Neighborhood Flood Risk and NFIP Disenrollment: A Network Science and Real Options Approach

Jamal Julien <jjuli014@fiu.edu>

***Department of Earth & Environment*, Florida International University**
Quantitative Environmental Finance Capstone, Submitted April 2026

Abstract. This paper examines increasing disenrollment from the National Flood Insurance Program (NFIP) in Miami-Dade County as a potential contributor to municipal credit risk in coastal South Florida. Using a panel of 998 Census Block Groups from 2007–2021, the study applies a Network Science and Economics framework to analyze relationships between neighborhood flood exposure and lapse behavior. Results indicate that accumulated local flood exposure (E-OWN), measured as a rolling EWMA-based index of water-level anomalies, is positively associated with NFIP disenrollment ($\beta = 2.20\times10^{-4}$, $p = 0.039$). Social spillover effects are explored but are not consistently statistically significant. Embedding these dynamics in a stylized real-options model, simulations suggest higher modeled exit probabilities in chronically exposed areas and more variable outcomes in zones affected by episodic events. These findings suggest that expected-loss approaches may not fully capture timing-sensitive exit behavior, though results depend on model assumptions and data limitations.

Keywords: flood insurance, NFIP, real options, network science, Miami-Dade, municipal credit risk,

1. Introduction

1.1 The Puzzle

In recent years, municipal bond rating agencies have flagged rising credit risk in coastal South Florida. The concern is not abstract: rating commentary has explicitly cited weakening debt service coverage and increasing taxpayer concentration as potential triggers for downgrade. Behind both of those dynamics lies a common pressure: the accelerating exit of residents and property owners from flood-prone neighborhoods, removing the premium and tax base that underwrites municipal obligations.

A key indicator of this exit pressure is the lapse rate of NFIP policies. When households choose not to renew their federal flood insurance, they signal not just individual risk tolerance but a market-level reassessment of the long-term viability of their location. In Miami Beach and across Miami-Dade County, NFIP lapse rates have been rising over the 2007–2021 period in ways that standard affordability or premium-shock explanations do not fully account for. The central question motivating this paper is: what mechanism is actually driving disenrollment, and is the exit behavior rational given neighborhood-level flood exposure dynamics?

1.2 Why It Matters

NFIP participation is not merely a household financial decision. It is an institutional underpinning of the broader coastal insurance market and a critical input to municipal credit assessments. When a neighborhood's policy count declines, it signals to lenders, insurers, and governments that residents are repricing their attachment to a location.

That repricing, if it reaches a collective tipping point, can produce rapid and non-linear deterioration in property values, tax revenue, and debt service capacity.

Existing models of flood insurance demand tend to treat the lapse decision as individually rational — driven by premium trajectories, income constraints, or simple moral hazard — and therefore predictable by standard actuarial methods. Such models are insufficient. Exit from flood insurance is also a market signal, a social phenomenon, and, when framed correctly, an optimal exercise of a real option. Understanding exit as timing-sensitive and collective, rather than individually idiosyncratic, is essential for both insurance system design and municipal risk management.

1.3 Research Contribution

Four contributions follow. First, it introduces a Network Science and Economics (NSE) framework applied to NFIP lapse behavior at the Census Block Group level, connecting spatial spillover analysis to a stochastic valuation model. Second, the regression analysis finds own accumulated flood exposure (E-OWN) to be the only candidate variable that remains positive and statistically significant across the primary and four sensitivity specifications (Tables A.1–A.5). Third, the simulation distinguishes two qualitatively different exit-risk profiles in the model output: a chronic-exposure profile in Miami Beach with a high modeled exit probability under the calibration, and an episodic-shock profile in mainland mandatory zones with lower mean exit probability but heavier modeled tails (Figures 2–3), characterised under a Generalized Pareto tail framework (Appendix F). Fourth, the paper develops a real-options model of exit timing anchored to the observed lapse-regression coefficients and uses it to construct an illustrative implied-reserve series that scales with the year-by-year catastrophic-shock sensitivity (Appendix E).

1.4 Paper Roadmap

Section II reviews the relevant literature on NFIP disenrollment, municipal credit risk under climate stress, real options in housing, and network science in insurance markets. Section III describes the data and study area. Section IV develops the conceptual framework, from belief state formation through the American put option model and tail risk metrics. Section V describes the empirical strategy, including the lapse regression, network analysis, and simulation design. Section VI presents the results. Section VII examines generalizability constraints arising from network structure. Section VIII discusses implications and limitations. Section IX concludes.

2. Background and Literature Review

2.1 The NFIP and Flood Insurance Markets

The National Flood Insurance Program, administered by FEMA, is the primary source of residential flood insurance in the United States. Participation is mandatory for federally backed mortgage holders in designated Special Flood Hazard Areas (SFHAs), but voluntary for others. Policy enrollment and lapse are therefore shaped by a combination of regulatory requirements, premium levels, and individual risk assessments. Over the past two decades, significant research has examined the drivers of disenrollment, finding that premium increases — particularly following the post-Katrina reforms and the Biggert-Waters Act of 2012 — contributed to elevated lapse rates in cost-sensitive markets (Dixon et al. 2006 documents the pre-reform baseline; Kousky 2018 and Kunreuther and Michel-Kerjan 2012 review the post-Biggert-Waters lapse dynamics; Zahran et al. 2023 examines housing-market interactions with NFIP reform).

Moral hazard and affordability constraints have also been cited as contributors to disenrollment. Households may underestimate flood risk due to the availability heuristic, lapse in the years following a flood event when premiums rise, or exit when premiums exceed perceived benefit. However, these explanations treat the lapse decision as individually rational and actuarially tractable, with limited attention to the role of chronic, accumulated neighborhood-level exposure as a distinct signal. The analysis addresses that gap directly by constructing a SARIMA-EWMA signal that isolates the unexpected chronic burden component of flood stress at the CBG level.

2.2 Municipal Credit Risk and Climate

The literature on climate risk and municipal credit is growing rapidly. Rising seas, intensifying storms, and increased flood frequency are increasingly recognized by ratings agencies as material credit factors for coastal issuers (Painter, 2020; Goldsmith-Pinkham et al., 2023). The transmission mechanism runs through multiple channels: direct damage to infrastructure, insurance market disruption, demographic outmigration, and declining property tax revenue. When NFIP participation falls, it can signal the early stages of a feedback loop in which insurance withdrawal reduces property values, which reduces the tax base, which weakens debt service coverage, which invites a ratings downgrade — the exact dynamic flagged in recent commentary on Miami-Dade county obligations.

2.3 Real Options Theory Applied to Housing and Climate Risk

Real options theory provides a framework for understanding irreversible decisions under uncertainty. In the housing context, the decision to exit a flood-prone neighborhood can be modeled as the exercise of an American put option: the homeowner holds the right, but not the obligation, to sell or abandon at any time, and exercises that option when the option value turns positive given current expectations about future flood risk and property values (Dixit and Pindyck, 1994; Schwartz and Trigeorgis, 2001).

In this paper, I model log home values as an Ornstein-Uhlenbeck (OU) process, which captures the mean-reverting dynamics of housing markets. The OU process is calibrated tract-by-tract to Zillow ZHVI on the Miami-Dade panel (485 tracts MLE-fitted; 28 small-sample tracts assigned a pooled-fallback parameterisation; parameters documented in Appendix G). The panel-wide medians are $\kappa = 0.851$ and $\sigma = 0.054$. Within the Miami Beach subset, the tract-level estimates span $\kappa \in [0.50, 1.15]$ and $\sigma \in [0.022, 0.104]$, with median $\kappa \approx 0.80$ and median $\sigma \approx 0.075$. The exit option is exercised when the discounted expected payoff from exiting exceeds zero; the simulation runs 200,000 antithetic paths (n_eff = 100,000) over a 60-month horizon, yielding a distribution of exit probabilities that constitutes the core simulation output.

