks: wealth provides a buffer that helps smooth consumption when losses occur. On the other hand, the asset used for saving is itself subject to destruction, so future resources depend on the

realization of shocks. As a result, accumulating assets both increases the ability to absorb losses and exposes households to the risk that part of the resources saved for future consumption may be destroyed.

To discipline the stochastic destruction process, 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. Therefore, I allocate observed county-year losses across census tracts using tract-level exposure shares to recover the distribution of local tail risk. The resulting data reveal pronounced spatial heterogeneity in both disaster frequency and upper-tail severity. I summarize this heterogeneity with two climate-risk zones (e.g., coastal and inland), each characterized by a distinct, discrete distribution of destruction shocks. These estimated processes are then embedded in the quantitative model. This framework allows me to quantify the extent to which localized climate tail risk affects household saving behavior and the distribution of wealth.

The model delivers three main results. First, precautionary saving responses to climate risk are highly nonlinear in wealth because wealth simultaneously serves as a buffer against shocks while being exposed to destruction risk. The balance between these forces depends on wealth levels. When wealth is low, a natural disaster in the following period may destroy a large share of the household's savings. In that case, the marginal utility of consuming today can exceed the expected marginal utility of saving, weakening saving incentives.

At higher levels of wealth, however, the standard precautionary saving mechanism dominates. Additional assets provide a meaningful buffer against destruction shocks, encouraging households to accumulate precautionary savings.

As climate risk increases, this divergence becomes more pronounced: households in high-risk zones save less at low wealth levels but more at intermediate wealth levels. Quantitatively, saving in high-risk zones can exceed that in safer areas by up to a factor of four 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 at the borrowing constraint rises sharply, from less than 0.5% to almost 30%.

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.

The analysis highlights the role of asset composition in economies exposed to rising climate risk. In the model, all wealth is subject to destruction shocks, and no safe asset is available. The absence of a safe asset weakens the ability of precautionary saving as a form of self-insurance and increases reliance on ex post insurance or public transfers following disasters. More broadly, the results suggest that financial innovation and public policies that shield household balance sheets from climate-related capital destruction may play a central role in limiting the distributional consequences of increasing climate risk.

This paper contributes to three strands of the literature: heterogeneous-agent models of climate risk, rare-disaster macroeconomics, and stochastic-return theories of wealth inequality. Its central contribution is to introduce empirically disciplined, spatially heterogeneous destruction risk into an incomplete-markets framework, demonstrating that partially uninsured climate risk operates primarily through a rather than through large-scale aggregate contractions.

Heterogeneous-Agent Climate Models. A growing literature studies the macroeconomic consequences of climate risk in quantitative heterogeneous-agent and spatial environments. Fried (2022) embeds storm risk in an Aiyagari model with regional heterogeneity and federal disaster aid, emphasizing adaptation incentives and redistribution across risk zones. Fried (2026) extends this framework to rising temperature exposure and heterogeneous income dynamics. Bilal and Rossi-Hansberg (2023) quantify local capital destruction and analyze migration and spatial reallocation in response to anticipated climate change, while Indaco and Ortega (2024) study local population adjustment under climate risk. These contributions emphasize migration, spatial equilibrium, adaptation, and policy responses as primary adjustment margins. By contrast, I abstract

from mobility and isolate a different mechanism: household self-insurance in incomplete markets when the saving instrument itself is exposed to destruction risk.

A closely related body of literature investigates the asset-pricing implications of incomplete disaster insurance coverage under heterogeneous climate-risk beliefs (Bakkensen and Barrage, 2022; Barrage and Furst, 2019; Gallagher, 2014; Woongchan, Barrage, and Walsh, 2025). Rather than modeling insurance demand directly, I investigate the limits of self-insurance of households: when savings are themselves exposed to destruction risk, and insurance is incomplete, private self-insurance becomes inherently limited.

Rare Disasters and Stochastic Depreciation. The paper relates to the rare-disaster literature pioneered by Barro (2009), Gourio (2012), and Gabaix (2012), in which low-probability, high-severity aggregate shocks shape asset prices and macroeconomic dynamics. Extensions incorporate heterogeneous beliefs and production heterogeneity to explain risk premia and valuation dynamics (Chen, Joslin, and Tran, 2012; Gomes, Grotteria, and Wachter, 2019; Miller, Paron, and Wachter, 2025). In these models, disasters are aggregate events affecting the entire economy simultaneously.

In contrast, I study localized and spatially heterogeneous capital-destruction shocks that leave aggregate productivity largely intact but alter the cross-sectional distribution of wealth. Quantitatively, aggregate capital declines modestly, yet the mass of households at the borrowing constraint rises sharply, and wealth inequality increases substantially. The contribution is to show that localized tail risk can generate first-order distributional effects even when aggregate effects are small.

Stochastic Returns and Wealth Inequality. A large literature studies multiplicative return risk and heterogeneous returns in incomplete-markets models (Angeletos, 2007; Benhabib, Bisin, and Luo, 2019; Benhabib, Bisin, and Zhu, 2011,1; Moll, Rachel, and Restrepo, 2022). In these frameworks, persistent heterogeneity in investment opportunities generates fat-tailed wealth distributions and shapes inequality dynamics.

I use this framework to study climate-induced capital destruction. The shocks are rare, highly skewed, and spatially heterogeneous rather than persistent return advantages. Because destruction scales with wealth, exposure increases mechanically with accumulation. This generates a novel distortion: precautionary saving can weaken at low wealth levels in high-risk regions, and wealth may cease to function as an effective buffer. In this sense, climate risk compresses the lower tail of the wealth distribution rather than fattening the upper tail.

2. Model

2.1. **Households.**

The household sector consists of a unit mass of infinitely-lived households, indexed by $i \in [0,1]$.[1] They are assumed to have identical CRRA preferences defined over expected lifetime utility from future consumption $\{c_t^i\}_{t\geq 0}$:

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

with $\beta = \frac{1}{1+\rho}$ and $\gamma \geq 1$.

Households face idiosyncratic capital-destruction risk, which generates heterogeneity in wealth holdings a_t^i and effective depreciation rates δ_t^i. Each period, they allocate resources between consumption and savings. Income is earned by supplying labor and capital to a representative firm. Households supply one unit of labor inelastically ($L_t = 1$, $\forall t$) and receive the wage w_t. They also accumulate savings in the form of physical capital, rented to the firm at the gross return r_t^K.

The wage rate w_t and capital return r_t^K are determined in equilibrium and are identical across households. In contrast, capital holdings are subject to idiosyncratic depreciation shocks reflecting household-specific disaster exposure. The resulting net return on capital for household i at time t is

$$r_t^i = r_t^K - \delta_t^i, \tag{2}$$

where the effective depreciation rate δ_t^i consists of a constant component $\delta > 0$, capturing standard physical depreciation, and an idiosyncratic stochastic process $s_t^i \geq 0$, capturing additional disaster-related capital destruction:

$$\delta_t^i = \delta + s_t^i. \tag{3}$$

The process s_t^i evolves according to a J-state Markov chain with state space $\mathcal{J} = \{1, \ldots, J\}$. Specifically,

$$s_t^i = \mu^\top x_t^i, \tag{4}$$

where $x_t^i \in \left\{e_1^{(J)}, \ldots, e_J^{(J)}\right\}$ is a selection vector corresponding to a column of the identity matrix $I_J \in \mathbb{R}^{J\times J}$. The vector $\mu = (\mu_1, \ldots, \mu_J)^\top \in \mathbb{R}^J$ collects the possible realizations

[1] In 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.

of s_t^i, with μ_j denoting the value of the shock in state $j \in \mathcal{J}$. Thus,

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

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

$$\mathbb{P}\left(s_{t+1}^i = \mu_{j'} \mid s_t^i = \mu_j\right) = \pi_{jj'}, \tag{5}$$

implying that each row of Π sums to one.

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

$$c_t^i + a_{t+1}^i = \left(1 + r_t^K - \left(\delta + s_t^i\right)\right) a_t^i + w_t, \tag{6}$$

together with a borrowing constraint:

$$a_{t+1}^i \geq -\phi. \tag{7}$$

The household's recursive problem is characterized by the value function $V(a, s)$, which satisfies the Bellman equation, subject to the budget (6) and borrowing (7) constraints:

$$\begin{aligned} V(a,s) &= \max_{a'} \frac{c^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}\left[V\left(a', s'\right) | s\right], \\ \text{s.t.} \quad & c + a' = \left(1 + r^K - (\delta + s)\right) a + w, \\ & a' \geq -\phi, \end{aligned} \tag{8}$$

where $\mathbb{E}\left[V\left(a', s'\right) | s\right] = \sum_{s'=1}^{J} \pi(ss') V\left(a', s'\right)$. The primes denote the value of the variable in the next period $t+1$, and s is the current realization of the idiosyncratic depreciation shock.

2.2. **Stationary Equilibrium Fix.** When capital destruction is proportional to individual wealth and sufficiently severe, the model may fail to generate a stable stationary wealth distribution. Since the only saving instrument is itself exposed to idiosyncratic destruction risk, precautionary motives may be weakened when destruction risk is large. In such cases, optimal asset accumulation can collapse toward the borrowing constraint, leading to a degenerate invariant distribution.

To ensure that the quantitative model delivers a non-degenerate stationary wealth distribution, I introduce a small idiosyncratic labor productivity component. Labor

income becomes

$$y_t^i = w_t \, \xi(x_t^i),$$

where $\xi(x_t^i) > 0$ is a state-dependent productivity shifter driven by the same Markov state as capital destruction. The productivity component is normalized to have stationary mean one, so that aggregate labor income remains equal to w_t in equilibrium (i.e., $\mathbb{E}(y_t^i) = w_t$). This modification introduces a precautionary savings motive without affecting aggregate quantities.[2]

2.3. **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 $\mathcal{Z} = \{1, \ldots, Z\}$, 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.[3] The climate zone assignment, therefore, captures persistent geographic exposure.

Climate zones are ordered such that both the frequency and severity of shocks increase with z, so that

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

This ordering reflects more serious average damage in more exposed regions, while remaining agnostic about differences in higher moments across zones.

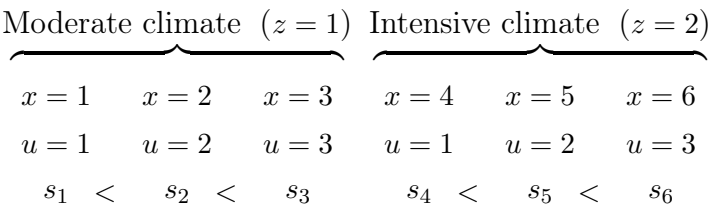
Within each climate zone $z \in \mathcal{Z}$, households face short-run weather realizations indexed by $u \in \mathcal{U} = \{1, \ldots, U\}$. Conditional on climate zone z, the regime evolves according to a climate zone-specific U-state Markov chain with transition matrix $\Pi^z = (\pi^z_{uu'})_{u,u' \in \mathcal{U}}$.

Let $\mu^z = (\mu^z_1, \ldots, \mu^z_U)^\top$ denote the vector of destruction realizations associated with regimes in zone z, where μ^z_u is the depreciation shock in regime u of climate zone z. The full vector of possible destruction realizations is obtained by concatenating the zone-specific vectors:

$$\mu = \left((\mu^1)^\top, \ldots, (\mu^Z)^\top\right)^\top \in \mathbb{R}^J, \qquad J = Z \times U.$$

[2]Details of the normalization are provided in Appendix A, Section A.2.

[3]One 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. I abstract from such climate-zone transitions in the baseline specification. For the fully general case, see Appendix A.3.



Notes: Organization of the $J = 6$ regimes used in the model. Each regime is indexed by $x \in \{1, \ldots, 6\}$ 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 weather state within each climate type. These weather states determine the damage intensity, with $u = 1$ indicating low damages, $u = 2$ medium damages, and $u = 3$ high damages. For any given zone z, the unconditional mean of damages are higher in the intensive climate ($z = 2$) than in the moderate climate ($z = 1$). Formally, $\mathbb{E}[s_t^i \mid z = 2] > \mathbb{E}[s_t^i \mid z = 1]$.

FIGURE 2.1. Regimes Organization

The complete climate-risk state for household i at time t is given by the pair $(z, u) \in \mathcal{Z} \times \mathcal{U}$. Each pair corresponds to a unique index $j \in \mathcal{J} = \{1, \ldots, J\}$.

As in the general formulation, the joint state is encoded by a selection vector $x_t^i \in \{e_1^{(J)}, \ldots, e_J^{(J)}\}$, which can be written as the Kronecker product

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

where $\zeta_t^i \in \{e_1^{(Z)}, \ldots, e_Z^{(Z)}\}$ denotes the zone selection vector and $\upsilon_t^i \in \{e_1^{(U)}, \ldots, e_U^{(U)}\}$ denotes the within-zone regime selection vector.[4] The destruction realization is still

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

so that each pair (z, u) corresponds to $s_t^i = \mu_u^z$.

In the following analysis, I consider the case $Z = 2$ climate zones and $U = 3$ weather regimes 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 $x \in \{1, \ldots, 6\}$. Figure 2.1 illustrates the organization of these regimes.

Heterogeneity therefore arises along two dimensions: across zones through differences in long-run exposure (persistent climate risk), and within zones through stochastic regime realizations (transitory weather risk). Importantly, the household's recursive problem is invariant to whether climate risk is represented directly as a J-state Markov chain or

[4]For example, when $Z = 2$ and $U = 3$, a household in zone 1 and regime 2 is represented by $\zeta_t^i = (1, 0)^\top$, $\upsilon_t^i = (0, 1, 0)^\top$, 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

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

.

decomposed into Z zones of U regimes. In the structured representation (Figure 2.1), the state variable remains the realized destruction level $s = \mu_u^z$. Conditional on zone z, households transition only across the zone-specific realizations collected in $\mu^z = (\mu_1^z, \ldots, \mu_U^z)$ according to Π^z. Thus, the optimization problem is unchanged; the zone–regime structure simply reorganizes the J states into economically interpretable blocks.

2.4. **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 on competitive input markets, at prices w and r^K respectively. The capital stock is endogenously determined by the aggregate capital holdings of households, such that $K_t = \int_0^1 a_t^i \, di$.

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

$$\max_{K_t, L_t} Y_t - r_t^K K_t - w_t L_t, \tag{9}$$

$$\text{s.t.} \quad Y_t = AK_t^\alpha L_t^{1-\alpha}. \tag{10}$$

The first-order conditions give the firm's demand for capital and labor:[5]

$$r_t^K = \alpha A K_t^{\alpha-1}, \tag{11}$$

$$w_t = (1-\alpha) A K_t^\alpha. \tag{12}$$

Note that prices are common across zones due to integrated capital and labor markets.

2.5. **The Recursive Competitive Equilibrium.** Households are permanently assigned to a climate zone $z \in \mathcal{Z}$. Let zone shares be $\{\lambda_z\}_{z=1}^Z$, with $\sum_{z=1}^Z \lambda_z = 1$. Within each climate zone, the idiosyncratic destruction regime $u \in \mathcal{U} = \{1, \ldots, U\}$ follows a zone-specific transition matrix $\Pi^z = (\pi_{uu'}^z)_{u,u' \in \mathcal{U}}$.

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:

[5] Given the normalization $L_t = 1$.

(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 (8), where expectations are taken with respect to the zone-specific transition matrix Π^z.[6]

(ii) **Firm.** Given prices, the representative firm behaves optimally and factor prices satisfy (11)–(12).

(iii) **Market clearing.** Given the distribution $m^z(a, u)$, markets clear:

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

(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, $m^z(a, s)$ is the stationary distribution.

$$m^z(a', s') = \sum_{s \in \mathcal{U}} \int_{a:\, a' = g_z(a,u)} m^z(da, u)\, \pi^z_{uu'}.$$

3. Data

This section constructs county- and census-tract-level measures of natural-disaster-induced capital depreciation. I first build county-level depreciation rates using SHELDUS losses and BEA capital stocks. I then explain why county averages mechanically understate localized destruction, motivating a disaggregation of losses to census tracts using FEMA's National Risk Index (NRI). The final dataset is a tract–year panel used to estimate the stochastic depreciation processes in the model.