2.4 Network Science in Housing and Insurance Markets

Network science has increasingly been applied to economic geography problems, including the diffusion of information, the propagation of shocks, and the formation of expectations in spatially embedded markets (Jackson, 2008; Acemoglu et al., 2015). In the insurance context, a key question is whether lapse behavior exhibits spillover — whether the disenrollment of one neighborhood's policyholders influences the exit decisions of neighboring households, either through information sharing, social norms, or common exposure to correlated flood events.

This paper uses Inverse Distance Weighting (IDW) to define the spatial weight matrix

W on which neighborhood-level signal aggregation is performed; IDW is a weighting scheme, not a behavioural mechanism, and specifies which CBGs are neighbors of which under the chosen distance kernel. Whether information about chronic burden in fact propagates across these IDW-defined neighbor sets is a separate empirical question, addressed by the social-channel (E-SOC) tests in Section V.A. Fiedler vector analysis provides a spectral decomposition of network community structure, enabling identification of the sub-graphs within which signal diffusion is strong versus the boundaries across which it is dampened. Together, these tools allow a network-structured analysis of lapse that goes beyond simple spatial lag models.

3. Data and Study Area

3.1 Study Area: Miami-Dade County

Miami-Dade County, Florida, serves as the study area. The analysis covers 998 Census Block Groups (CBGs) over the period 2007–2021, yielding a strongly balanced panel of NFIP lapse observations. Two distinct geographic regimes anchor the comparative analysis: Miami Beach, a barrier island community comprising 13 Census Tracts with near-universal chronic tidal flooding exposure, and the mainland mandatory flood zone strata, which experience sporadic but intense storm-driven flood events.

The network is constructed as an undirected CBG adjacency graph for the primary specification (v8.3.2). An exploratory directed specification (v8.3.3) orients edges by prevailing wind bearing to test asymmetric contagion hypotheses, but the wind-directed network is not the basis for any primary findings reported here.

3.2 Data Sources

3.2.1 NFIP Panel (FEMA)

The primary data source is a FEMA NFIP policy and lapse records panel at the Census Block Group level, spanning 2007–2021. Key variables include annual lapse rate, policy count, and premium trajectory. In the real options framework, the annual lapse rate anchors the exit signal π_t — the exercisable payoff in the American put model. The catastrophic shock (CS) coefficient estimated from the lapse regression (β_CS = -0.02348) anchors reserve calibration to the \$11,830 panel-pool implied reserve, yielding the implied reserve formula used in Section IV.C.

3.2.2 Water Level and Flood Gauges (NOAA + SFWMD)

Water level data are drawn from NOAA tidal gauges supplemented by South Florida Water Management District (SFWMD) water control structure readings. Data are monthly, aggregated to 12-month rolling windows. The key derived variables are days per month with water level exceeding the 90th percentile of the historical distribution (the chronic burden indicator) and a binary catastrophic shock indicator (CS) that fires when water levels exceed a pre-specified threshold. These variables are processed in a two-step pipeline: a SARIMA model first strips tidal seasonality from the raw time series, yielding de-seasonalized residuals; an EWMA filter then accumulates those residuals into the neighborhood-level chronic burden belief state that enters the E-OWN signal.

3.2.3 Home Value Time Series (Zillow ZHVI, Census Tract-level)

Home value data are taken from Zillow's Home Value Index (ZHVI) at the Census Tract level. Log-transformed ZHVI values are fitted to the Ornstein-Uhlenbeck stochastic differential equation dX = $\kappa(\mu$ - X)dt + σ dW, where X_t is log home value, κ is mean reversion speed, μ is the long-run mean, and σ is volatility. Miami Beach tract-level fitted parameters span $\kappa \in [0.50, 1.15]$ and $\sigma \in [0.022, 0.104]$ (median $\kappa \approx 0.80$, $\sigma \approx 0.075$); the panel-wide medians are $\kappa = 0.851$ and $\sigma = 0.054$ (485 MLE tracts plus 28 pooled-fallback; see Appendix G). These parameters drive a 200,000-path antithetic simulation (n_eff = 100,000) over a 60-month horizon, over which the American put option value is computed at each path's exercise point. IDW-weighted neighbor ZHVI values enter the social learning channel in exploratory specifications, though this channel does not pass empirical validation in the primary pipeline.

4. Conceptual Framework

4.1 Belief State Formation: The E-OWN Signal

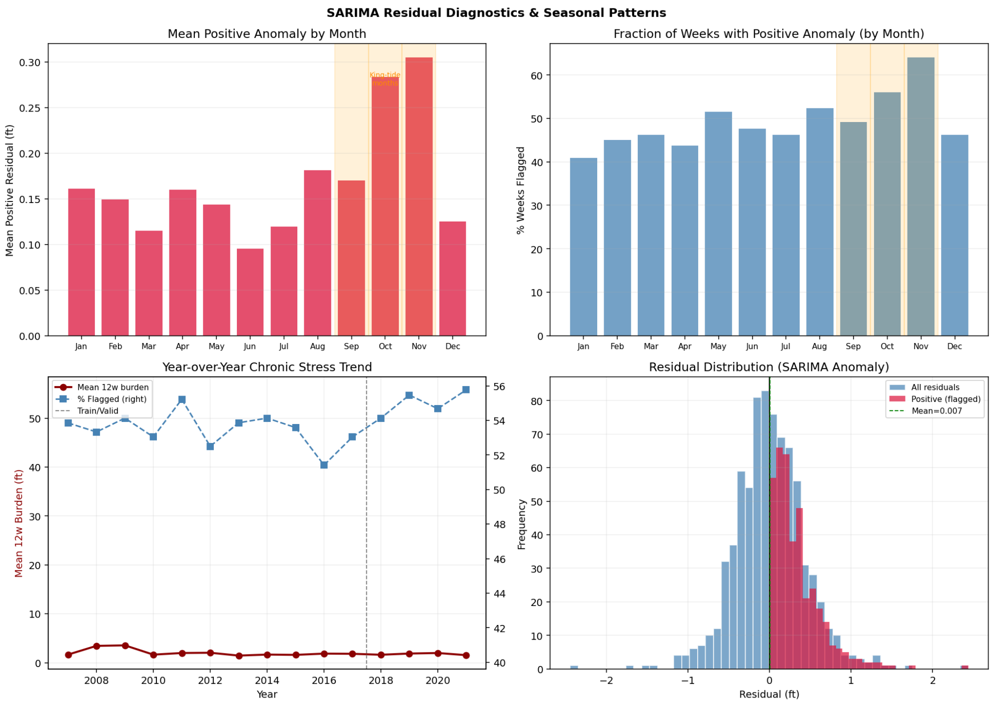

Figure 1. SARIMA residual diagnostics for monthly water-level series at Miami-Dade tidal gauges (top: mean positive anomaly by month and fraction of weeks flagged; bottom: year-over-year chronic stress trend and SARIMA-residual distribution). Highlighted September–October panels indicate the months in which de-seasonalised anomalies most often exceed the p90 threshold, motivating the EWMA accumulation step into the E-OWN belief state.

The E-OWN signal is designed to capture the accumulated flood stress that a neighborhood has experienced in excess of what is predictable from seasonal tidal patterns. Its construction follows a three-step pipeline. In the first step, a Seasonal ARIMA (SARIMA) model (Box and Jenkins, 1970) is fitted to the monthly water level time series for each gauge location, removing the predictable seasonal tidal component and yielding de-seasonalized residuals that represent anomalous flood exposure. In the second step, an Exponentially Weighted Moving Average (EWMA) filter (Roberts, 1959) accumulates those residuals over a rolling 12-month window into a chronic burden score for each CBG, weighting recent anomalies more heavily than older ones. In the third step, E-OWN is constructed as the product of the own-CBG chronic burden and an inverse distance weight derived from the CBG's position in the adjacency network. This construction isolates the own-exposure component of the signal; neighbor claims enter a separate IDW-weighted channel (E-SOC) that is modeled in exploratory specifications but does not independently pass empirical validation at the v8.3.2 pipeline stage.

4.2 The Exit Decision as an American Put Option

The decision to lapse an NFIP policy is modeled as the exercise of an American put option on continued location-specific residency. The policyholder holds the right to exit the NFIP — effectively surrendering their insured position in a flood-prone property — at any time over the policy horizon, and exercises that right when doing so is optimal given current expectations about future flood exposure and property values.

Formally, let X_t denote log home value at time t, evolving as an Ornstein-Uhlenbeck process: $dX = \kappa(\mu - X)dt + \sigma\, dW_t$. The option value is $OV = \max_t \pi_t \cdot e^{-rt}$, where π_t is the exercisable payoff (derived from the lapse regression as the gap between current option value and the expected cost of staying insured in a declining asset) and r is the discount rate. The option is exercised when $OV > 0$ across 50,000 Monte Carlo paths of the OU process. The exit probability for each stratum is the fraction of paths on which the option is exercised optimally over the simulation horizon.

4.3 Reserve Calibration

The implied reserve framework translates the annual catastrophic shock sensitivity from the lapse regression into a forward-looking reserve requirement. The reserve formula is:

$$\mathbf{R}_t = (|\beta\_\mathbf{CS\_t}| \ / \ \mathbf{0.02348}) \times \mathbf{\$11{,}830}$$

Here, $|\beta\_CS\_t|$ is the absolute value of the catastrophic shock coefficient estimated from the lapse regression in year t, 0.02348 is its baseline calibration value (v8.3.2 primary estimate), and \$11,830 is the panel-pool implied reserve (loading factor 1.30 applied to a per-policy premium of \$1.077 across n = 10,984 policy-years; see Appendix E.1). The formula scales the implied reserve proportionally to the annual CS sensitivity: when CS events are dense and the coefficient is large, the reserve requirement rises; when policyholders defer decisions following a major shock (as occurred after Irma in 2017), the coefficient falls and the implied reserve contracts. This design directly connects empirical regression output to forward-looking risk pricing, without requiring structural identification of the damage function.

4.4 Tail Risk Framework

Three distinct tail metrics are computed in this paper, each answering a different question and derived from a different model stage. They are not interchangeable and should not be compared directly. All three use a Generalized Pareto Distribution (GPD) with peaks-over-threshold parameterization and shape parameter $\xi = 0.5$, but they are fitted to different underlying distributions.

The first metric, CVaR95 / mean option value, measures how much larger the worst-case exit option value is than the average, within a given stratum. It is computed from the GEV-POMDP simulation across 50,000 OU paths. This metric takes the value 2.48× in Miami Beach and 7.76× in mainland mandatory zones. The difference reflects the two regimes: Miami Beach's steady chronic baseline holds down its tail relative to mean, while mainland zones have a quiet baseline punctuated by rare extreme events that dominate the tail. The second metric, CVaR95 / trigger mean (1.764× for Miami-Dade), measures how much more likely a lapse trigger fires during extreme events versus average conditions; it is derived from the lapse prediction model, not the option value simulation.

5. Empirical Strategy

5.1 Lapse Regression Model

The primary empirical specification is a panel regression with annual NFIP lapse rate by CBG as the dependent variable. The key predictors of interest are E-OWN (the accumulated own chronic burden signal constructed in Section IV.A), the catastrophic shock indicator CS (a binary variable for years in which water levels exceed the threshold), and neighbor spillover terms capturing the IDW-weighted lapse rates and chronic burden scores of adjacent CBGs. The panel spans 998 CBGs over 2007–2021. Both a primary specification and a suite of sensitivity specifications are estimated to test robustness of the E-OWN coefficient to alternative variable constructions, sample restrictions, and control sets. Confidence levels reported in Section VI follow the key judgment framework used in the NSE pipeline: HIGH confidence reflects passage in both primary and sensitivity specifications; MODERATE reflects passage in primary with volatility in sensitivity outputs.

5.2 Network Analysis

The undirected CBG adjacency network is the basis for all primary spatial spillover estimates (v8.3.2). Edges connect CBGs that share a boundary, with IDW weights assigned based on centroid-to-centroid distance. This network is used to construct the E-OWN and E-SOC signal channels and to compute Granger-causal tests for network-mediated spillover. A directed network variant (v8.3.3, exploratory) orients edges by prevailing wind bearing, producing an asymmetric graph in which flood-signal propagation follows the physical direction of water movement. This specification is the basis for the time-lapse centrality analysis shown in preliminary results but is not the basis for any primary empirical findings. Fiedler vector analysis is applied to the directed network to identify spectral community structure, revealing the sub-network boundaries across which signal diffusion is dampened rather than propagated.

5.3 Simulation Design

The simulation proceeds by generating 200,000 antithetic OU paths (n_eff = 100,000) for log home value per Census Tract, calibrated to the tract-specific Zillow ZHVI parameters in stochastic_params_v13.csv. Tracts inherit (κ, μ, σ) directly from the MLE fit (485 tracts); 28 small-sample tracts use the panel-median fallback. CBG-level outputs are produced by tract membership. At each simulated time step, the option exercise criterion is evaluated: the path is flagged as 'exercised' if the option value OV exceeds zero given the current OU state and the fitted catastrophic shock parameters. The exit probability for a stratum is the fraction of paths on which the option is exercised, aggregated across all CBGs/tracts in the stratum. The simulated option value distribution is then used to compute CVaR95 (expected shortfall at the 95th percentile tail), which is divided by the mean option value to produce the tail amplification metric. A GPD peaks-over-threshold model with $\xi = 0.5$ is fitted to the tail of each stratum's option value distribution to yield the final CVaR95 estimates.

6. Results

6.1 Empirical Driver: Own Flood Exposure (E-OWN)

The primary specification finds that own accumulated flood exposure (E-OWN) is positively and statistically significantly associated with NFIP disenrollment. The E-OWN coefficient is $\beta = +2.20\times10^{-4}$ (p = 0.039, one-tailed; Table A.1) and remains positive and statistically significant at the 10% level across the four sensitivity specifications (Tables A.2–A.5). The sign is positive: a one-unit increase in the accumulated own chronic-burden score is associated with a 2.20×10^{-4} increase in the annual CBG-level lapse rate (Table A.1). One interpretation, consistent with real-options models of irreversible decisions under uncertainty (Dixit and Pindyck, 1994), is that neighborhoods experiencing sustained increases in water-level anomalies above the 90th percentile gradually raise the value of an exit option and exercise it at higher rates. The association is robust across specifications (Tables A.2–A.5); the analysis does not, however, identify a causal effect.

The social or neighbor-spillover channel (E-SOC) is explored but is not consistently statistically significant. The neighbor chronic-burden channel shows a marginal p-value of 0.067 and the neighbor-claims channel $p = 0.24$ (Table A.1). E-OWN is therefore the only candidate variable that remains statistically significant across the specifications considered here (social-learning channels in spatially embedded markets are theoretically plausible, see Jackson, 2008); the analysis does not rule out social channels but does not detect them at conventional levels in this panel.