3.1. Construction of County-Level Depreciation Rates.

To estimate the stochastic depreciation processes s, I construct a county–year panel dataset that captures capital losses from natural disasters. The empirical counterpart to the model's depreciation shock is the loss rate

$$d^i_t \equiv \frac{\text{NDL}^i_t}{K^i_{t-1}}, \tag{13}$$

where NDL^i_t denotes direct economic losses from natural hazards in county i and year t, and K^i_{t-1} is the corresponding county-level capital stock at the beginning of the period.

Capital Stocks. The Bureau of Economic Analysis (BEA) provides annual national estimates of the current-cost net stocks of fixed assets and consumer durable goods,

[6] A full description of the solution method used to numerically solve the model can be find in Appendix A.1.

denoted K_t^N.[7] In the absence of subnational capital stock data, I allocate national capital proportionally to county GDP shares. Under a Cobb–Douglas technology with common factor shares, this corresponds to the assumption of constant capital–output ratios across counties. Let GDP_t^N and GDP_t^i denote national and county GDP, respectively. I set[8]

$$K_t^i = \frac{GDP_t^i}{GDP_t^N} K_t^N. \tag{14}$$

Disaster Losses. County-level disaster losses are drawn from the Spatial Hazard Events and Losses Database for the United States (SHELDUS; CEMHS). SHELDUS reports county-level events from 1960–2022 and records direct economic losses disaggregated into property and crop components. I use property losses as the baseline measure of capital destruction and exclude crop losses because cultivated assets are not capitalized in BEA fixed-asset accounts.[9] The sample window is 2002–2022 to align disaster losses with county-level macroeconomic controls. Dollar values are in current dollars.

SHELDUS distinguishes 18 hazard types. For parsimony, I group these into eight categories following IRDR (2014). The mapping is reported in Appendix B.2, Table B.1.

Geographic Harmonization. Counties are matched across datasets using Federal Information Processing Standards (FIPS) codes. Because BEA local-area statistics aggregate some counties under composite codes, I apply the same aggregation to the SHELDUS loss data to ensure consistency. I also harmonize counties affected by boundary changes over the sample period. Appendix B.3 and Table B.2 document all adjustments.

3.2. **Limits of County-Level Depreciation Rates for Measuring Household Shocks.** County-level depreciation rates capture geographic exposure but cannot identify the intensity of capital destruction faced by households. Because disaster damages are observed at the county level while disasters typically affect limited areas within counties, aggregation mechanically smooths measured depreciation rates and masks within-county heterogeneity.

Table 3.1 reports the mean and selected percentiles of county-level capital depreciation rates conditional on positive reported damages. Two features stand out. First,

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

[8]It is worth noting that this assumption underestimates losses in low-intensive-capital counties and overestimates losses in high-intensive-capital counties.

[9]Soloveichi (2021) provides estimates of cultivated assets up to 2019. Appendix B.4 reports robustness checks incorporating crop losses.

average depreciation rates from natural disasters are quantitatively small: the mean county-level depreciation rate is 0.14%, and even at the 95th percentile, losses amount to only 0.23% of the county's beginning-of-period capital stock. This reflects the fact that natural disasters typically affect a limited geographic area within a county, whereas the denominator includes the entire county-level capital stock. As a result, aggregation substantially smooths measured depreciation rates, masking the intensity of localized destruction and the distribution of household-level wealth impacts.

TABLE 3.1. Distribution of County-Level Capital Depreciation Rates

Obs.>0	Mean (%)	Median (%)	P90 (%)	P95 (%)	P99 (%)
47417	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 mean, the median, and percentiles (P90, P95, and P99) are computed over all county-year observations with positive reported damages. All values are expressed in percentage points.

Second, the distribution is strongly right-skewed: rare extreme events disproportionately influence the mean. Consistent with this interpretation, the mean exceeds the 90th percentile, indicating that a small number of very large depreciation realizations drive the average.

In the county-level panel, without differentiating by hazard, 73% of county-year observations exhibit positive reported disaster losses. Under county-level measurement, all capital, hence the entire population, within those counties is treated as exposed to disaster-related depreciation. This measure, therefore, captures geographic exposure at the county level rather than direct household-level impacts.

Survey evidence illustrates this gap. Using the Federal Reserve's Survey of Household Economics and Decisionmaking (SHED; Board of Governors of the Federal Reserve System (2024)), between 10–20% of respondents report any financial impact from natural disasters in the previous year, and roughly 2% report substantial impacts (see Table B.3 in Appendix B.5). The contrast between widespread county-level exposure and concentrated household-level impacts motivates a disaggregation of county-level losses to census tracts while preserving national coverage.

3.3. **Construction of Census-Tract-Level Depreciation Rates.** Census tracts are relatively permanent subdivisions of counties. I adopt 2020 Census tract boundaries and

treat them as time-invariant over 2002–2022.[10] Tracts contain, on average, roughly 4,000 residents.[11]

Within-County Exposure Shares. To discipline within-county heterogeneity, I use FEMA's National Risk Index (NRI) tract-level measures of expected annual losses (EAL), social vulnerability (SVI), and community resilience (CRI) (Federal Emergency Management Agency, 2025). For tract i located in county j, I define exposure shares based on expected annual losses:

$$w_{i|j} = \frac{EAL_i}{\sum_{k \in j} EAL_k}, \qquad \sum_{i \in j} w_{i|j} = 1. \tag{15}$$

Let $Records_t^j$ denote the realized number of disaster events in county j during year t. Conditional on $Records_t^j$, events are independently allocated across census tracts according to these exposure shares:

$$(Records_t^i)_{i \in j} \mid Records_t^j \sim \text{Multinomial}\big(Records_t^j, \{w_{i|j}\}_{i \in j}\big)\,. \tag{16}$$

A tract is classified as "hit" if it experiences at least one event.[12]

Figure 3.1 illustrates how the tract-level hit probabilities implied by the allocation procedure map into FEMA's baseline risk classifications. The figure plots the conditional composition of census tracts by NRI rating across the distribution of hit probabilities. Tracts with low hit probabilities are predominantly classified as "Very Low" or "Relatively Low" risk, whereas tracts with higher hit probabilities are increasingly represented by "Relatively High" and "Very High" risk categories (see Section 4.1 for the risk classification). This monotonic shift reflects the role of tract-level expected annual losses in distributing county-level disaster intensity across space: tracts with higher baseline exposure receive a larger share of disaster intensity and therefore exhibit a higher probability of being hit. The figure thus provides a transparent interpretation of how ex ante risk measures translate into heterogeneous local exposure while preserving consistency with observed county-level disaster activity.

[10] Census tract shapefiles are obtained from U.S. Census Bureau (2020) using the `tigris` package in `R`.
[11] Details on tracts and the construction of the depreciation rates at the tract level can be found in Appendix B.6.
[12] Appendix B.6 provides the equivalent Poisson representation and formal derivations.

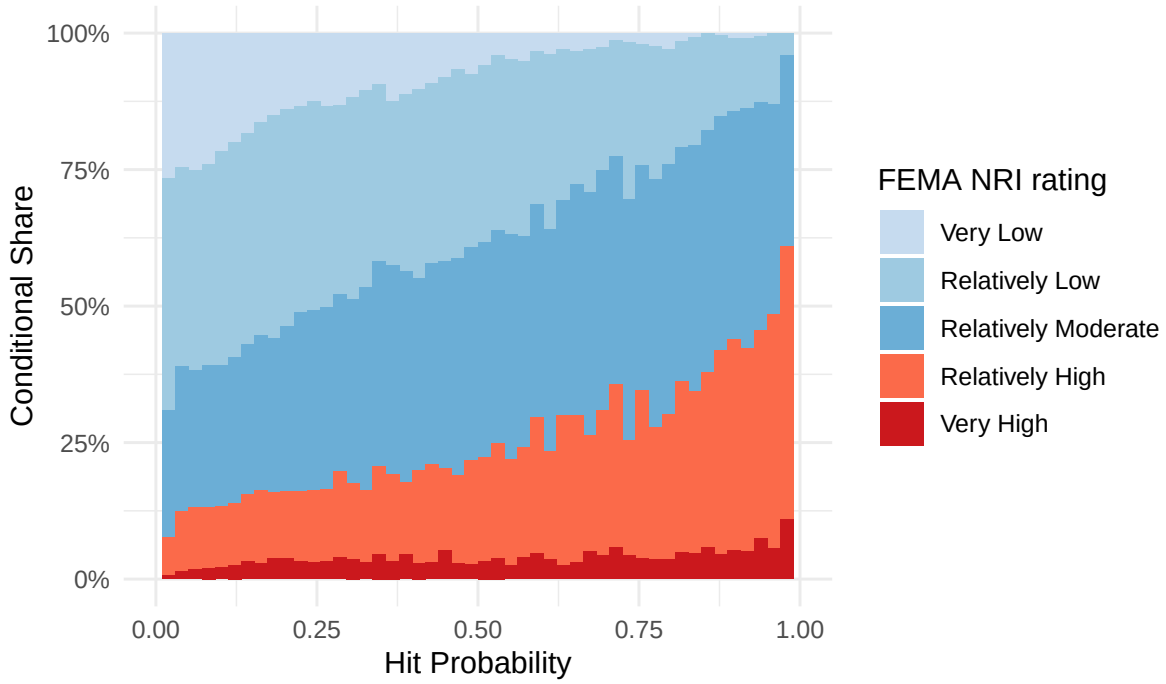


FIGURE 3.1. Census Tract Hit Probability by FEMA National Risk Index Rating

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 event counts to tracts in proportion to each tract's expected annual losses (EAL). Appendix B.6 provides the equivalent Poisson representation and derivations. For each value of the hit probability on the horizontal axis, the stacked areas show the share of tracts in each NRI rating category.

Capital and Effective Capital. Tract-level physical capital is obtained by allocating county capital using building replacement value shares:

$$\kappa_{i|j} = \frac{BV_i^{2020}}{\sum_{k \in j} BV_k^{2020}}, \qquad k_t^i = \kappa_{i|j} K_t^j. \tag{17}$$

Effective capital incorporates tract-level vulnerability and resilience wedges:

$$\tilde{k}_t^i = k_t^i \, \omega_i^{SVI} \, \omega_i^{CRI}, \tag{18}$$

where

$$\omega_i^{SVI} = \exp(-\theta_S s x^i), \qquad \omega_i^{CRI} = \exp(-\theta_C c d^i). \tag{19}$$

I impose $\theta_C < \theta_S$, so that vulnerability governs the first-order discount in effective capital while resilience acts as a secondary modifier.

The implications of these wedges for the cross-sectional distribution of capital are illustrated in Figure 3.2, which compares physical capital k_t^i and effective capital $\tilde{k}_t^i$ at the tract-year level. The wedges shift the distribution leftward and increase dispersion, reflecting lower and more heterogeneous capital services once local structural conditions are accounted for.

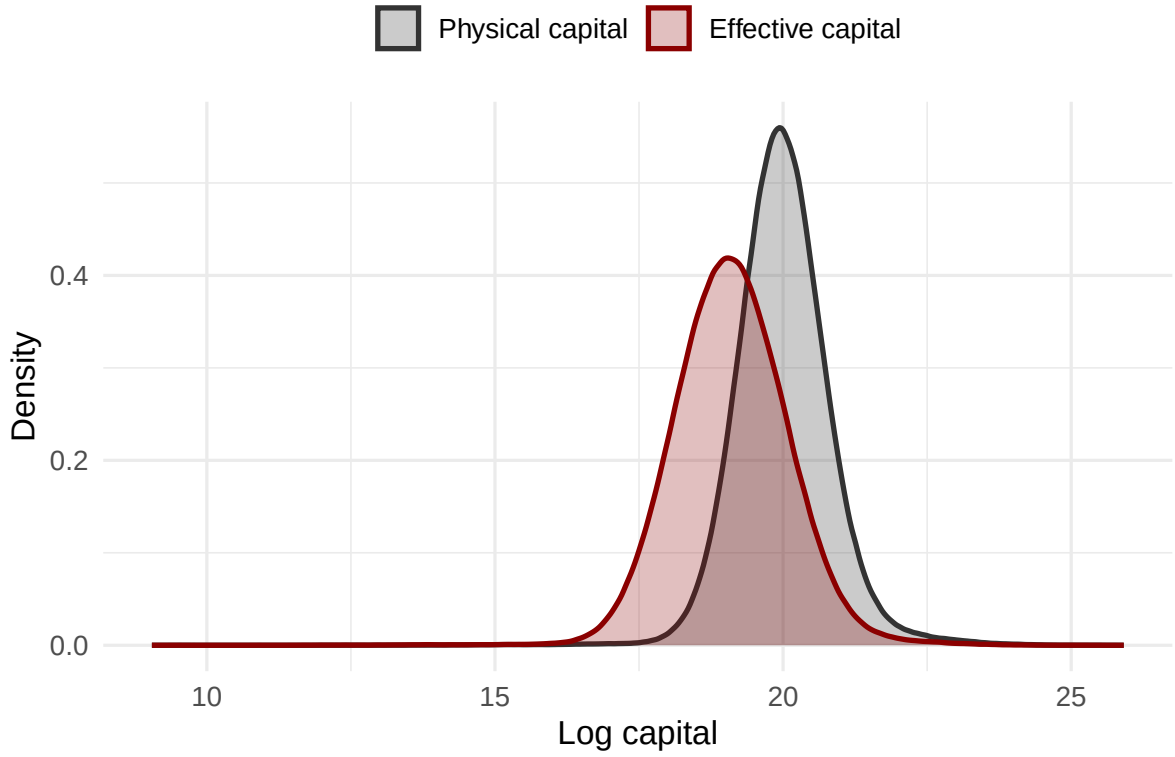


FIGURE 3.2. Distribution of Tract-Level Capital

Notes: The figure plots the panel density of tract-level physical capital k_t^i and effective capital $\tilde{k}_t^i$ over the 2002–2022 period. Effective capital is defined as $\tilde{k}_t^i = k_t^i \omega_i^{SVI} \omega_i^{CRI}$, where ω_i^{SVI} and ω_i^{CRI} denote the vulnerability and resilience wedges derived from the normalized Social Vulnerability Index (SVI) and from one minus the Community Resilience Index (CRI), respectively. Each wedge takes the exponential form $\exp(-\theta_X X)$, for $X \in \{SVI, CRI\}$ with parameters $\theta_S = 1$ and $\theta_C = 0.8$. They scale physical capital to reflect systematic differences in exposure, structural fragility, and recovery capacity across census tracts. The horizontal axis reports log capital, and the vertical axis the corresponding kernel density.

Effective capital captures the idea that identical physical stocks deliver different levels of capital services across locations. Tracts with higher vulnerability or lower resilience experience a proportional reduction in capital services, so that the same physical asset yields fewer effective units. By construction, the wedges are multiplicative and bounded above by one, implying that local characteristics can only reduce, but not increase, the effective value of capital.

Allocation of County-Level Losses. Observed county-year losses NDL_t^j are allocated across hit census tracts using an aggregation-consistent rule designed to generate rare but severe localized destruction and smaller impacts in other affected tracts.[13] The tract-level depreciation rate is defined as

$$d_t^i = \frac{\ell_t^i}{\tilde{k}_t^i}, \tag{20}$$

where ℓ_t^i denotes allocated loss in tract i during year t. By construction, tract-level losses aggregate to observed county-year losses.

This allocation mechanism yields a highly right-skewed distribution of positive depreciation rates, as shown in Figure 3.3. Most affected tract-year observations exhibit small losses relative to capital, while a small fraction experience severe destruction consistent with localized disaster epicenters. The procedure therefore reproduces the fat right tail observed in micro-level disaster data, in which extreme capital losses are concentrated geographically rather than evenly distributed.