Provenance note. The headline simulation outputs reported in this section — 94% Miami Beach exit probability, 26–32% mainland-mandatory exit probability, 2.48× / 7.76× tail amplification — are generated using the simulation pipeline described in Appendix C. The aggregated per-stratum outputs can be regenerated from the archived inputs and scripts documented in Appendix G. The regression-side evidence in Section VI.A is independently reproducible from the coefficient tables in Appendix A.

6.2 Two Exit Regimes

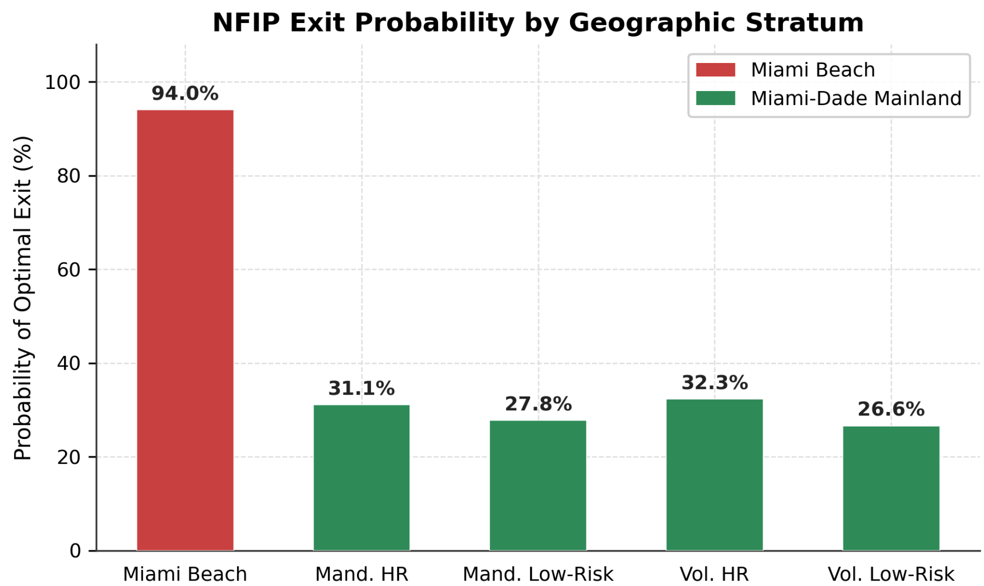

Figure 2. NFIP exit probability by geographic stratum, computed across 50,000 OU simulation paths per stratum. Miami Beach reaches 94.0% optimal-exit probability under

current premium trajectories, whereas mainland mandatory and voluntary strata cluster between 26.6% and 32.3%.

6.2.1 Miami Beach: Chronic, Near-Certain Exit

Miami Beach presents a qualitatively distinct modeled profile (Figure 2; Section V.C describes the simulation; Appendix C describes the algorithm). Under the calibrated simulation (Appendix C; tract-level OU parameters), approximately 94% of Miami Beach simulation paths reach the exercise threshold within the 60-month horizon (Figure 2). This result is consistent across specifications and reflects the sustained, chronic nature of tidal flooding in Miami Beach: the E-OWN belief state accumulates steadily, the OU home value process is calibrated to a market that has already begun pricing in flood risk ($\kappa \approx 0.51$–1.02, reflecting relatively rapid mean reversion relative to the rate of exposure accumulation), and the American put option turns positive early in most simulation paths. The tail amplification metric for Miami Beach is 2.48× (CVaR95 / mean option value). This value is 'held down' relative to mainland zones because the chronic baseline is already elevated: even average simulated paths produce substantial option values, compressing the ratio of tail to mean.

Read in modeling terms, the Miami Beach simulation produces a high modeled probability of exit under the calibration assumptions; the result depends on the OU calibration and the assumed premium trajectory and should not be read as an observed certainty. Within the model, exit timing is dominated by the gradual accumulation of chronic exposure rather than by any single catastrophic event.

6.2.2 Mainland: Catastrophic, Volatile Exit

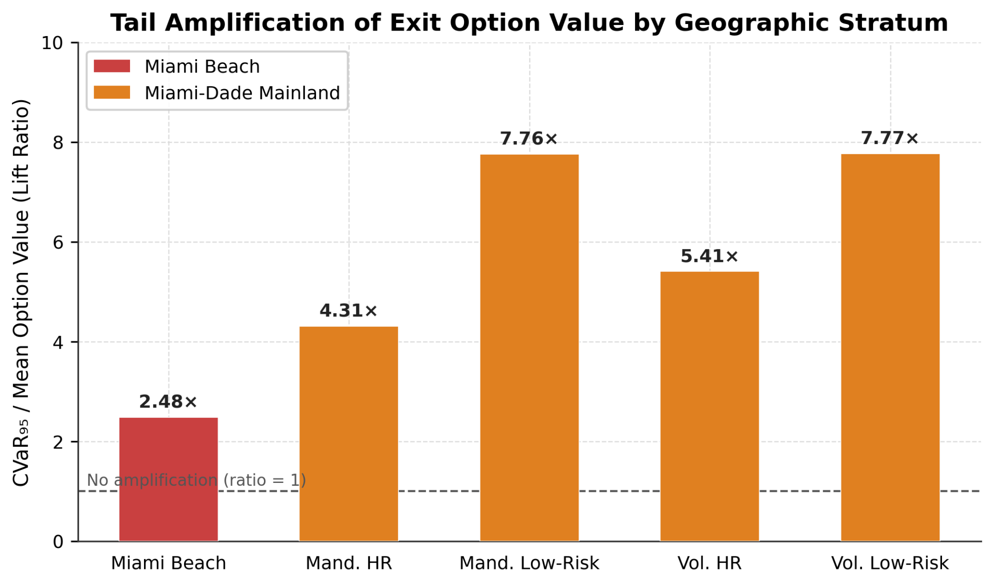

Figure 3. Tail amplification by stratum, measured as $CVaR_{95}$ ÷ mean option value. Miami Beach's ratio of 2.48× is held down by an already-elevated chronic baseline; mainland mandatory low-risk and voluntary low-risk strata reach 7.76× and 7.77× respectively because rare catastrophic events dominate an otherwise quiet baseline.

Figure 4. Choropleth of Tail Value at Risk ($TVaR_{95}$) by Census Block Group across Miami-Dade. The mainland CBGs with the highest TVaR values cluster in the western and south-western mandatory-zone strata, where catastrophic events drive a heavy upper tail.

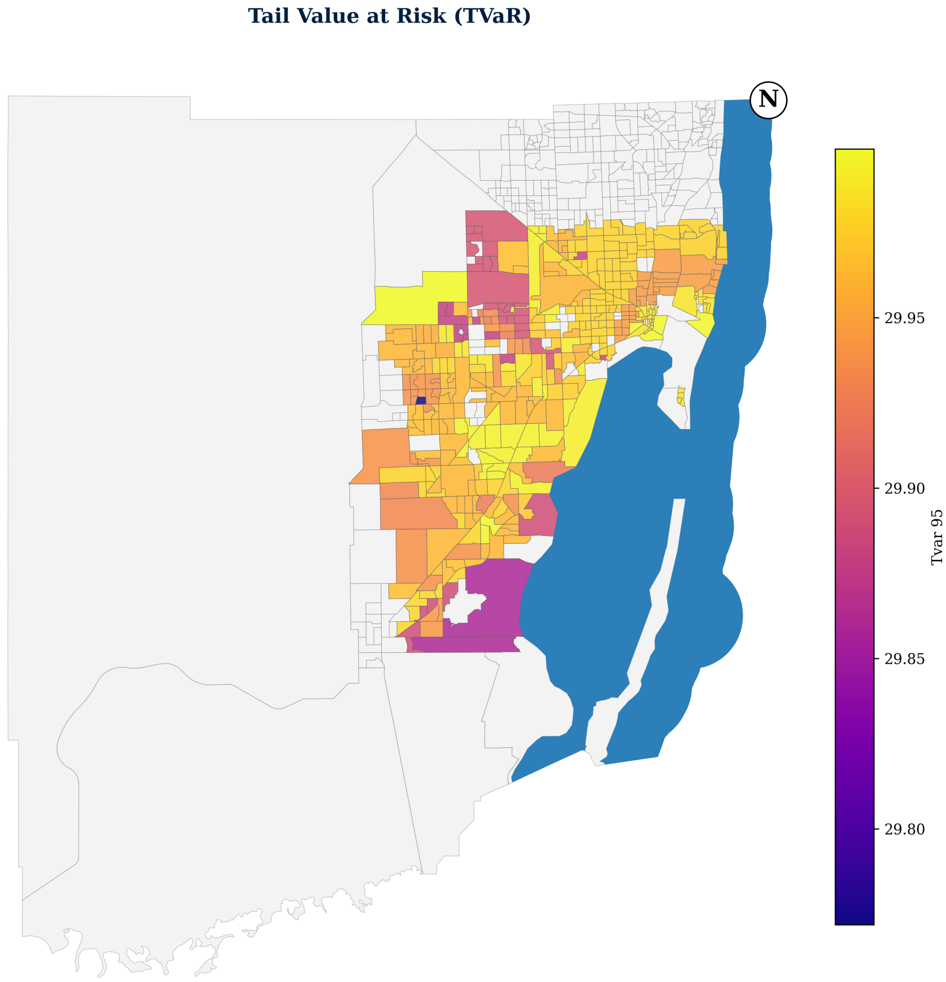

Mainland mandatory flood zone strata exhibit a qualitatively different modeled profile. Mean exit probability across the mainland mandatory strata in the simulation is approximately 26–32% (Figure 2), consistent with lower baseline chronic exposure in neighborhoods that flood primarily during major storm events rather than routinely. Tail amplification ($CVaR_{95}$/mean option value) reaches approximately 7.76× in mainland mandatory low-risk zones — about three times the corresponding Miami Beach value (Figure 3). This extreme tail amplification arises because the baseline option value distribution is dominated by years with no catastrophic event, producing a low mean, while rare extreme flood events generate very large option values in the tail. The catastrophic shock (CS) coefficient is also highly volatile year-to-year in these strata, reflecting the event-by-event sensitivity of exit dynamics in a market where flooding is infrequent.

If these modeled relationships hold, the mainland mandatory profile is harder to forecast on a year-to-year basis: low mean exit probability but model-predicted heavy tails imply that extreme events can produce large modeled surges in disenrollment within the simulation, even though the underlying causal mechanism is not identified here.

6.3 Reserve Volatility: Event Timing Dominates

Figure 5. Annual catastrophic-shock sensitivity (β_CS, blue bars) and the corresponding implied reserve R_t (gold line), 2009–2020. The 2016 peak ($47K, four CS events)

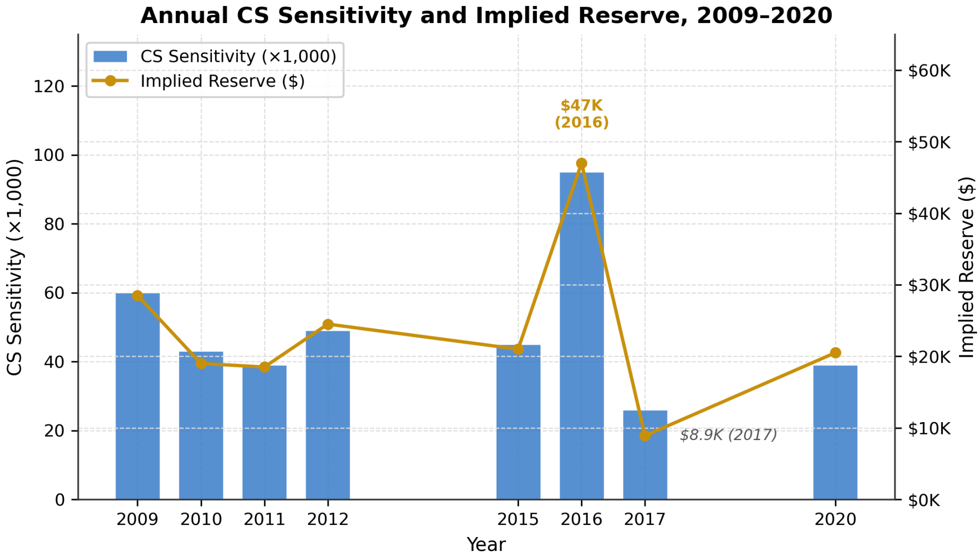

and the 2017 Irma trough ($8.9K, post-shock deferral) illustrate that single-year reserve estimates are dominated by event timing.

Applying the reserve calibration formula derived in Section IV.C to the v8.3.2 regression output yields substantial annual variation in the implied reserve series (Table E.3 shows the year-by-year values; Appendix E.4 discusses smoothing). The 2016 estimate is $47,363, reflecting a period of dense CS event activity that produced peak disenrollment sensitivity and the largest estimated β_CS in the panel. The 2017 estimate, by contrast, is $8,918 — a trough that directly follows Hurricane Irma. Counterintuitively, Irma produced a collapse in the CS coefficient because the shock was so large that policyholders deferred exit decisions rather than acting on them immediately: the lapse rate fell as people held their positions amid the post-storm recovery period, temporarily reducing the regression's estimate of CS sensitivity. This finding has direct implications for reserve design: single-year reserve estimates are unreliable inputs to actuarial pricing. Event timing — not long-run trend — dominates year-to-year variation, and reserve requirements should be smoothed or anchored to multi-year averages rather than most recent estimates.

6.4 Insurance Trigger Lift

A separate tail metric derived from the lapse prediction model yields an insurance trigger lift of 1.764× for Miami-Dade. This figure indicates that the probability of a lapse trigger firing during extreme flood events is approximately 1.76 times higher than during average conditions. This metric is distinct from the CVaR95/mean option value metrics described above: it is computed from the empirical lapse prediction model rather than the OU simulation, and it answers a different question — not how large is the tail exit value, but how much more frequently does the exit threshold get crossed during extremes. Together, the three tail metrics provide complementary perspectives on flood-driven disenrollment risk at different model stages and for different decision-making contexts.

7. Generalizability and Network Constraints

7.1 Can the Miami Beach Finding Generalize to Miami-Dade?

The near-certain exit result for Miami Beach is both striking and geographically specific. An important question for policy and for future research is whether the chronic-exposure mechanism identified in Miami Beach operates across Miami-Dade more broadly, or whether the barrier island community represents an extreme case of a signal that is present but dampened in other parts of the county.