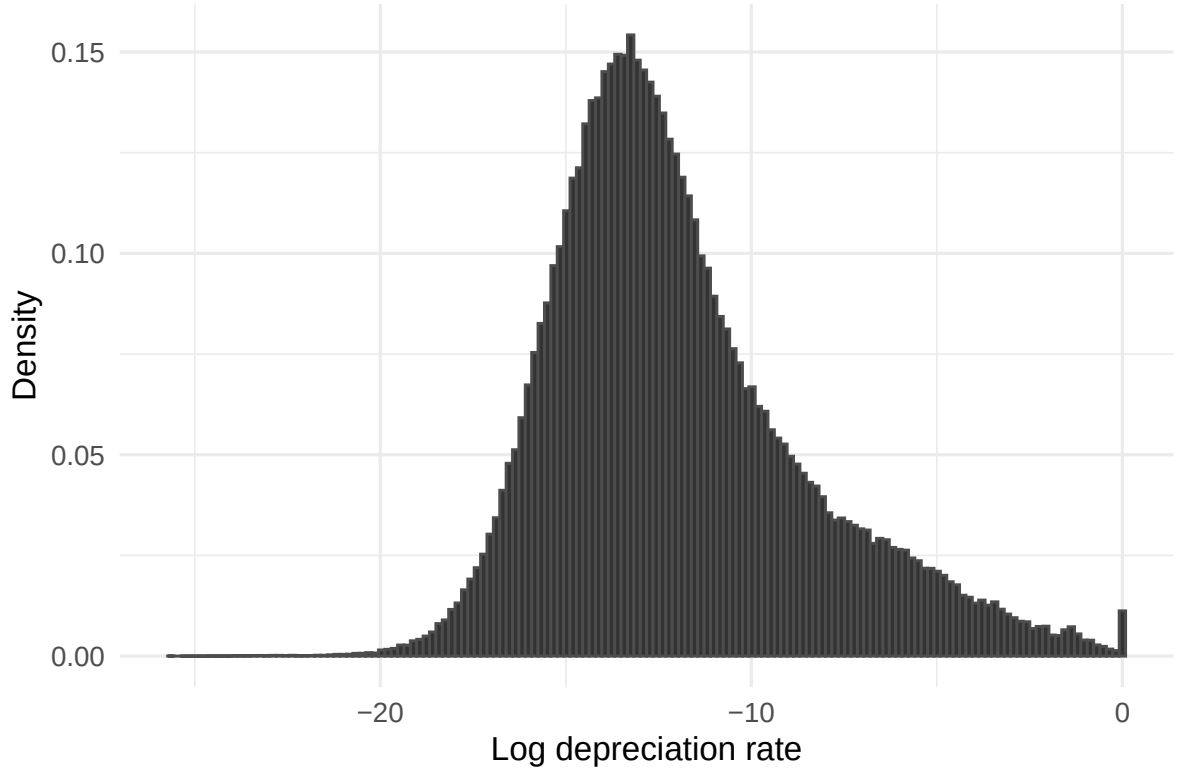


FIGURE 3.3. Distribution of Census-Tract Depreciation Rates

Notes: The figure plots the distribution of log tract-level depreciation rates for observations with positive losses. Depreciation is defined as realized allocated losses divided by census tract effective capital. The distribution is strongly right-skewed: most tract-year observations exhibit small depreciation rates, while a thin upper tail reflects localized severe destruction concentrated in a small number of tracts.

[13]The full allocation algorithm and parameter choices are described in Appendix B.6.

Summary statistics confirm this concentration of destruction. Table 3.2 quantifies the asymmetry visible in Figure 3.3. While the mean depreciation rate remains modest, the upper tail is economically significant. In particular, the 99th percentile indicates that a small fraction of tract-year observations experience large capital losses relative to effective capital. Although the allocation rule permits up to 20% baseline depreciation at the epicenter, such magnitudes are seldom reached: the 99th percentile is approximately 13.5%, indicating that county-level losses are typically insufficient to achieve a full 20% capital reduction across an entire tract, a unit that represents, on average, about 4,000 individuals. Extreme tract-level destruction is rare and mechanically constrained by aggregate losses.

The construction, therefore, serves to map aggregate county-level losses into a micro-level depreciation process that preserves aggregate discipline while allowing for spatial concentration. Without such a mapping, county averages would mechanically dilute extreme local shocks, understating the effective capital risk faced by households located in directly affected tracts.

TABLE 3.2. Distribution of Tract-Level Capital Depreciation Rates

Obs.>0	Mean (%)	Median (%)	P90 (%)	P95 (%)	P99 (%)
278183	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 mean, the median, and percentiles (P90, P95, and P99) are computed over all tract-year observations with positive reported damages. All values are expressed in percentage points.

The resulting dataset is a balanced tract–year panel for 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 4.

3.4. **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. (13).[14] 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 3.3, "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.[15] Wildfires similarly generate elevated conditional losses in the counties where they occur, despite a smaller national count of events.[16]

The aggregate distribution of depreciation rates d_t^i is depicted in Figure 3.4, 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 3.4 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

[14] In 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.

[15] Major 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.7.

[16] 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. 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 3.3. 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.

upper tail of the depreciation distribution of d, reaching 23%, whereas most other hazards 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 *percentile*:		
		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 3.4. 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.

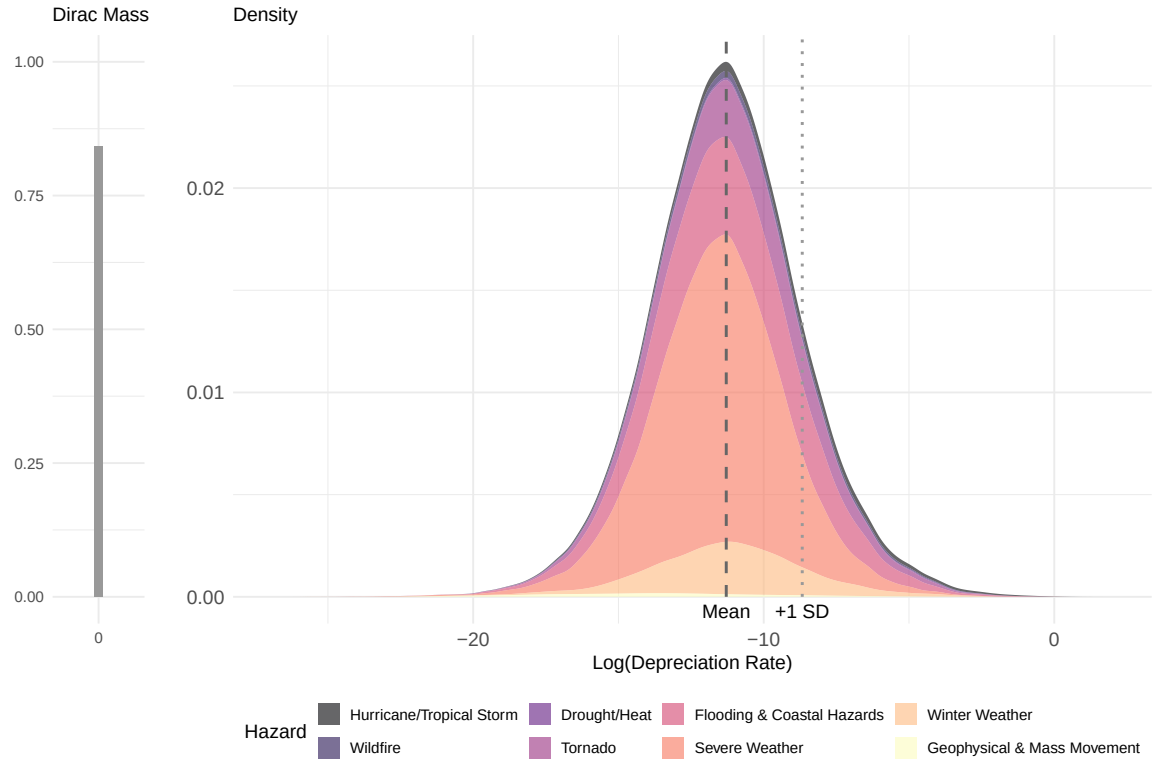


FIGURE 3.4. Disasters' Depreciation Rate Distribution

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.

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.

To illustrate this spatial heterogeneity, Figure 3.5 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 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

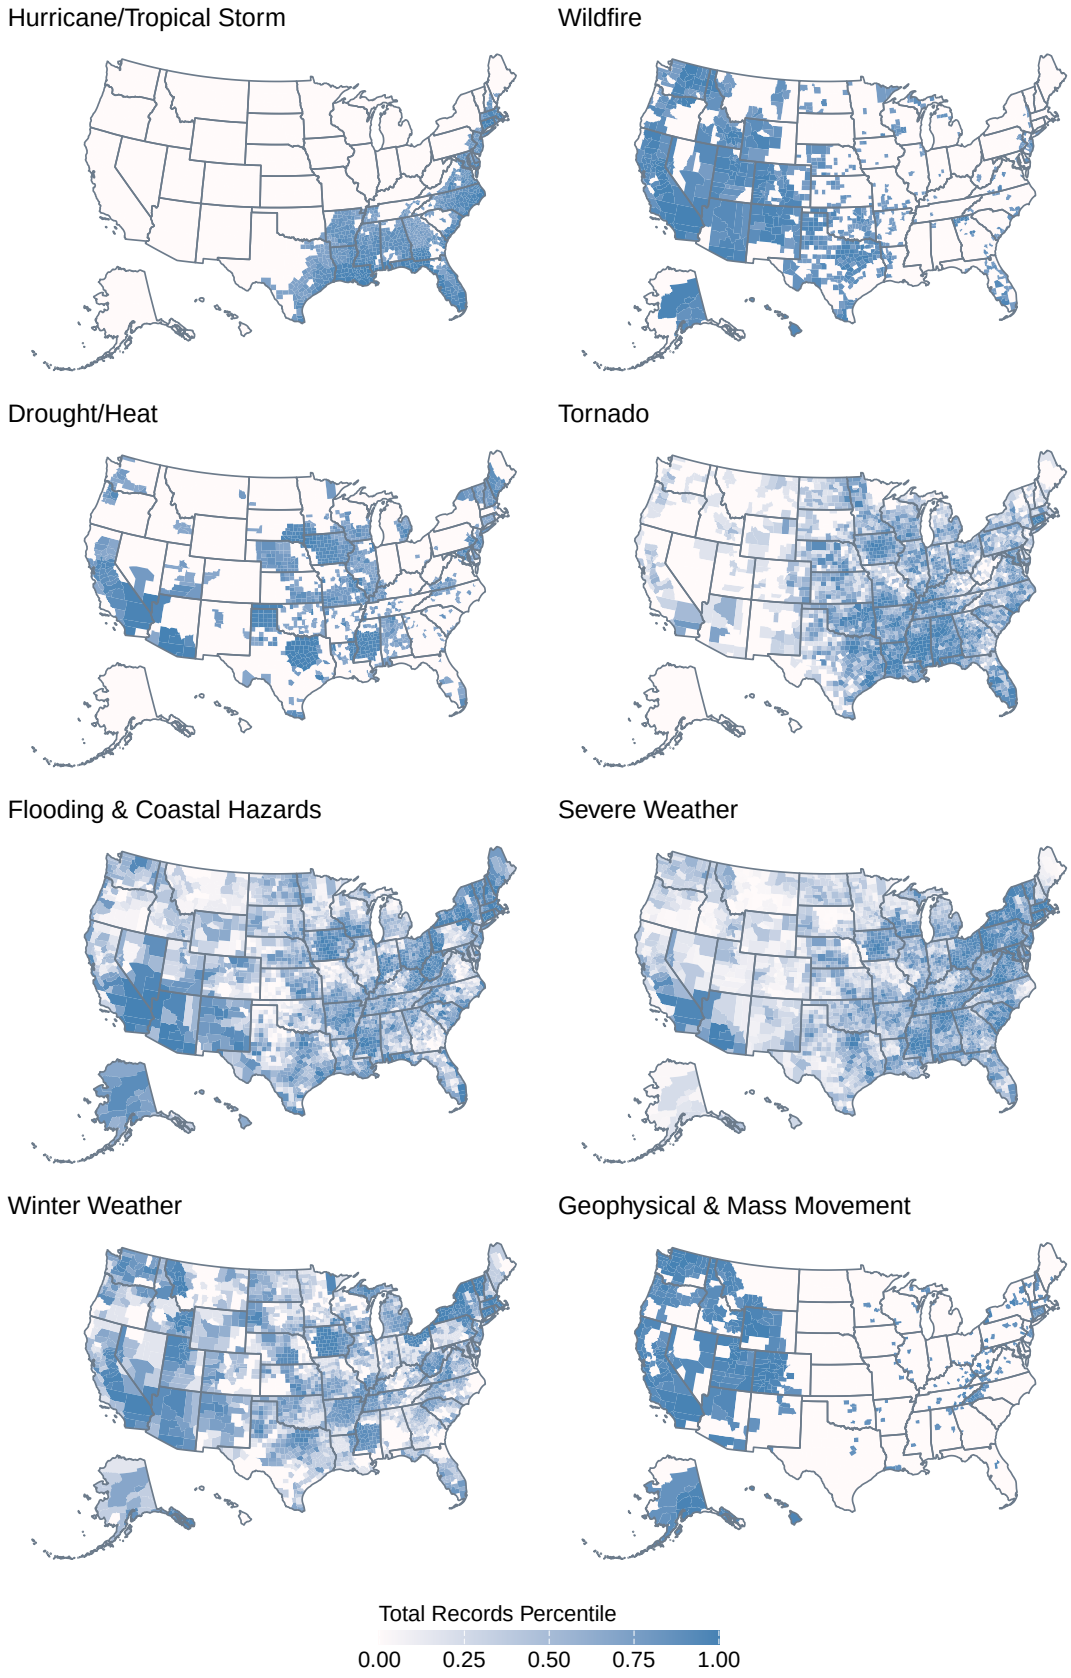


FIGURE 3.5. Spatial Distribution of Natural Hazard Records across U.S. Counties

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.

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 display a similar coastal pattern 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 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. Natural disasters not only differ across hazard types but also in their spatial concentration. This variation supports empirical strategies that allow the distribution of shocks to differ across locations. In particular, it motivates a framework in which agents are exposed to distinct climate zones, each characterized by its own distribution of disaster realizations.

4. Climate Change Process

This section outlines the empirical strategy for estimating the Markov process governing disaster-induced depreciation rates across census tracts. The objective is to identify (i) regime-specific depreciation intensities and (ii) transition probabilities across latent disaster regimes. In the next section, I embed this Markov process in my macroeconomic model to study the implications of disaster risk on saving behavior.

For tractability, I abstract from cross-regional interdependence and estimate the Markov process separately for each climate zone $z \in \mathcal{Z}$. Tracts are first assigned to their respective climate zone, and estimation is conducted independently within each group using only observations from that climate zone. For notational simplicity, the climate zone subscript z is omitted throughout, although all expressions are understood to be conditional on the corresponding climate zone.

4.1. **Estimation Strategy.**

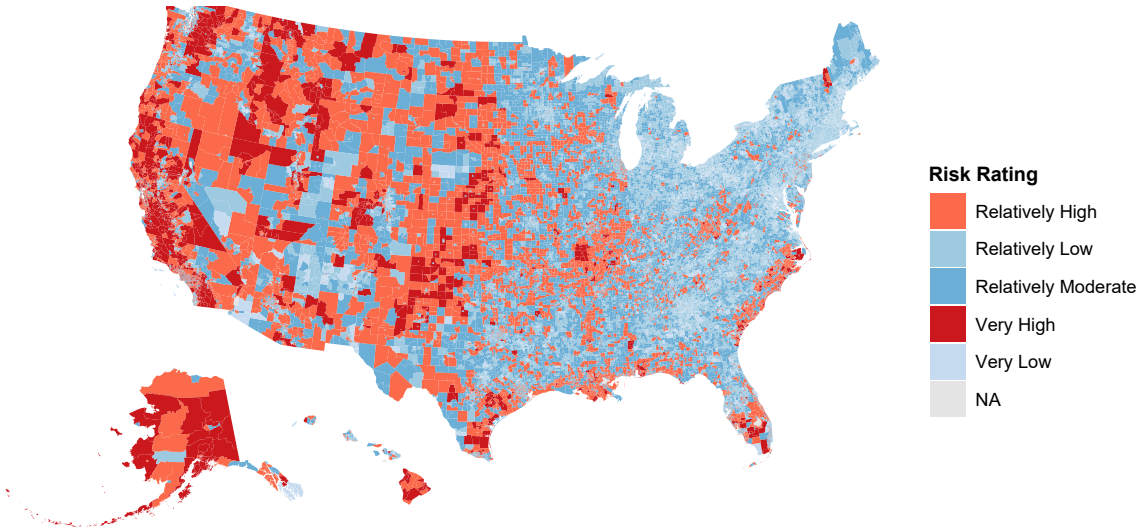


FIGURE 4.1. Climate Zones

Notes: The figure displays the geographic distribution of FEMA's National Risk Index (NRI) rating at the census-tract level. Census tracts classified as *Relatively High* or *Very High* are grouped into the *intensive* climate zone (shown in red), while tracts classified as *Very Low*, *Relatively Low*, or *Relatively Moderate* form the *moderate* climate zone (shown in blue). Gray areas correspond to missing values.