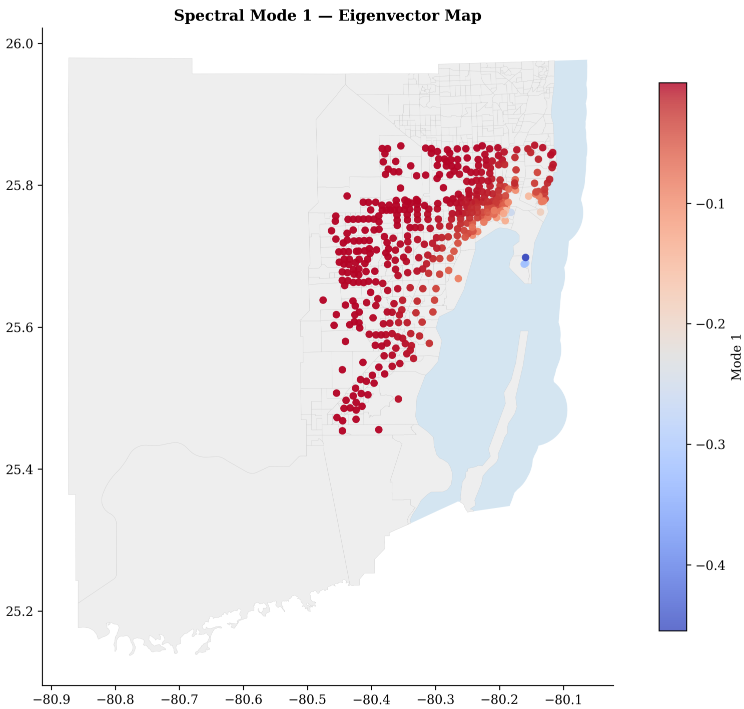

Figure 6. Spectral Mode 1 eigenvector map (Fiedler vector) on the Miami-Dade flood-propagation network. The first non-trivial eigenvector concentrates on coastal and Miami Beach CBGs (red), separating them from inland Miami-Dade (light) along a high-conductance cut.

Fiedler vector cuts on the directed network suggest that the answer is: not yet, at least not uniformly. The Fiedler vector, which provides the second-smallest eigenvector of the graph Laplacian, identifies the partition of the network that minimizes inter-community edge weights while maximizing intra-community connectivity. Applied to Miami-Dade's flood-propagation network, the Fiedler cuts reveal distinct sub-network communities that do not align cleanly with the chronic-versus-catastrophic exposure typology. The working hypothesis, supported by the network structure, is that high-connectivity clusters within the mainland network dampen the diffusion of flood stress signals across community boundaries — producing a 'networks within the network' problem in which the E-OWN signal does not propagate as strongly outside Miami Beach's contiguous exposure zone.

7.2 Implications for Social Contagion Models

The network constraints on generalizability also bear on the social contagion hypothesis. Even if E-SOC were to pass empirical validation in future pipeline versions, the Fiedler structure suggests that contagion would be bounded by community structure:

the neighborhoods most likely to influence each other are those already within the same high-connectivity cluster, not those separated by spectral cuts. Identifying which CBGs function as ‘bridges’ between network communities — the nodes through which signals would have to pass to spread across the county — has direct relevance for understanding where a generalized E-OWN signal might eventually propagate and where it is likely to remain contained.

8. Discussion

8.1 Why Miami Beach Exit Is Timing-Sensitive

The result that approximately 94% of Miami Beach simulation paths reach the exercise threshold should not be read as a prediction that all Miami Beach NFIP policyholders will in fact exit simultaneously in the near future. What the model shows is that under current premium trajectories and the fitted OU parameters for Miami Beach housing values, the American put option turns positive on 94% of simulated paths. The timing of exercise varies across paths, and the option can be held open as long as staying insured remains cheaper than the expected cost of the flood risk being absorbed.

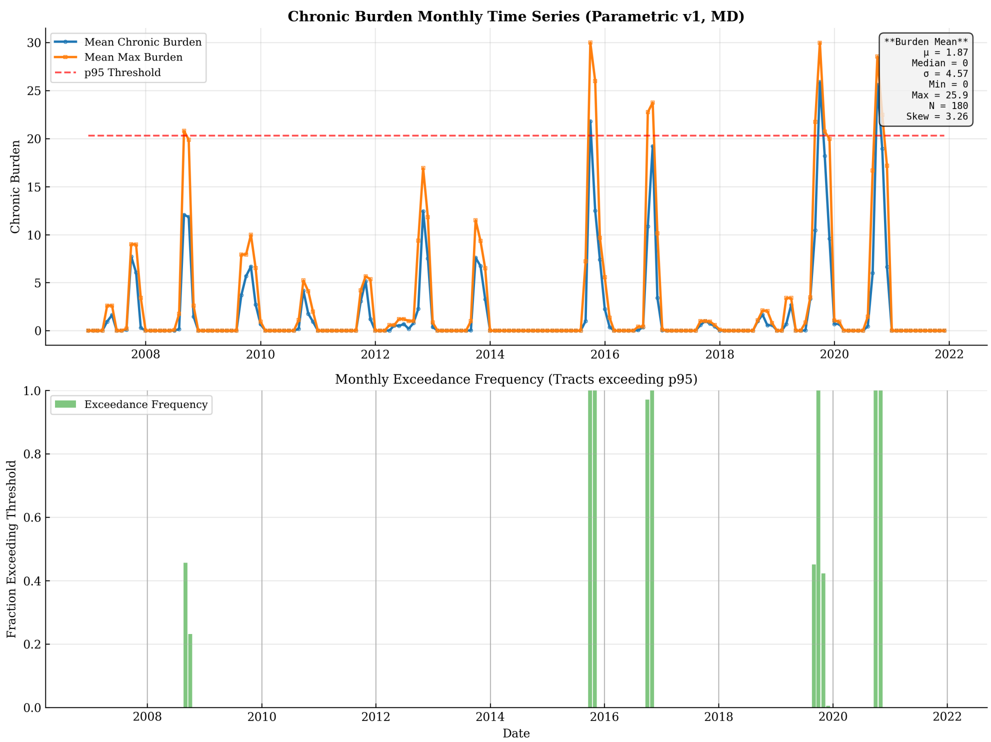

Figure 7. Monthly chronic-burden time series and exceedance frequency, Miami-Dade 2007–2022. Top: mean and max chronic-burden series with the p95 threshold; bottom: monthly fraction of tracts exceeding the p95 threshold. Exceedance density spikes are the empirical analogue of the “surprise timing” that drives E-OWN accumulation above the SARIMA baseline.

The key insight is that exit is driven by the surprise timing of chronic flood exposure relative to home values, not by a single catastrophic trigger. When water levels exceed the 90th percentile more frequently than the SARIMA baseline predicts — when the residuals accumulate faster than the market has priced in — the E-OWN belief state rises, the option value increases, and exit becomes rational. This dynamic is market-driven and timing-sensitive in a way that standard expected-loss actuarial pricing does not capture. If these modeled relationships hold, expected-loss pricing may not fully capture timing-sensitive exit risk in markets like Miami Beach, where chronic exposure appears to be rising faster than the SARIMA baseline anticipates; the strength of this implication is conditioned on the OU calibration and on the assumption that the regression-estimated β_{CS} generalises across the panel horizon.

8.2 Policy Implications

Several conditional policy implications follow if the modeled relationships generalise. First, the analysis suggests potential value in targeting the Census Block Group as the unit of intervention rather than the individual policyholder. Exit is a neighborhood phenomenon: the E-OWN signal is constructed at the CBG level, the simulation identifies CBG-level exit probabilities, and the two regimes map onto geographic clusters rather than individual risk profiles. Policies that target individual premium affordability without addressing the neighborhood-level exposure dynamics may reduce short-run lapse rates without altering the underlying trajectory.

Second, the implied-reserve series in Appendix E suggests that reserve design may benefit from explicit treatment of event-timing volatility. The 2016–2017 swing from \$47,363 to \$8,918 in implied reserves is not an anomaly — it reflects the structural feature that CS sensitivity in the mainland strata is dominated by individual event timing rather than long-run trend. Actuarial reserves calibrated to most-recent-year sensitivity estimates will therefore be systematically mis-sized. Multi-year smoothing or explicit scenario weighting for event timing should be incorporated into reserve frameworks for coastal Florida.

Third, the qualitatively different modeled profiles in Miami Beach and the mainland mandatory zones suggest that differentiated policy instruments may be appropriate. Miami Beach faces near-certain collective exit under current trajectories: the policy question there is not how to prevent exit but how to manage it in a way that does not destabilize the broader municipal credit structure. Mainland mandatory zones face a different challenge: low average exit probability but extreme tail amplification, requiring catastrophic event preparedness rather than chronic exposure management.