Moderate and Intensive Climate Zones. The constructed census-tract panel is partitioned into two climate zones: a *moderate* zone (M) and an *intensive* zone (I), such that $\mathcal{Z} = \{M, I\}$. The classification is based on FEMA's National Risk Index (NRI) rating at the census-tract level (Federal Emergency Management Agency, 2025).

The NRI rating transforms underlying quantitative risk measures into qualitative categories. The construction of the rating is designed to satisfy two principles: (i) census tracts within a given category are as similar as possible in terms of composite risk exposure, and (ii) census tracts across categories are as distinct as possible. The resulting classification consists of five ordered categories: *Very Low*, *Relatively Low*, *Relatively Moderate*, *Relatively High*, and *Very High*.

I define the *intensive* climate zone as the set of census tracts classified as *Relatively High* or *Very High*, while the remaining tracts constitute the *moderate* zone. This aggregation preserves FEMA's ordinal risk ranking while allowing for a parsimonious two-zone structure suitable for structural estimation.

Figure 4.1 displays the geographic distribution of the two zones. Intensive tracts, which represent 16% of census tracts in the sample, are predominantly located along coastal

regions and in areas historically exposed to wildfires and hurricanes. This spatial pattern is consistent with the hazard exposure documented in Figure 3.5 and the summary statistics reported in Table 3.3.

Hamilton Filter and Maximum Likelihood Estimation. Within each climate zone, census tracts are assumed to be independent and identically distributed. Let $i = 1, \ldots, N$ index tracts and $t = 1, \ldots, T$ index years. The observed depreciation rate in tract i and year t is denoted by $d_t^i \geq 0$.

Let $\upsilon_t^i \in \{e_1^{(U)}, \ldots, e_U^{(U)}\} \subset \mathbb{R}^U$ denote the unobserved regime indicator, where $e_u^{(U)}$ is the selection vector corresponding to the u-th column of the U-dimensional identity matrix. The latent regime follows a first-order Markov chain with transition matrix $\Pi = (\pi_{uu'})$ satisfying

$$\mathbb{P}\left(\upsilon_{t+1}^i = e_{u'}^{(U)} \mid \upsilon_t^i = e_u^{(U)}\right) = \pi_{uu'}, \qquad \sum_{u'=1}^{U} \pi_{uu'} = 1. \tag{21}$$

Conditional on the latent regime, positive depreciation rates follow an inverse-logistic distribution:

$$\log\left(\frac{d_t^i}{1 - d_t^i}\right) \mid \upsilon_t^i \sim \mathcal{N}\left(\mu^\top \upsilon_t^i, \sigma^2\right), \tag{22}$$

where $\mu = (\mu_1, \ldots, \mu_U)'$ collects regime-specific means and σ is a common standard deviation.

Zero Depreciation and Mixture Structure. A substantial fraction of tract-year observations exhibits zero depreciation, corresponding to years without realized disaster damage. To account for this feature, I introduce a two-component mixture. Conditional on regime u, depreciation is zero with probability $1 - p_u$ and strictly positive with probability p_u.

Let $p = (p_1, \ldots, p_U)'$. The conditional density is

$$f(d_t^i \mid \upsilon_t^i; \mu, \sigma, p) = (1 - p^\top \upsilon_t^i)\, \mathbf{1}_{\{d_t^i = 0\}} + (p^\top \upsilon_t^i)\, \mathbf{1}_{\{d_t^i > 0\}}\, \phi\left(\log\left(\frac{d_t^i}{1 - d_t^i}\right); \mu^\top \upsilon_t^i, \sigma\right), \tag{23}$$

where $\phi(\cdot)$ denotes the normal probability density function.

Panel Likelihood. For a panel dataset, let $D_t = \{d_t^i\}_{i=1}^N$ denote the N-dimensional vector of depreciation rates and $U_t = \{(\upsilon_t^i)'\}_{i=1}^N$ the corresponding stacked matrix of latent states.

Conditional on the latent state vector U_t, inverse-logistic transformed positive depreciation rates are independent across tracts and normally distributed:

$$\log\left(\frac{D_t}{1-D_t}\right)|X_t \sim \mathcal{N}_N\left(U_t\mu, \mathbf{\Sigma_t}\right), \tag{24}$$

where $\mathbf{\Sigma_t} = \text{diag}\left(\sigma^2\right)$ is an $N \times N$ diagonal variance-covariance matrix.

In the general specification, the algorithm estimates the parameters

$$\theta = \left[\{\mu_j\}_{l=1}^{J}, \sigma, \{p_j\}_{l=1}^{J}, \{\pi_{lm}\}_{l,m=1}^{J-1}\right],$$

for every zone independently using a Hamilton filter (Hamilton, 1994) and solving with maximum likelihood estimation.[17]

However, in the structural implementation described in the following section, I impose additional parameter restrictions to ensure a transparent mapping between latent regimes and economically interpretable disaster states.

Parameter Estimation. Within each climate zone, the depreciation process is modeled with $U = 3$ latent regimes. In the quantitative depreciation-shock Aiyagari model, disaster risk is summarized by regime-dependent destruction intensities and a transition matrix governing persistence. To ensure a transparent mapping between the empirical process and the structural model, I impose $p = (0, 1, 1)^\top$. The first regime, therefore, corresponds to a no-disaster state with degenerate zero depreciation, while the remaining two regimes generate strictly positive destruction.

The third regime is disciplined to capture rare but severe events by targeting the 99.5th percentile of the empirical depreciation distribution. This ensures that the model reproduces economically meaningful tail risk without overfitting extreme realizations.

To ensure that heterogeneity is captured primarily by discrete regime differences rather than within-regime dispersion, I impose a common standard deviation σ across states within each zone and introduce a mild penalty on large values of σ in the likelihood function. This regularization discourages excessive within-regime smoothing and favors sharper regime separation.

The motivation is structural: in the depreciation-shock Aiyagari model, disaster risk is represented by a finite vector of depreciation shocks and a transition matrix. Large within-regime dispersion would blur the mapping between the regime-switching estimation and the discrete shock structure used in the quantitative model. By disciplining σ,

[17] Implementation details are provided in Appendix C.

the estimation recovers well-defined peaks corresponding to economically interpretable depreciation states.

Estimated parameters for both climate zones are reported in Table 4.1.

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

Common Parameters			
p	0	1	1
σ		1	

Moderate Zone

	s_{low}	s_{medium}	s_{high}	π
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}	π
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 4.1. 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.

4.2. **Estimation Results.** Table 4.1 shows that the intensive zone is characterized by both higher regime-specific means and greater persistence of high-damage states. In particular, the severe regime exhibits a substantially larger destruction intensity and a higher probability of remaining in that state relative to the moderate zone. This implies that households in intensive tracts face both larger conditional losses and more persistent exposure to elevated disaster risk.

The common within-regime dispersion remains limited due to the penalization of large σ, ensuring that cross-state heterogeneity is primarily driven by differences in regime means and transition probabilities rather than excessive smoothing within states.

Figure 4.2 evaluates the estimated mixture against the empirical distribution. Panel A shows the aggregate distribution, while Panels B and C report the decomposition by climate zone. The left bars display the Dirac mass at zero depreciation, and the right panels compare the empirical kernel density (dashed line) with the estimated mixture density (solid line).

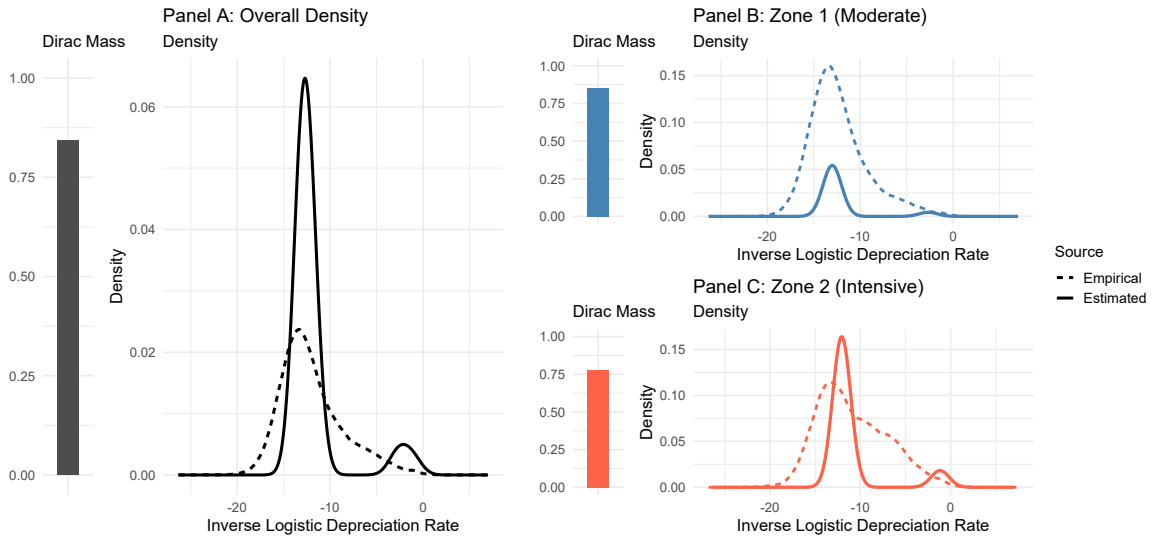


FIGURE 4.2. Climate Zones

Notes: Panel A displays the overall distribution of inverse-logistic depreciation rates. The left bar reports the mass of zero depreciation (Dirac mass), computed as the empirical share of infinite inverse-logistic observations. The right panel shows the continuous density of strictly positive depreciation rates. The dashed line corresponds to the empirical kernel density, while the solid line represents the estimated mixture distribution.
Panels B and C repeat the same decomposition separately for Zone 1 (Moderate) and Zone 2 (Intensive). In each case, the left bar reports the zone-specific Dirac mass, and the right panel displays the continuous component of the distribution.
In all panels, the estimated densities and discrete spike components are multiplied by their respective zone shares λ_z. Formally, each estimated mixture weight π_k^z and its associated density component are scaled by λ_z, so that the total contribution of zone z to the overall distribution equals its population share. This ensures that the combined distribution integrates to the aggregate mass and that each zone contributes proportionally to its weight in the data.

Importantly, the objective of the estimation is not to obtain a perfectly smooth fit of the empirical kernel density. Rather, the mixture specification is designed to recover a small number of economically interpretable peaks corresponding to distinct disaster regimes. Consistent with this objective, the estimated densities exhibit pronounced modal structure reflecting the discrete regime means, while abstracting from minor fluctuations in the empirical kernel.

The intensive zone (Panel C) displays both a smaller Dirac mass and a thicker right tail relative to the moderate zone (Panel B), consistent with higher destruction intensity and greater exposure to severe states. The estimated mixture captures these features through higher regime means and more persistent high-damage states.

Figure 4.3 further illustrates how the estimated regimes map into the structural environment. The left bar decomposes the total Dirac mass into zone-specific contributions

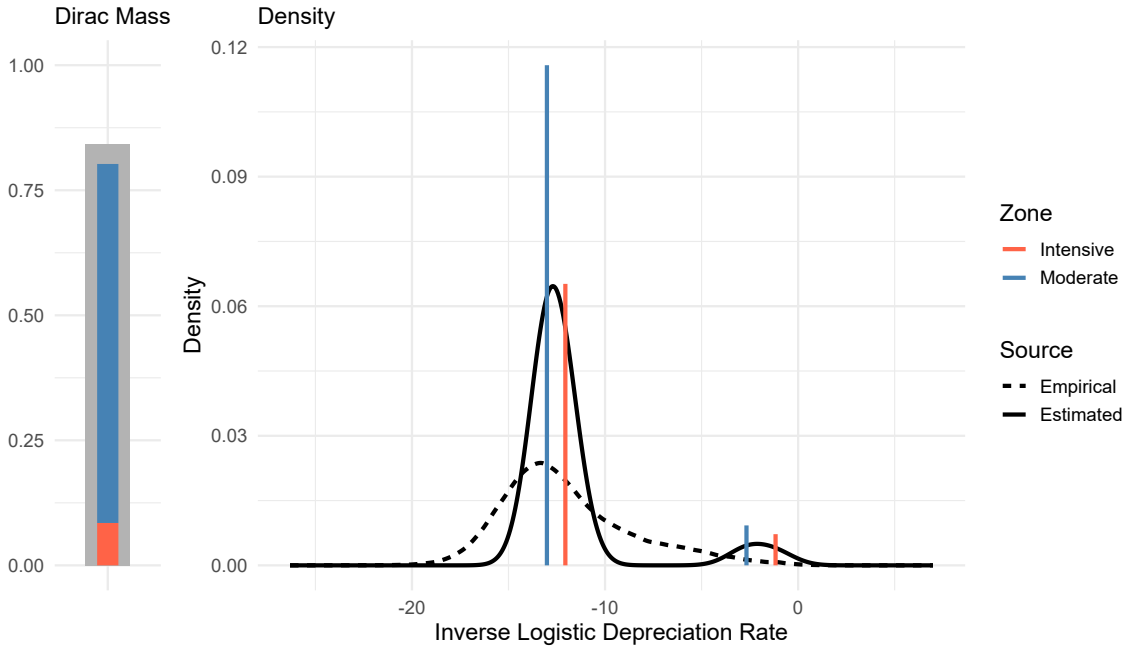


FIGURE 4.3. Climate Zones

Notes: This figure presents the overall distribution of inverse-logistic depreciation rates under the estimated mixture specification. The left bar reports the total Dirac mass at zero depreciation. The colored segments within the bar decompose this mass into the contributions of the Moderate (blue) and Intensive (red) zones. Each segment equals $\lambda_z \pi_0^z$, where π_0^z denotes the estimated probability of zero depreciation in zone z. The right panel displays the continuous component of the distribution. Vertical colored lines represent the discrete mixture spikes located at the estimated support points μ_k^z, with heights equal to $\lambda_z \pi_k^z$. The smooth black curve corresponds to the estimated continuous density, which is scaled by the zone shares, λ_z, prior to aggregation. Thus, both the discrete and continuous components of the estimated mixture are weighted by their respective zone shares λ_z, so that the share-weighted sum across zones gives the overall distribution.

$\lambda_z \pi_0^z$. The right panel displays the regime-specific spikes located at the estimated support points μ_k^z, with heights proportional to $\lambda_z \pi_k^z$. The resulting distribution is explicitly multi-modal by construction.

This discrete peak structure is intentional: in the structural model, each peak corresponds to a distinct depreciation shock that households internalize when forming expectations. The estimation, therefore, prioritizes identifying economically meaningful regimes over matching every local feature of the empirical density.

5. Quantitative Results

Calibration. I embed the results of my regime-switching process into my depreciation-shock Aiyagari model. Table 5.1 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 match the stationary mean of effective labor equal to 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
Productivity States (normalized)		
ξ_1	High productivity	1.02
ξ_2	Middle productivity	0.93
ξ_3	Low productivity	0.83

TABLE 5.1. 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.