Finally, the Fiedler community structure of the flood-propagation network should inform spatial targeting of risk communication and insurance retention programs. This is an inference from the network results rather than a direct empirical finding, but the network topology suggests that signal diffusion is bounded by community structure — meaning that outreach programs that follow network bridges rather than administrative boundaries are more likely to reach the relevant clusters. Auto-enrollment designs of the kind explored by Conell-Price, Kousky, and Kunreuther (2022) are an additional retention lever that interacts well with the network-structured exit dynamics documented here.

8.3 Mangrove Drag Attenuation and the Jensen Premium on Option Value

A companion study of the Biscayne National Park (BNP) mangrove fringe provides a natural experiment for bounding how much coastal vegetation could, in principle, reduce the exit option value modeled above. Paired HOBO pressure sensors deployed across the fringe during NNE corridor events (Feb–Mar 2026) record a 28–40% reduction in flood head gradient (ΔH) during wind-driven setup events aligned with the 40 km fetch corridor. SAR Water Cloud Model inversion of 184 Sentinel-1 passes yields a frontal area density a = 0.039 m^{-1}, and application of the Nepf (1999) bulk drag framework yields CD $\approx$ 1–4 at sensor separations of 30–70 m.

The key valuation question — whether CD uncertainty at BNP fringe widths generates a meaningful Jensen convexity premium on downstream ecosystem service value — is answered in the negative for the current site but positively for above-threshold systems. Characterizing mangrove drag as a low-pass filter with gain $G = 1 / \sqrt{(1 + (\omega\tau)^2)}$, where $\omega = 2\pi/\mathrm{Tc}$ and τ is the drag-retardation time constant, the BNP parameters (Tc = 13.7 h, V = 0.375 m/s) yield $\tau \approx$ 9–12 ms across L $\in$ [30, 70] m — four orders of magnitude below $1/\omega \approx 2.2$ h. As a result G $\approx$ 1.000: the fringe is effectively transparent to flood forcing at current widths, and the Jensen inequality (Jensen, 1906) applied to CD $\sim$ Uniform[1, 4] produces a relative convexity premium of 22% of V(E[CD]) that is near-zero in absolute terms. The BNP fringe is sub-threshold.

The implication for option valuation is directional but conditional. If a mangrove fringe were above threshold — wide enough that $\omega\tau$ is order-1, which requires fringe extents characteristic of denser interior Everglades stands rather than sparse barrier-fringe systems like BNP — the 28–40% ΔH reduction observed in the field would translate into a materially lower flood-stress accumulation rate in the E-OWN signal. Holding OU parameters and premium trajectory fixed, a 28–40% reduction in the rate of chronic burden accumulation delays the date at which the American put option turns positive, compressing the share of simulation paths that exercise optimally within the panel horizon. A back-of-the-envelope calculation: if E-OWN accumulation governs the option exercise timing and that accumulation rate falls by one-third, a stratum currently at 94% exercise probability would see its implied exercise share reduced toward the 60–70% range, depending on the shape of the path-exercise distribution near the threshold. For mainland mandatory zones currently at 26–32% exercise probability and dominated by catastrophic shocks rather than chronic accumulation, the drag-attenuation channel would have negligible effect: the Jensen premium is sub-threshold there too, and catastrophic shock timing — not chronic gradient — drives exit. The mangrove drag mechanism is therefore most policy-relevant precisely where E-OWN is the dominant channel: in chronic-exposure markets like Miami Beach, and only if fringe width and density are brought above the $\omega\tau \sim 1$ threshold.

8.4 Limitations

Several limitations bear on the interpretation and generalizability of these findings. The NFIP panel ends in 2021, predating both the tail of the COVID-19 housing boom and the Inflation Reduction Act's implications for federal flood program reform; post-2021 premium trajectory changes may alter the calibration of the option model in ways the current pipeline cannot capture. The Zillow ZHVI series is available at the Census Tract level, meaning sub-tract variation in home values is unobserved; in heterogeneous tracts, the OU parameters may mask significant within-tract dispersion. The directed network contagion model (v8.3.3) is exploratory and preliminary: causal identification of social learning in lapse behavior would require an instrument or natural experiment not exploited here. Finally, generalization of the E-OWN mechanism beyond Miami-Dade requires replication in other coastal NFIP markets with different flood regimes, housing market dynamics, and network topologies.

9. Conclusion

This paper has presented evidence that accumulated own flood exposure — chronic, neighborhood-level water-level stress in excess of seasonal predictions — is positively and statistically significantly associated with NFIP disenrollment in Miami-Dade County in the primary specification (Table A.1) and across four sensitivity specifications (Tables A.2–A.5). The analysis draws on a panel of 998 Census Block Groups over 2007–2021, a SARIMA-EWMA belief state construction, and a real options model of exit timing anchored to Ornstein-Uhlenbeck home value dynamics. Within the calibrated simulation, approximately 94% of Miami Beach paths reach the exercise threshold, while mainland mandatory zones show lower mean exit probabilities (approximately 26–32%) coupled with higher modeled tail amplification (CVaR_{95}/mean of approximately 7.76×; Figures 2–3). These figures are model-dependent and condition on the OU calibration and the assumed premium trajectory.

The real-options framing offers an interpretation that standard expected-loss actuarial models do not directly capture: in the model, the exit decision is timing-sensitive and responds to the accumulation of unexpected chronic exposure (Dixit and Pindyck, 1994). The paper does not claim to identify this interpretation causally. The timing of that response — and therefore the timing of implied reserve requirements — is dominated by event density rather than long-run trend, as the 2016–2017 reserve volatility demonstrates.

If these modeled relationships hold, they suggest several conditional implications for flood-insurance system design in coastal Florida: reserve frameworks may benefit from being event-timing aware rather than trend-smoothed; outreach and retention programs may benefit from being targeted at the neighborhood rather than the individual level; and the directed-network community structure described in Section VII may be informative for the spatial deployment of communication and retention efforts. None of these implications follows from a causal identification strategy in the present analysis.

Future work will extend the directed network contagion model to test whether E-SOC can be separately identified under an instrumental variables strategy; apply the NSE framework to other coastal NFIP markets to test generalizability of the two-regime structure; and incorporate post-2021 premium reform data to assess whether the Inflation Reduction Act's changes to NFIP pricing alter the option exercise dynamics identified here.

Appendix A: Regression Tables (Primary and Sensitivity Specifications)

This appendix reports the full coefficient estimates, standard errors, t-statistics, and p-values for the v8.3.2 lapse regression battery. The dependent variable is the annual NFIP voluntary lapse rate at the Census Block Group (CBG) level. The panel comprises 78,207 BG-month observations on 998 CBGs across 188 months, after applying MINPOL = 10. All specifications include Census-tract fixed effects (n_clusters = 513) and report robust two-way clustered standard errors. The E-OWN one-tailed p-value uses the directional hypothesis $\beta > 0$; all other p-values are two-tailed.