Tail Depreciation Risk and Convergence. When depreciation is proportional to individual wealth, precautionary saving is strongly distorted: higher asset holdings mechanically increase exposure to catastrophic loss. In an environment where the only saving instrument is subject to a risk of destruction, households optimally limit wealth accumulation. Absent additional risk, the stationary equilibrium collapses toward the borrowing constraint, with the wealth distribution concentrating at zero.

To obtain a non-degenerate stationary equilibrium, I introduce a small productivity shock to labor income w that is common across zones. The values of the shock ξ are shown in Table 5.1. This specification is consistent with empirical evidence documenting income losses and labor market disruptions following natural disasters (Pleninger, 2022). The productivity component restores a standard consumption-smoothing motive, partially offsetting the endogenous disincentive to save created by proportional capital destruction.

Figure 5.1 reports the stationary consumption policies, saving schedules, and wealth distributions across climate zones under this calibration. Although productivity risk prevents collapse to the corner solution, households remain heavily concentrated near the borrowing constraint, particularly in the intensive climate zone (Zone 2), where exposure to tail depreciation risk is greatest.

The central mechanism is isolated in Figure 5.2. The figure compares Zone 1 and Zone 2 in the regime with no contemporaneous depreciation shock (i.e., the 0-shock

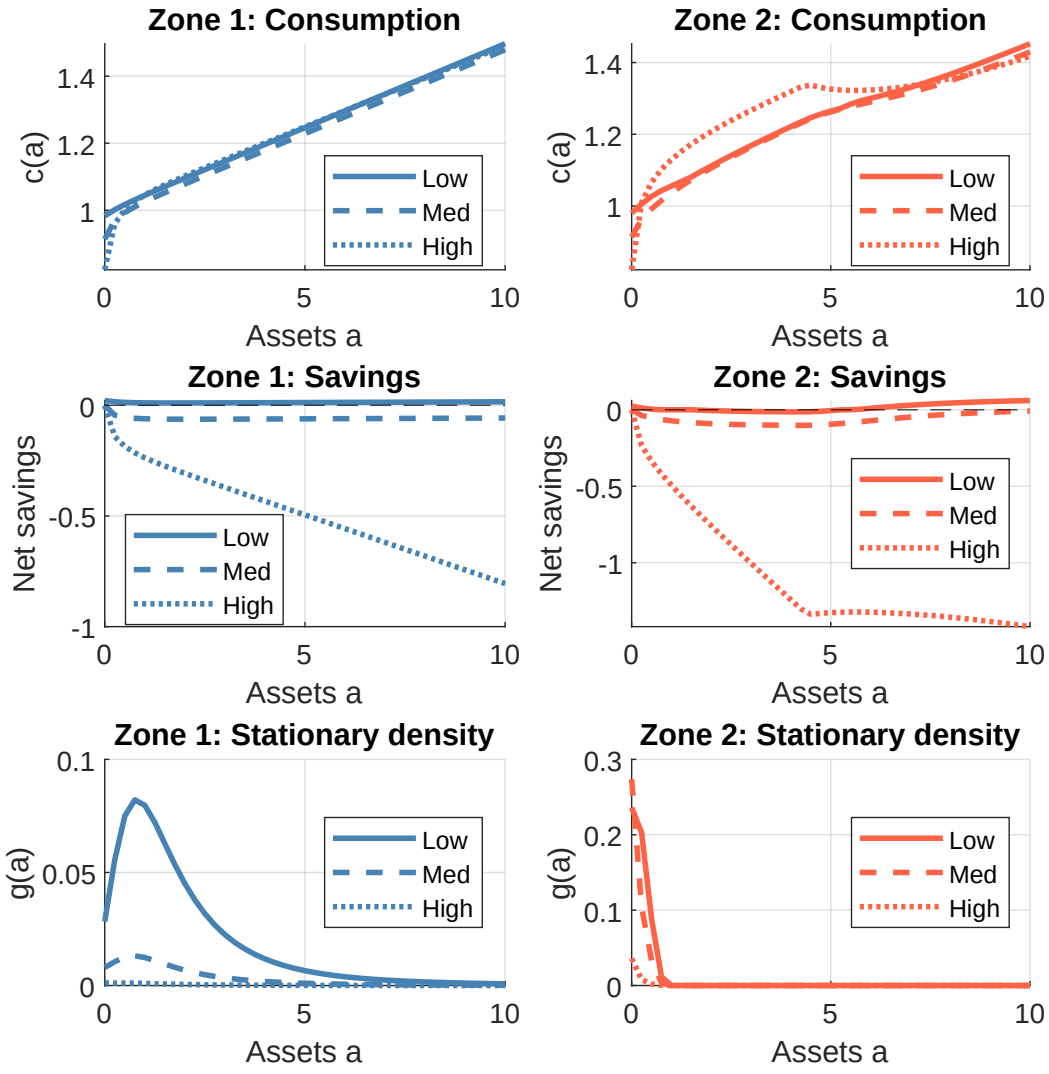


FIGURE 5.1. Stationary Policies and Wealth Distributions by Climate Zone

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 $\xi = [1.02, 0.93, 0.84]$. Differences across zones arise solely from cross-zone differences in the distribution of idiosyncratic depreciation risk. *Low*, *Med*, and *High* correspond to the household state and their 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% for Zone 1 and Zone 2, respectively.

state). Productivity risk is identical in realization; current labor income and preferences are therefore the same. In this state, there is no realized destruction of wealth.

Despite identical current incomes and the absence of contemporaneous disaster shocks, households choose different saving policies. In the intensive climate zone, saving is weaker

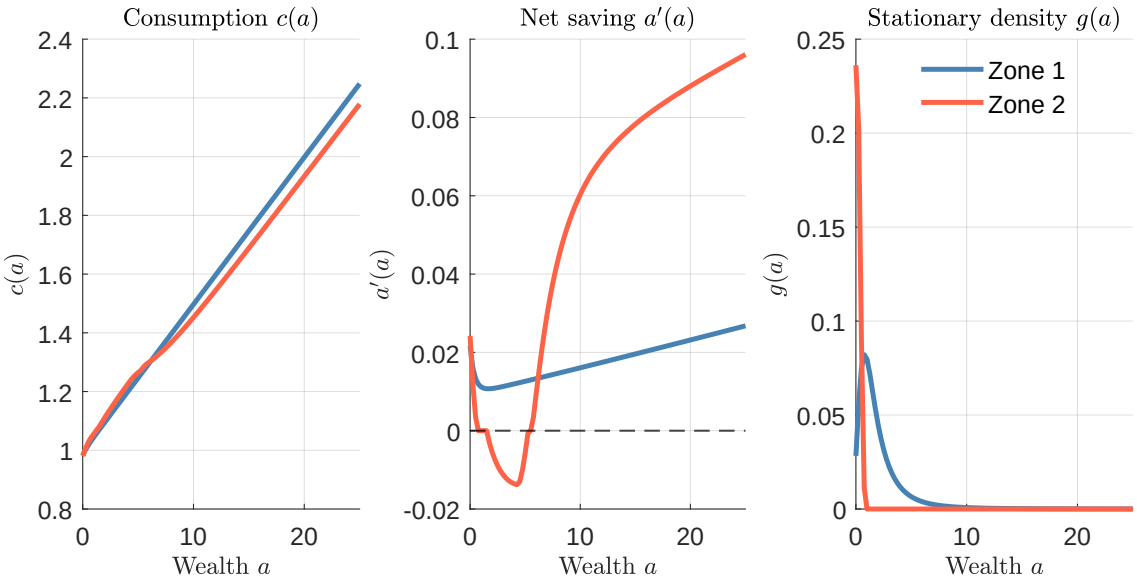


FIGURE 5.2. Identical Realization, Different Risk Distribution: Low-State Comparison

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.

at low wealth levels, and the stationary distribution places more mass at the borrowing constraint. At intermediate wealth levels, households self-insure more aggressively, reflecting their more adverse distribution of future depreciation shocks.

These differences arise solely from cross-zone heterogeneity in expected tail risk. The saving gap, therefore, measures endogenous insurance against a riskier environment: even when realized shocks are identical, households facing a more severe distribution of future destruction risk adjust their wealth accumulation behavior.

Aggregate and Distributional Effects. To disentangle mechanisms, I consider four model variants: M1 and M3 include only productivity risk ($\xi = [1.02, 0.93, 0.84]$ and $\xi = [1.04, 0.86, 0.69]$, respectively), common across zones, while M2 and M4 additionally incorporate idiosyncratic depreciation risk.[18]

[18]The distinction between M1 and M3 serves two purposes. First, it isolates the incidence of income dispersion by varying the spread of productivity shocks while holding climate risk absent. Second, it establishes a benchmark for assessing how differences in income risk shape the effects of disaster-induced depreciation once climate risk is introduced. Comparing M1 to M3, therefore, isolates the effect of income dispersion alone, allowing the interaction between productivity and depreciation risk to be identified when both are jointly present.

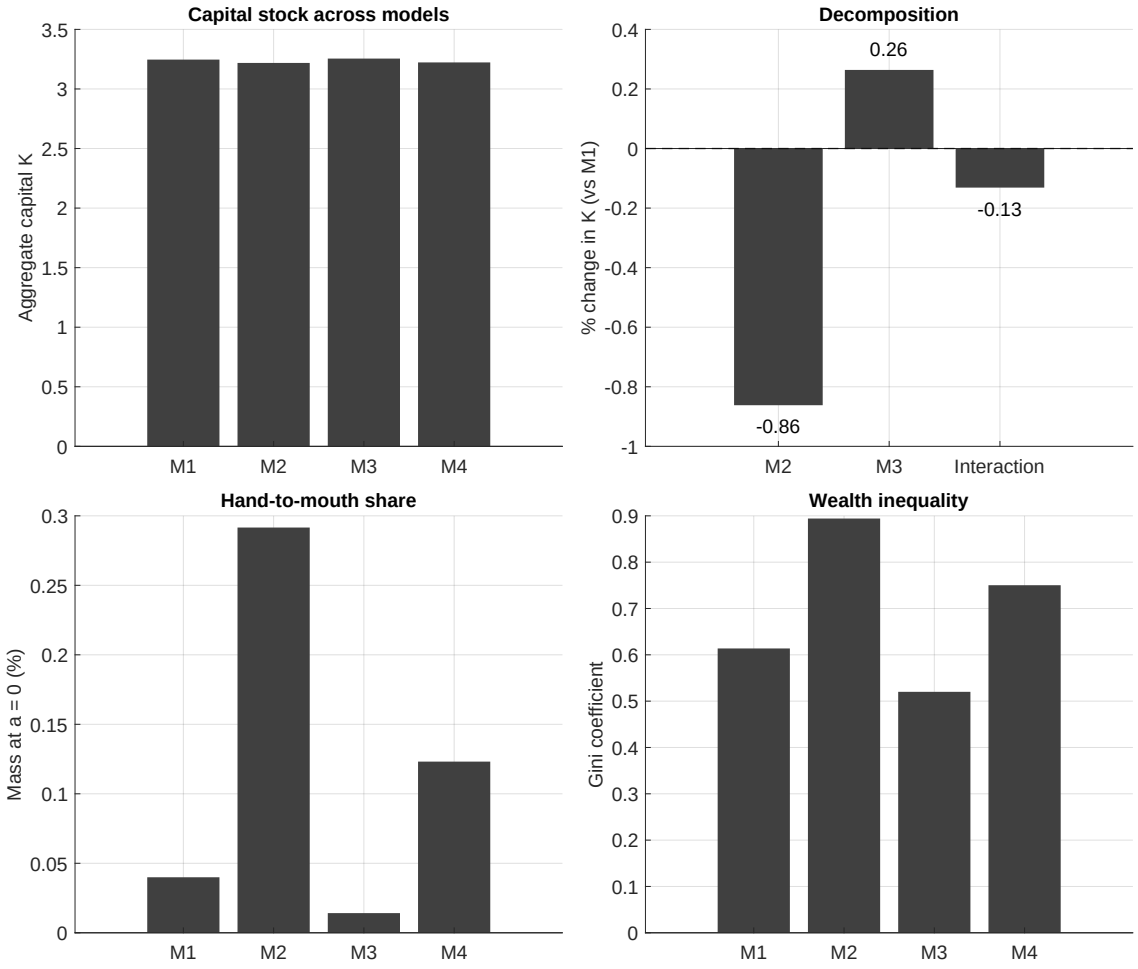


FIGURE 5.3. Aggregate Capital and Distributional Effects Across Model Variants

Notes: M1 and M3 include only productivity risk ($\xi = [1.02, 0.93, 0.84]$ and $\xi = [1.04, 0.86, 0.69]$, respectively), common across zones. M2 and M4 additionally include idiosyncratic depreciation risk. The upper-left panel reports aggregate capital. The upper-right panel decomposes the percentage change in capital relative to M1 into the effect of depreciation risk (M2), productivity risk (M3), and their interaction effect defined as $\Delta K^{\mathrm{M4}} - \left(\Delta K^{\mathrm{M2}} + \Delta K^{\mathrm{M3}}\right)$, with $\Delta K^{\mathrm{M}i}$ denoting the percentage change relative to M1. The lower panels report the mass at the borrowing constraint and the Gini coefficient of the stationary wealth distribution. The Gini coefficient is computed from the invariant wealth distribution $g(a)$ using the discrete formula $G = (N^2\bar{a})^{-1} \sum_{i=1}^{N} (2i - N - 1) a_{(i)}$, where $a_{(i)}$ denotes sorted wealth and $\bar{a}$ its mean.

Figure 5.3 summarizes the aggregate and distributional effects of tail depreciation risk. The upper-left panel shows that aggregate capital changes only modestly across specifications. Relative to M1, aggregate capital declines by approximately 0.9% in M2. The decomposition in the upper-right panel indicates that the interaction between productivity and depreciation risk is sub-additive: the combined effect is smaller in magnitude than the sum of the individual effects. Even under higher productivity dispersion (M4), the aggregate response remains below 1%.

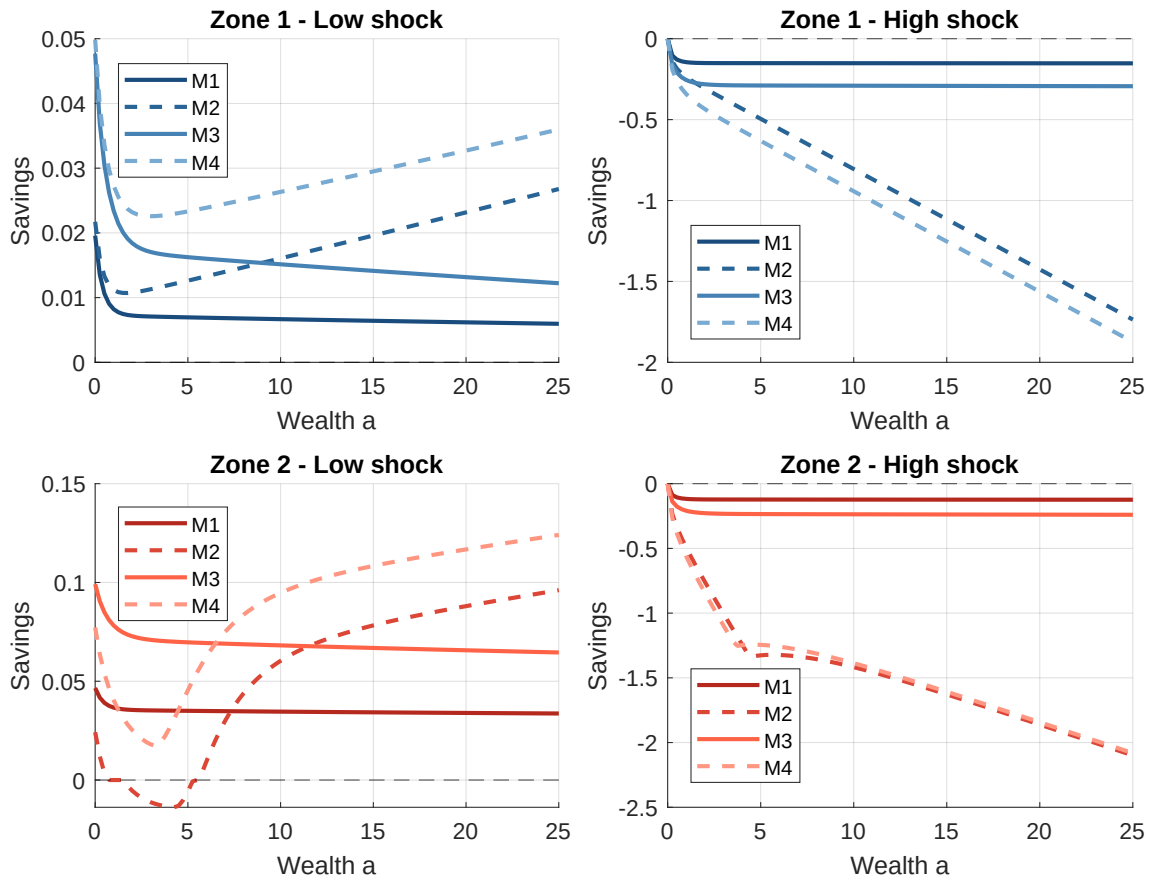


FIGURE 5.4. State-Contingent Saving Policies Across Climate Zones

Notes: The figure reports saving policies for different productivity shocks ($\xi_{1,2} = [1.02, 0.93, 0.84]$ and $\xi_{3,4} = [1.04, 0.86, 0.69]$), common across zones. The *Low* state corresponds to no additional shock (0-shock), while *High* corresponds to an additional depreciation rate of 6.4% and 23.7% for Zone 1 and Zone 2, respectively. Depreciation risk generates strong asymmetry: dissaving in the high-shock regime is steep, while precautionary saving in the low-shock regime is substantially flattened, particularly in the intensive climate zone.

In contrast, the lower panels reveal large distributional adjustments. Introducing depreciation risk sharply increases the mass of households at the borrowing constraint. Wealth inequality rises substantially, as reflected in the Gini coefficient.

Taken together, the figure shows that tail destruction risk operates primarily through a reallocation channel rather than a large aggregate contraction. The dominant effect is a compression of the wealth distribution toward the lower tail.

State-Contingent Saving Behavior. Figure 5.4 highlights the asymmetry between saving and dissaving across states. In the high-shock regime, households dissave aggressively in both zones. In the low-shock regime, however, precautionary saving is substantially muted when destruction risk is proportional to wealth.

The mechanism is straightforward: the marginal return to saving is endogenously affected by exposure to catastrophic loss. As wealth increases, so does potential destruction. This flattens saving schedules and generates a stationary equilibrium with a large mass of households at the borrowing constraint.

To quantify the magnitude of these policy distortions, Table 5.2 reports average net saving rates by wealth group, climate zone, and productivity state. Two patterns emerge. First, in the intensive climate zone, average net saving at low wealth levels is substantially lower than in the moderate zone, even in states without contemporaneous destruction. Second, the gap widens in adverse productivity states, where dissaving is considerably stronger under wealth-proportional destruction.

Wealth (a)	Moderate Zone (Z1)	Intensive Zone (Z2)	Relative Saving (Z2/Z1)
1	0.011	0.000	0.0
5	0.013	-0.006	-0.5
10	0.016	0.060	3.8
25	0.027	0.096	3.6
50	0.045	0.134	3.0

TABLE 5.2. Net Saving by Climate Zone at Selected Wealth Levels (Model M2)

Notes: The table reports net saving $a'(a) - a$ at selected wealth levels in Zone 1 and Zone 2, conditional on the zero-depreciation regime. Productivity realizations are identical across zones ($\xi = 1.02, 0.93, 0.84$), so differences arise solely from cross-zone variation in future depreciation risk. The first two columns report optimal savings at wealth level a in each zone. The third column reports the savings ratio (Z2/Z1), highlighting the non-linear response to climate risk exposure. At low wealth levels, households in the intensive climate zone save less than those in the moderate zone. As wealth increases, however, saving in the intensive zone rises substantially. In a stationary equilibrium, aggregate net saving equals zero, but saving policies differ markedly across zones and wealth levels.

These magnitudes confirm that the mechanism operates primarily through expectations: even when current income and realized depreciation are identical, differences in the distribution of future tail risk generate economically meaningful differences in saving behavior.

Concave Depreciation Risk. To assess whether wealth-dependent exposure is central to the results, I consider a concave destruction technology,

$$\tilde{d}(a) = \epsilon a^{\tau}, \qquad \tau = 0.5,$$

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

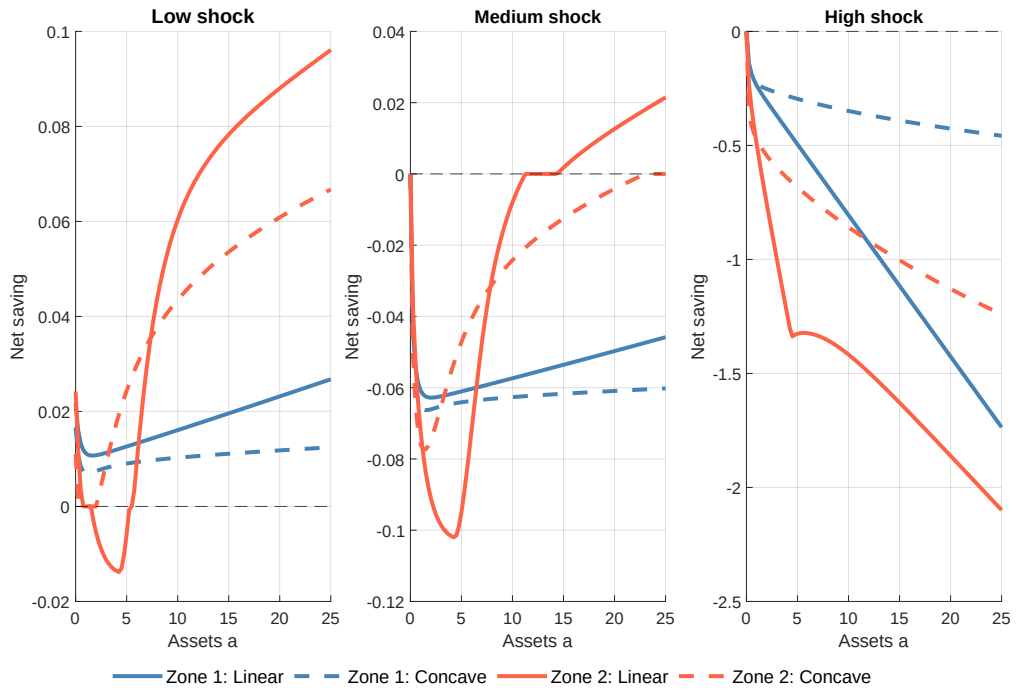


FIGURE 5.5. Linear versus Concave Depreciation Risk

Notes: The figure compares saving policies under a linear ($\tau = 1$) and concave ($\tau = 0.5$) depreciation technology with a three-state productivity shock ($\xi = [1.04, 0.86, 0.69]$). Concavity reduces marginal exposure to destruction at higher wealth levels. In the low-shock regime (0-shock), precautionary saving is restored relative to the linear specification. In the high-shock regime, dissaving persists but is substantially less abrupt.

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

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 proportionally increase effective exposure to destruction.

This specification contrasts with the baseline linear case, where exposure scales proportionally with wealth. While the empirical section 3.3 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 5.5 shows that concavity restores precautionary incentives. In the low-shock regime, saving schedules become upward sloping even at low wealth levels, particularly

in the intensive climate zone. Dissaving in the high-shock regime remains present but is less abrupt than under linear destruction.

However, concavity may generate excessive accumulation at higher wealth levels, as households build buffers to shield themselves from catastrophic states. Hence, the degree of curvature in the destruction technology critically shapes both precautionary motives and wealth dispersion.

Summary. Four results emerge:

(1) When destruction risk is proportional to wealth, precautionary saving is distorted because higher accumulation increases exposure to catastrophic loss. Even under identical realized shocks, differences in the distribution of future risk generate different saving behavior across climate zones.
(2) Aggregate capital responds modestly (below 1%), but the wealth distribution adjusts sharply: the mass of hand-to-mouth households increases substantially, and inequality rises.
(3) The distortion operates asymmetrically across states: dissaving in high-shock regimes is steep, while precautionary saving in low-shock regimes is muted.
(4) Reducing marginal exposure to destruction through concavity restores precautionary saving incentives, highlighting the central role of wealth-dependent risk.

Overall, tail climate risk operates predominantly through a distributional reallocation channel rather than an aggregate contraction channel, compressing the wealth distribution toward the lower tail while leaving aggregate capital largely unchanged

6. 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 fundamentally differs 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 the importance of asset composition and financial structure. In an economy where all wealth is exposed to destruction and no safe asset is available, self-insurance through accumulation is inherently limited. Policies or financial innovations that partially decouple wealth from physical destruction risk—through insurance design, resilient asset structures, or risk-sharing mechanisms—can therefore have disproportionately large distributional effects even if aggregate output changes little.

As climate change increases the frequency and severity of extreme events, understanding how tail risk interacts with household balance sheets becomes central. The evidence presented here suggests that rising climate risk may not primarily manifest as large aggregate contractions, but rather as a persistent reshaping of the wealth distribution and an expansion of financial fragility among exposed households.

An important direction for future research is to introduce a risk-free asset that is not exposed to climate destruction. Such an extension would allow households to hedge wealth-dependent exposure and would clarify how asset composition mitigates the distributional effects of localized tail risk. In equilibrium, a safe asset could be supplied either through inter-household trade when destruction risk is sufficiently dispersed or through financial intermediaries that can diversify risk at scale. Embedding these risk-sharing mechanisms in the present framework would help quantify the extent to which financial structure can attenuate climate-induced wealth compression.

Such an extension would allow a quantitative decomposition of the climate protection gap: how much risk households can absorb through precautionary accumulation, how much can be diversified through private insurance and financial intermediation, and how much must ultimately be borne by the public sector.

References

Angeletos, G.-M. (2007). Uninsured Idiosyncratic Investment Risk and Aggregate Saving. *Review of Economic dynamics 10*(1), 1–30.

Bakkensen, L. A. and L. Barrage (2022). Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics. *The Review of Financial Studies 35*(8), 3666–3709.

Barrage, L. and J. Furst (2019). Housing Investment, Sea Level Rise, and Climate Change Beliefs. *Economics letters 177*, 105–108.

Barro, R. J. (2009). Rare Disasters, Asset Prices, and Welfare Costs. *American Economic Review 99*(1), 243–264.

Benhabib, J., A. Bisin, and M. Luo (2019). Wealth Distribution and Social Mobility in the US: A Quantitative Approach. *American Economic Review 109*(5), 1623–1647.

Benhabib, J., A. Bisin, and S. Zhu (2011). The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents. *Econometrica 79*(1), 123–157.

Benhabib, J., A. Bisin, and S. Zhu (2015). The Wealth Distribution in Bewley Economies with Capital Income Risk. *Journal of Economic Theory 159*, 489–515.

Berkeley Earth (2026, January). Global Temperature Report for 2025. `https://berkeleyearth.org/global-temperature-report-for-2025/`. Accessed February 17, 2026.

Bilal, A. and E. Rossi-Hansberg (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.

Center for Emergency Management and Homeland Security (2023). Spatial hazard events and losses database for the united states (sheldus®).

Chen, H., S. Joslin, and N.-K. Tran (2012). Rare Disasters and Risk Sharing with Heterogeneous Beliefs. *The Review of Financial Studies 25*(7), 2189–2224.

Federal Emergency Management Agency (2025, December). 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.

Gomes, J. F., M. Grotteria, and J. A. Wachter (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.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.

Indaco, A. and F. Ortega (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*. Beijing: Integrated Research on Disaster Risk. 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.

Miller, M., J. D. Paron, and J. Wachter (2025). Sovereign default and the decline in interest rates. Technical report, National Bureau of Economic Research.

Moll, B., L. Rachel, and P. Restrepo (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.

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.

Woongchan, J., L. Barrage, and K. Walsh (2025). Pricing Climate Risks: Evidence from Wildfires and Municipal Bonds. Technical report, Economics Working Paper Series.

Appendix A. Model

A.1. **Model Solution.** The solution to the household's optimization problem (8) is characterized by a value function $V^R(a,z)$ and an associated policy function for next-period assets $a'(a,z)$ that satisfy the following first-order and complementary slackness conditions.

$$u'(c(a,z)) = \lambda(a,z), \tag{25}$$

$$\lambda(a,z) = \beta \sum_{z' \in \mathcal{Z}^R} \Pi^R(z'|z)\, \lambda(a',z') \big[1 + r^K - \delta^{z'}\big], \tag{26}$$

$$a'(a,z) \geq -\phi, \quad \lambda(a,z) \geq 0, \quad \lambda(a,z)\big[a'(a,z) + \phi\big] = 0. \tag{27}$$

Substituting the first-order condition into the Euler equation yields the standard **Euler condition**:

$$c(a,z)^{-\gamma} = \beta\, \mathbb{E}_{z'|z}\Big[c(a',z')^{-\gamma}\big(1 + r^K - \delta^{z'}\big)\Big], \tag{28}$$

where the **consumption policy** $c(a,z)$ is implied by the budget constraint:

$$c(a,z) = \big(1 + r^K - \delta^z\big)a + w - a'(a,z). \tag{29}$$

The pair of functions $\{a'(a,z), c(a,z)\}$ jointly solve the dynamic programming problem, and the stationary distribution $\mu^R(a,z)$ satisfies:

$$\mu^R(a',z') = \sum_{z \in \mathcal{Z}^R} \Pi^R(z'|z)\, \mu^R\big(a, z : a'(a,z) = a'\big), \tag{30}$$

ensuring consistency between individual decisions and aggregate outcomes.

A.2. **Technical Regularization.** When capital destruction is proportional to individual wealth and sufficiently severe, the model may fail to generate a stable stationary wealth distribution. Since the only saving instrument is itself exposed to idiosyncratic destruction risk, precautionary motives may be weakened when destruction risk is large. In such cases, optimal asset accumulation can collapse toward the borrowing constraint, leading to a degenerate invariant distribution.

To restore a standard consumption-smoothing motive, I introduce a small idiosyncratic component of labor productivity driven by the same Markov state x_t^i that governs capital destruction. Labor income becomes

$$y_t^i = w_t\, \xi(x_t^i),$$

where $\xi(x_t^i) > 0$ denotes a state-dependent productivity shifter.

Let Π denote the $J \times J$ transition matrix of the Markov chain, and let $\bar{\pi} = (\bar{\pi}_1, \ldots, \bar{\pi}_J)^\top$ be its stationary distribution satisfying $\bar{\pi}^\top = \bar{\pi}^\top \Pi$ and $\sum_{j \in \mathcal{J}} \bar{\pi}_j = 1$. Let $\tilde{\xi} = (\tilde{\xi}_1, \ldots, \tilde{\xi}_J)^\top$

denote the raw productivity levels associated with each state. The normalized productivity vector $\xi = (\xi_1, \ldots, \xi_J)^\top$ is defined as

$$\xi_j = \frac{\tilde{\xi}_j}{\sum_{k \in \mathcal{J}} \bar{\pi}_k \, \tilde{\xi}_k},$$

which implies

$$\sum_{j \in \mathcal{J}} \bar{\pi}_j \, \xi_j = 1.$$

This normalization ensures that aggregate labor income remains equal to w_t in the stationary equilibrium. The productivity component therefore introduces cross-sectional dispersion without altering aggregate quantities or the firm's problem. In the quantitative implementation, the model converges to a non-degenerate stationary wealth distribution.

To regularize the wealth dynamics, I introduce a small three-state earnings shifter common to both zones. Let the raw earnings vector be

$$\tilde{e} = (1 + \varepsilon, \, 1, \, 1 - \varepsilon)^\top,$$

with $\varepsilon > 0$ small. Each Markov state jointly determines a destruction realization and an earnings realization.

Because the transition matrices may differ across zones, the stationary probability distribution over states generally differs as well. Let π_z denote the stationary distribution of the Markov chain in zone $z \in \{1, 2\}$, and let λ_z denote the population share of zone z, normalized so that $\lambda_1 + \lambda_2 = 1$.

The economy-wide stationary mean of the raw earnings shifter is therefore

$$\bar{e} = \lambda_1 \, \pi_1^\top \tilde{e} + \lambda_2 \, \pi_2^\top \tilde{e}.$$

I normalize the earnings vector once and for all according to

$$e = \frac{\tilde{e}}{\bar{e}}.$$

By construction,

$$\lambda_1 \, \pi_1^\top e + \lambda_2 \, \pi_2^\top e = 1,$$

so that the stationary mean of labor income in the entire economy equals w_t. The earnings shifter therefore introduces cross-sectional dispersion without altering aggregate labor income, equilibrium prices, or the firm's problem.

A.3. **General Framework for Z-zone and U-regimes.**

Appendix B. Data

B.1. **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.[19]

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 18 hazard types. For parsimony, I group these into eight categories following IRDR (2014).[20]

B.2. **Hazard Categories Mapping.**

B.3. **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 Table B.2.

BEA groupings — grouped by *to* FIPS

From	To	Aggregate Name
[15005, 15009]	15901	Kalawao + Maui
[51003, 51540]	51901	Albemarle + Charlottesville
[51005, 51580]	51903	Alleghany + Covington

[19] Soloveichi (2021) provides estimates of cultivated assets up to 2019. Appendix B.4 reports robustness checks incorporating crop losses.

[20] Appendix Figure B.1 documents the correspondence between original SHELDUS hazard types and the eight aggregated categories.

[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.

B.4. **Property and Crop losses.**

B.5. **SHED.**

B.6. **Census Tract and Synthetic Depreciation Rates.** 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.[21]

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 that captures 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

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

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

TABLE B.3. Financial Impact Across SHED Survey Waves

	Wave 2024	Wave 2023	Wave 2022	Wave 2021
Panel A. Number of respondents				
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		
Panel B. Share of respondents in percent				
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.

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,

$$\text{NRI}_i = \text{EAL}_i \times f\left(\frac{\text{SVI}_i}{\text{CRI}_i}\right), \tag{31}$$

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.[22]

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.

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,

$$Records_t^j \sim \text{Poisson}(\lambda_j), \tag{32}$$

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,

$$\lambda_j = \frac{1}{T}\sum_{t=1}^{T} Records_t^j, \tag{33}$$

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

$$w_{i|j} = \frac{EAL_i}{\sum_{k\in j} EAL_k}, \tag{34}$$

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

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,

$$Records_t^i \sim \text{Poisson}(\lambda_i), \qquad \text{where} \qquad \lambda_i = w_{i|j}\lambda_j. \tag{35}$$

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

$$\mathbb{P}(hit^i) = 1 - \exp(-\lambda_i). \tag{36}$$

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

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, $\omega_{i|j} = \mathbb{P}(hit^i \mid hit^j)$. 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

$$\left(Records_t^i\right)_{i \in j} \mid Records_t^j \sim \text{Multinomial}\left(Records_t^j, \{\omega_{i|j}\}_{i \in j}\right). \tag{37}$$

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.[23]

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 3.1. Repeating the exercise over 200 independent draws yields a remarkably stable pattern across the full panel. At the annual level, exposure exhibits slightly greater

[23] Indeed, the multinomial distribution implies $\mathbb{E}\left(Records_t^i \mid Records_t^j\right) = \omega_{i|j} Records_t^j$.

variation, with the share of the population residing in hit tracts ranging from 13% to 22% across years.[24]

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.

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:[25]

$$\kappa_{i|j} \equiv \frac{BV_i^{2020}}{\sum_{k\in j} BV_k^{2020}}, \qquad \sum_{i\in j} \kappa_{i|j} = 1. \tag{38}$$

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

$$k_t^i = \kappa_{i|j}\, K_t^j. \tag{39}$$

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

$$\tilde{k}_t^i = k_t^i\, \omega_i^{SVI}\, \omega_i^{CRI}, \tag{40}$$

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.

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

$$\omega_i^{SVI} = \exp(-\theta_S sx^i), \tag{41}$$

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.

[24]Population is held fixed at 2020 Census levels for all census tracts.

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

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

$$\omega_i^{CRI} = \exp(-\theta_C cd^i), \tag{42}$$

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.

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:

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

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

$$d_t^i = \frac{\ell_t^i}{\tilde{k}_t^i}, \tag{43}$$

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

Figure B.2 disaggregates this distribution by hazard type and by major versus minor disaster designation, as classified under NCEI and PDD declarations. For each hazard, the left panel shows the relative share of observations with $d_t^i = 0$ (the Dirac mass), while the right panel plots kernel densities of $\log(d_t^i)$ conditional on $d_t^i > 0$, separately for major (red) and minor (blue) events.

Several regularities emerge: First, in most hazard types, the Dirac mass is higher for major events, suggesting that they indeed represent rather tail events (i.e., high severity but low frequency). However, the opposite occurs for hurricanes, tropical storms, wildfires, and droughts, as shown in Table 3.3, where the majority of disasters are labelled as major.

Second, the distribution of positive depreciation rates tends to shift to the right for major events. That is, conditional on $d_t^i > 0$, major disasters are associated with higher average losses and heavier upper tails. This pattern is most pronounced for the first three panels. An exception is Severe Weather, where the conditional density curves for major and minor events overlap to a great extent, suggesting weaker differentiation in economic damage despite classification differences. One explanation is that severe weather can occur at the same time as another severe hazard, such as a hurricane, but they are not the ones causing the larger share of losses.

- talk about the capital at 0 for 08014.

B.7. **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).[26]

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:

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

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

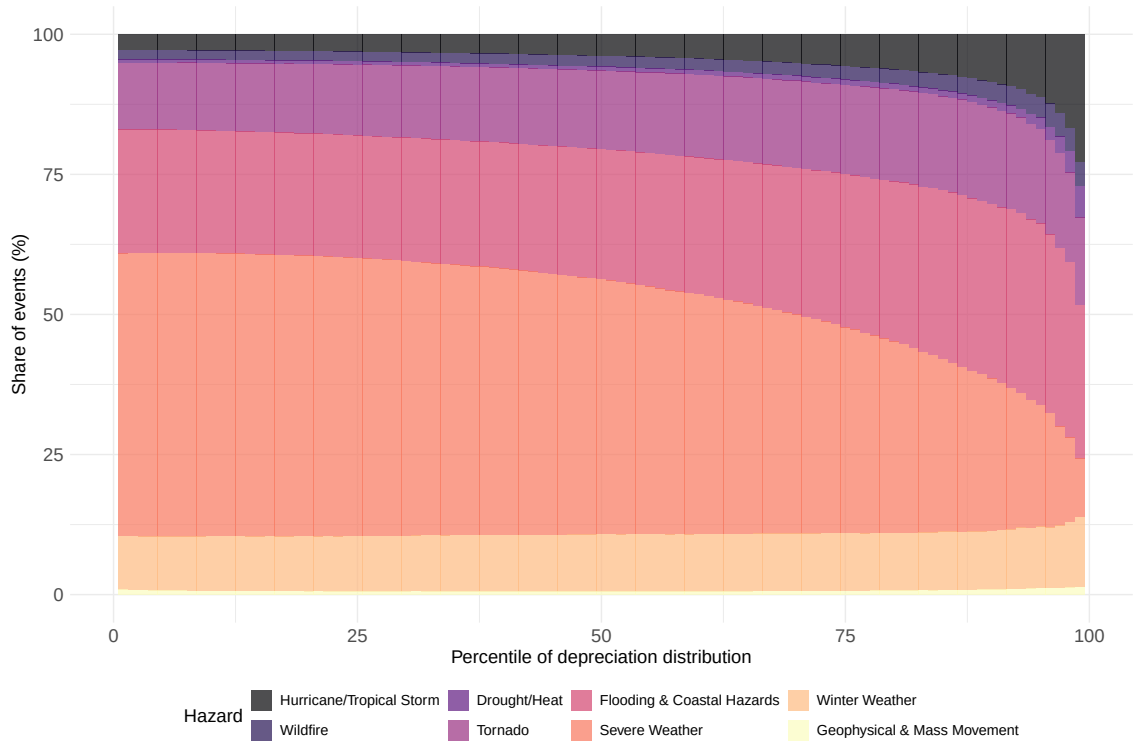


FIGURE B.1

Notes:

This procedure preserves all NCEI events while supplementing them with additional PDD observations that are absent from the NCEI-major set.[27,28]

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. Figure B.3 in Appendix B.7 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.

[27]In the assembled dataset, this reconciliation adds roughly 20% more county–year–hazard major-event observations relative to relying on NCEI alone.

[28]As the SHELDUS dataset is updated, discrepancies between PDD and SHELDUS decline over time; the PDD supplementation primarily affects the early 2000s (see also Gallagher 2023).

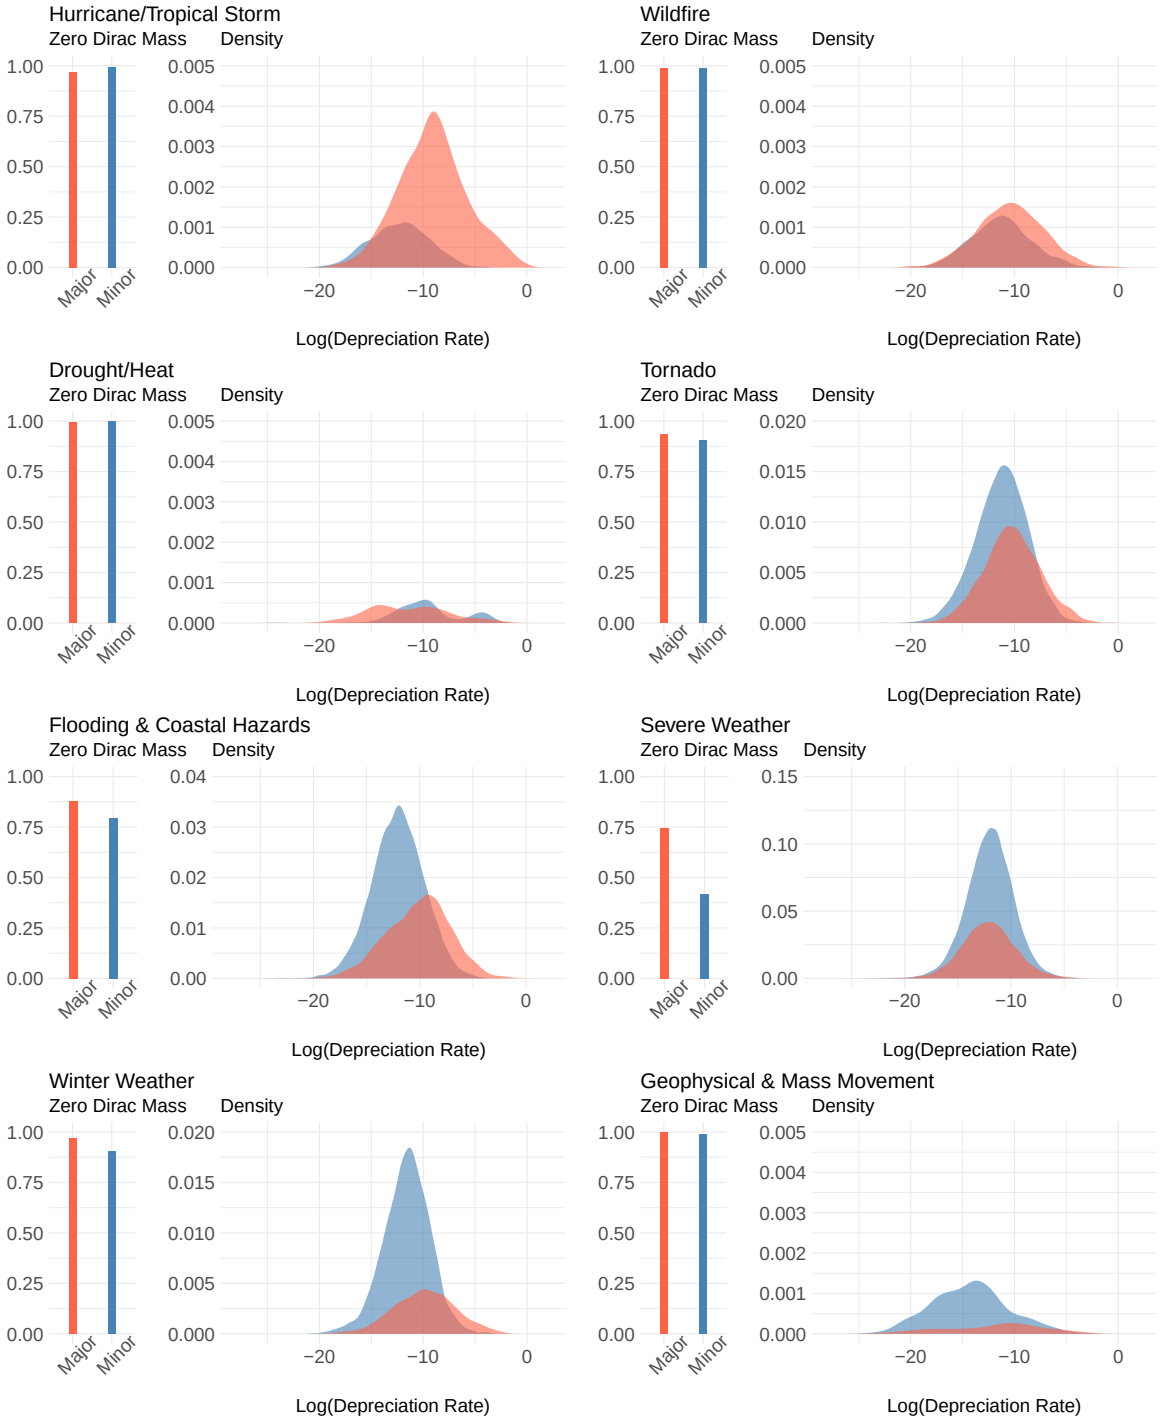


FIGURE B.2. Disasters' Depreciation Rate Distribution

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.

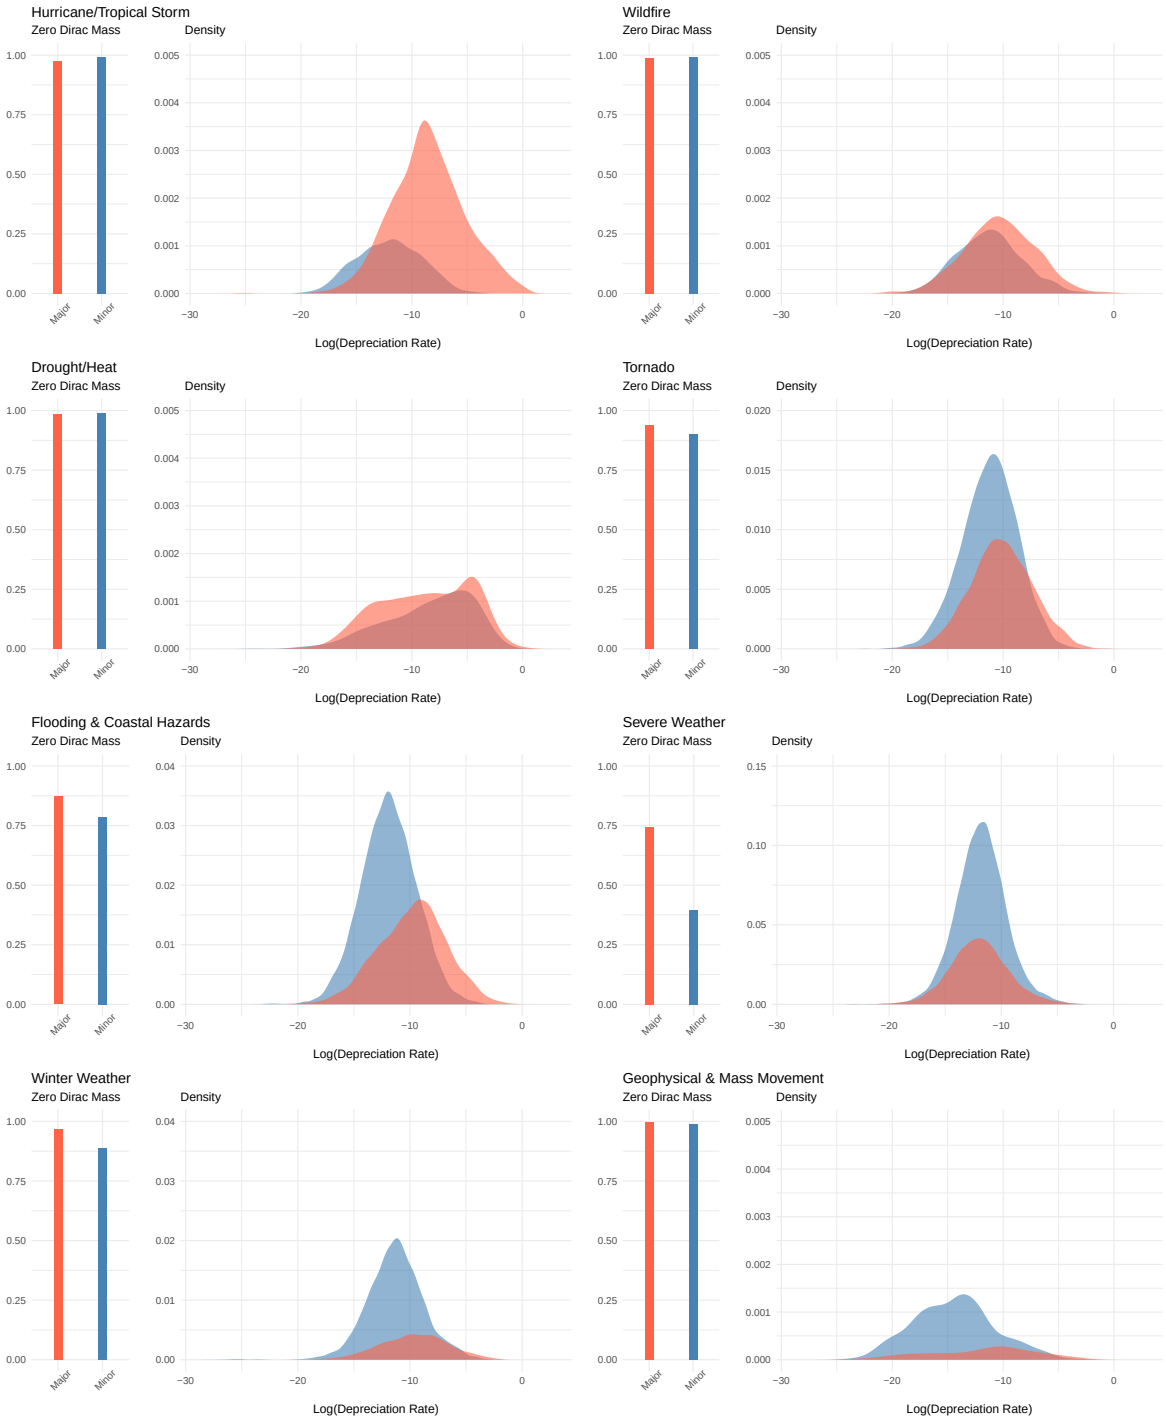


FIGURE B.3. Disasters' Depreciation Rate Distribution with Crop Damages

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.

Appendix C. Estimation

$$\log(d_t^i) = e_t^i \bar{\delta} + \sigma \varepsilon_{i,t}, \quad \varepsilon_{i,t} \sim \mathcal{N}(0,1), \tag{44}$$

where $z_{i,t}$ is a k-dimensional indicator vector such that The indicator vector $z_{i,t}$ allows for a compact representation of the regime-specific parameters such that the realized depreciation rate is drawn from a mixture of k log-normal distributions, each corresponding to a different latent state.

$$z_{i,t} = e_{s_{i,t}} \in \mathbb{R}^k, \tag{45}$$

where $e_{s_{i,t}}$ denotes the canonical unit vector, equal to 1 for the unobserved latent state $s_{i,t}$.

The column vectors $\bar{\delta} = \left\{\bar{\delta}^j\right\}_{j=1}^k$ and $\sigma = \left\{\sigma^j\right\}_{j=1}^k$ denote the regime-specific means and standard deviations, respectively.

Conditional on the latent state, the logarithm of the observed depreciation rate follows a Gaussian distribution:

$$\log(\delta_{i,t}) \,|\, z_{i,t} \sim \mathcal{N}(z_{i,t}'\bar{\delta}, (z_{i,t}'\sigma)^2). \tag{46}$$

In the data, a large fraction of the observed depreciation rates are equal to zero, reflecting years in which a county does not experience a natural disaster. To account for this, I incorporate a two-component mixture model in which the observed depreciation rate is drawn either from a Dirac mass at zero (with probability $1-p$) or from a log-normal distribution (with probability p). The resulting density is given by:

$$f(\delta_{i,t}|z_{i,t};\bar{\delta}_{s_t},\sigma_{s_t},p_{s_t}) = (1-p_{s_t})\,\mathbb{1}_{\{\delta_{i,t}=0\}} + p_{s_t}\,\mathbb{1}_{\{\delta_{i,t}>0\}} \cdot \phi_{LN}(\delta_{i,t};\bar{\delta}_{s_t},\sigma_{s_t}), \tag{47}$$

where $\phi_{LN}(\cdot)$ denotes the log-normal density with parameters determined by the latent regime active at time t.

Estimation of the model parameters—including regime-specific means, variances, transition probabilities, and the mixture weight p—is performed using maximum likelihood, incorporating a Hamilton filter to account for latent regime dynamics. Following Hamilton (1994), the general filter corresponds to

$$\begin{aligned} \hat{z}_{i,t|t} &= \frac{\hat{z}_{i,t|t-1} \odot \eta_{i,t}}{\mathbf{1}_3' \left(\hat{z}_{i,t|t-1} \odot \eta_{i,t}\right)}, \\ \text{with} \quad \hat{z}_{i,t+1|t} &= \Pi \hat{z}_{i,t|t}, \end{aligned} \tag{48}$$

where $\hat{z}_{i,t|t}$ are the filtered probabilities for county i at time t, knowing all information up to time t (i.e., $\underline{\delta}_{i,t} = [\delta_{i,t}, \delta_{i,t-1}, ..., \delta_{i,1}]'$) such that

$$\hat{z}_{i,t|t} = \left\{\mathbb{P}\left(s_{i,t} = j | \underline{\delta}_{i,t}; \bar{\delta}, \sigma, p\right)\right\}_{j=1}^k, \tag{49}$$

and $\eta_{i,t}$, the vector of densities for each state of the economy defined as

$$\eta_{i,t} = \left\{ f(\delta_{i,t}|s_{i,t} = j; \bar{\delta}_j, \sigma_j, p_j) \right\}_{j=1}^{k}. \tag{50}$$

Note that $\mathbf{1}_K$ corresponds to a column vector of ones of size $K \times 1$, and $\odot$ to the Hadamard-product.

For a panel dataset, I define $\boldsymbol{\delta_t} = [\delta_{1,t}, ..., \delta_{N,t}]'$, and $Z_t = \left[z'_{1,t}, ..., z'_{N,t}\right]'$, the conditional joint density function is

$$\boldsymbol{f}(\boldsymbol{\delta}_t|Z_t; \bar{\delta}_{s_t}, \sigma_{s_t}, p_{s_t}) = (1 - p_{s_t})\, \mathbb{1}_{\{\boldsymbol{\delta}_t = 0\}} + p_{s_t}\, \mathbb{1}_{\{\boldsymbol{\delta}_t > 0\}} \cdot \phi_{LN}(\boldsymbol{\delta}_t; \bar{\delta}_{s_t}, \sigma_{s_t}), \tag{51}$$

where the log of the positive depreciation rates follow the joint Gaussian distribution

$$\log\left(\boldsymbol{\delta_t}\right)|Z_t \sim \mathcal{N}_N\left(Z_t\bar{\delta}, \boldsymbol{D_t}\right), \tag{52}$$

where $\boldsymbol{D_t} = \text{diag}\left((Z_t\sigma)^2\right)$ is an $N \times N$ diagonal variance-covariance matrix.

With $H_t = [\eta_{1,t}, ..., \eta_{N,t}]'$, the filter of the panel mixture model becomes

$$\begin{aligned} \hat{Z}_{t|t} &= \left[\hat{Z}_{t|t-1} \odot H_t\right] \oslash \left[\left(\hat{Z}_{t|t-1} \odot H_t\right) \mathbf{1}_{3\times 3}\right], \\ \text{with} \quad \hat{Z}_{t+1|t} &= \hat{Z}_{t|t}\Pi' \end{aligned} \tag{53}$$

The loglikelihood is therefore[29]

$$\sum_t \mathbf{1}'_N \log\left(\left(\hat{Z}_{t|t-1} \odot H_t\right) \times \mathbf{1}_3\right) \tag{54}$$

I apply a numerical maximum likelihood estimation on the vector

$$\theta = \left[\bar{\delta}, \sigma, \{\pi_{jl}\}_{j,l=1}^{k-1}, \{p_j\}_{j=1}^{k}\right],$$

using "Nelder-Mead" and "nlminb" optimization algorithm.

Appendix D. Results

Table D.1. Aggregate and Distributional Statistics

	K	r	w	Mean Wealth	Gini	Mass at 0 (%)	Mean C
M1	3.2500	0.1500	0.9900	5.6900	0.6100	3.9900	1.2600
M2	3.2200	0.1500	0.9900	3.7800	0.8900	29.1600	1.1600
M3	3.2600	0.1500	0.9900	5.3600	0.5200	1.4100	1.2300
M4	3.2200	0.1500	0.9900	3.8300	0.7500	12.3200	1.1600

Notes: The table reports aggregate capital, the Gini coefficient, and the mass of households at the borrowing constraint in the stationary equilibrium. The Gini coefficient is computed from the invariant wealth distribution $g(a)$ using the discrete formula $G = (N^2\bar{a})^{-1} \sum_{i=1}^{N} (2i - N - 1) a_{(i)}$, where $a_{(i)}$ denotes sorted wealth and $\bar{a}$ its mean.

Wealth Distribution. Figure D.1 displays cumulative wealth distributions. Adding depreciation risk shifts the CDF sharply upward near zero wealth, indicating a substantial increase in the fraction of households with negligible assets. This effect is strongest under

[29] The log is taken element-wise.

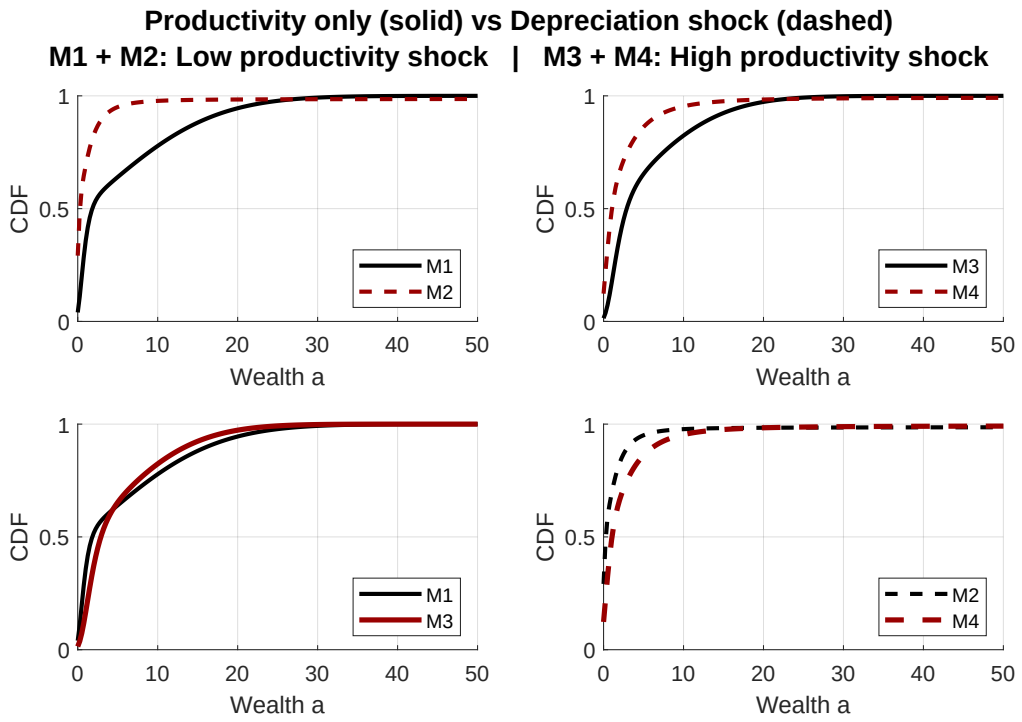


FIGURE D.1. Identical Realization, Different Risk Distribution: Low-State Comparison

Notes: The figure compares Zone 1 and Zone 2 in the low productivity state within Model M2. Since the realized productivity shock is identical across zones, differences in consumption, saving, and the stationary wealth distribution are driven solely by cross-zone differences in the distribution of idiosyncratic depreciation risk. The intensive climate zone exhibits weaker precautionary saving incentives and a larger mass of households near the borrowing constraint.

low productivity dispersion (M2), where the distribution becomes highly concentrated near the borrowing constraint.

Importantly, aggregate capital changes little across specifications, implying that the increase in inequality is primarily bottom-driven. The distribution becomes polarized because a large mass of households remains hand-to-mouth, not because of a substantial expansion of the upper tail.

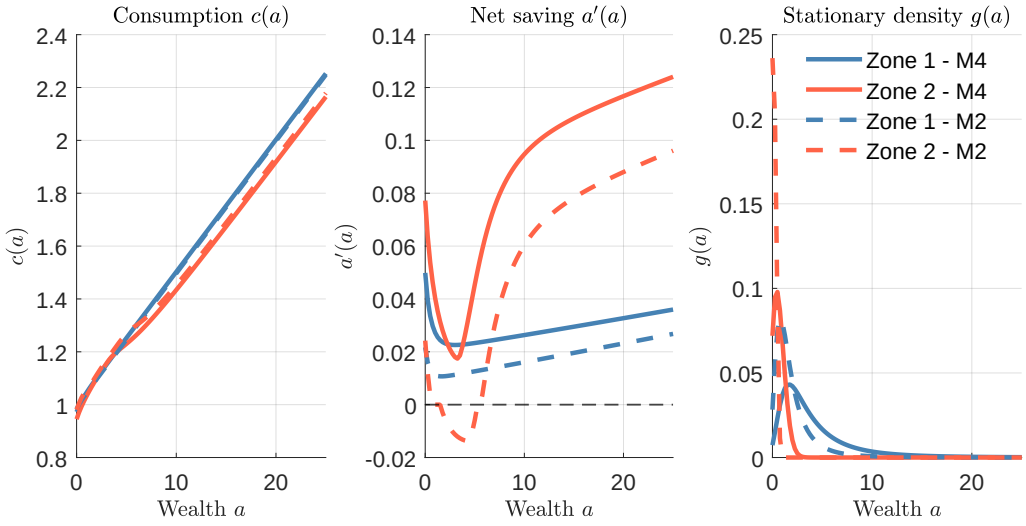


FIGURE D.2. Identical Realization, Different Risk Distribution: Low-State Comparison

Notes: The figure compares Zone 1 and Zone 2 in the low productivity state within Model M2. Since the realized productivity shock is identical across zones, differences in consumption, saving, and the stationary wealth distribution are driven solely by cross-zone differences in the distribution of idiosyncratic depreciation risk. The intensive climate zone exhibits weaker precautionary saving incentives and a larger mass of households near the borrowing constraint.