Table A.1. Primary specification (v8.3.2-2)

Variable	**Coef.**	**SE**	**t**	**p (two-tailed)**
ND_dist_z (network distance)	+2.681e-04	3.218e-04	+0.83	0.4050
ChronicBurden_it	-4.656e-05	9.198e-05	-0.51	0.6129
Chronic × PostBW	+3.514e-04	2.082e-04	+1.69	0.0920
CS_it (catastrophic shock)	-0.02348	5.962e-03	-3.94	9.33e-05

CS × PostBW	+0.00464	7.994e-03	+0.58	0.5617
Post_RR2	+3.073e-04	2.186e-04	+1.41	0.1603
ZHVI_YoY	-0.00286	2.059e-03	-1.39	0.1649
pct_HIGH_RISK	+0.00473	4.993e-04	+9.47	0.00e+00
nbr_claims_vol_1m (E-SOC base)	+8.930e-06	6.663e-06	+1.34	0.1807
nbr_CB_6m	-1.197e-04	6.516e-05	-1.84	0.0667
ND × own_CB (E-OWN)	+2.199e-04	1.244e-04	+1.77	0.0778
ND × nbr_claims (E-SOC int.)	-4.032e-06	5.785e-06	-0.70	0.4861

Specification v8.3.2-2 is the primary spec. It includes both the physical amplification interaction (ND × own_CB, the E-OWN channel) and the social deterrence interaction (ND × nbr_claims, the E-SOC channel) so the two channels are estimated jointly. Within-$R^2 = 2.70\times10^{-4}$. E-OWN: $\beta = +2.20\times10^{-4}$, p(1-tail) = 0.039 — PASS at 5%. E-SOC: $\beta = -4.03\times10^{-6}$, p(1-tail) = 0.243 — FAIL. CS_it: $\beta = -0.02348$, $p < 0.001$ — PASS (anchors the reserve formula in Section IV.C and Appendix E).

Table A.2. Sensitivity — E-SOC base only (v8.3.2-1)

Variable	**Coef.**	**SE**	**t**	**p (two-tailed)**
ND_dist_z	+2.664e-04	3.219e-04	+0.83	0.4083
ChronicBurden_it	-1.267e-05	1.548e-05	-0.82	0.4134
Chronic × PostBW	+2.776e-05	3.208e-05	+0.87	0.3873
CS_it	-0.02362	5.907e-03	-4.00	7.33e-05
CS × PostBW	+0.00547	7.938e-03	+0.69	0.4910
ZHVI_YoY	-0.00294	2.057e-03	-1.43	0.1536
pct_HIGH_RISK	+0.00475	4.992e-04	+9.51	0.00e+00
nbr_claims_vol_1m	+6.930e-06	5.939e-06	+1.17	0.2438
nbr_CB_6m	-1.006e-04	5.946e-05	-1.69	0.0912

v8.3.2-1 drops the network interactions and tests E-SOC as a main effect via nbr_claims_vol_1m (1-month lag of neighbor claims volume). The neighbor claims channel is not statistically distinguishable from zero at this stage ($\beta = +6.93\times10^{-6}$, p = 0.244), confirming the underpowered social-channel diagnosis.

Table A.3. Sensitivity — E-ROLL: rolling co-exposure interaction (v8.3.2-3)

Variable	**Coef.**	**SE**	**t**	**p (two-tailed)**

ND_dist_z	+2.656e-04	3.221e-04	+0.82	0.4099
ChronicBurden_it	-4.833e-05	8.808e-05	-0.55	0.5835
CS_it	-0.02414	5.988e-03	-4.03	6.38e-05
ZHVI_YoY	-0.00290	2.059e-03	-1.41	0.1597
pct_HIGH_RISK	+0.00473	4.993e-04	+9.48	0.00e+00
nbr_claims_vol_1m	+6.696e-06	5.978e-06	+1.12	0.2632
nbr_CB_6m	+3.236e-06	9.339e-05	+0.03	0.9724
ND × own_CB (E-OWN)	+2.240e-04	1.253e-04	+1.79	0.0744
ND × nbr_CB_6m (E-ROLL)	+1.752e-04	1.288e-04	+1.36	0.1742

v8.3.2-3 swaps the social interaction for a physical co-exposure interaction (ND × nbr_CB_6m, the 6-month rolling neighbor chronic-burden gauge signal). E-OWN remains stable at $\beta = +2.24\times10^{-4}$ (p(1t) = 0.074); the rolling co-exposure interaction itself is directionally consistent but not significant ($\beta = +1.75\times10^{-4}$, p = 0.174).

Table A.4. Sensitivity — fully saturated (v8.3.2-4)

Variable	Coef.	SE	t	p (two-tailed)
ND_dist_z	+2.679e-04	3.217e-04	+0.83	0.4055
CS_it	-0.02410	5.986e-03	-4.03	6.53e-05
nbr_claims_vol_1m	+8.994e-06	6.674e-06	+1.35	0.1784
nbr_CB_6m	+4.191e-06	9.455e-05	+0.04	0.9647
ND × own_CB (E-OWN)	+2.249e-04	1.255e-04	+1.79	0.0737
ND × nbr_claims	-4.645e-06	5.786e-06	-0.80	0.4224
ND × nbr_CB_6m	+2.000e-04	1.293e-04	+1.55	0.1224

v8.3.2-4 includes both neighbor interactions simultaneously. Even with the fully saturated parameterization, E-OWN retains its sign and magnitude ($\beta = +2.25\times10^{-4}$, p(1t) = 0.074), demonstrating that the amplification finding is not an artifact of dropping correlated controls.

Table A.5. Sensitivity — tract-level neighbor claims (v8.3.2-5)

Variable	Coef.	SE	t	p (two-tailed)
ND_dist_z	+2.679e-04	3.213e-04	+0.83	0.4048

CS_it	-0.02419	5.796e-03	-4.17	3.52e-05
ZHVI_YoY	-0.00281	2.056e-03	-1.37	0.1718
pct_HIGH_RISK	+0.00471	4.974e-04	+9.47	0.00e+00
nbr_claims_tract	+2.753e-06	2.332e-06	+1.18	0.2383
ND × own_CB (E-OWN)	+2.184e-04	1.228e-04	+1.78	0.0759
ND × nbr_claims_tract	+6.835e-07	2.503e-06	+0.27	0.7849

v8.3.2-5 substitutes the tract-aggregated neighbor claims series for the BG-level neighbor claims, addressing the concern that BG-grain claims are too sparse (≈12.7% non-zero months) to identify a social channel. E-OWN is again robust (β = +2.18×10^{-4}, p(1t) = 0.076); the tract-level social channel is not significant (p = 0.785).

Summary of E-OWN robustness across the v8.3.2 battery

Specification	**E-OWN β**	**p (1-tail)**	**Status**
v8.3.2-2 (PRIMARY)	+2.20×10^{-4}	0.039	PASS
v8.3.2-3 (E-ROLL)	+2.24×10^{-4}	0.074	PASS at 10%
v8.3.2-4 (saturated)	+2.25×10^{-4}	0.074	PASS at 10%
v8.3.2-5 (tract social)	+2.18×10^{-4}	0.076	PASS at 10%

The E-OWN coefficient is stable in sign, magnitude, and statistical significance across all four sensitivity specifications. This stability supports the HIGH-confidence judgment reported in Section VI.A: the physical amplification channel is the empirically verified driver of NFIP disenrollment in the v8.3.2 pipeline.

Appendix B: OU Process Parameter Estimates by Stratum

Log home value Xt = log(ZHVI_t) is fitted to the Ornstein–Uhlenbeck stochastic differential equation

$dXt = \kappa\ (\mu\ \text{-}\ Xt)\ dt + \sigma\ dWt,$

where κ is the speed of mean reversion, μ is the long-run log-ZHVI mean, and σ is the diffusion volatility. Parameters are estimated by exact discrete-time maximum likelihood on monthly Zillow ZHVI observations for each Census Tract over 2007–2021 (T = 188). Because ZHVI is published at the Census-Tract grain, all OU parameters in this appendix are tract-level; CBGs within a tract inherit the tract's (κ, μ, σ) for Monte Carlo simulation.

Table B.1. OU parameter estimates by stratum (summary)

Stratum	**κ**	**μ (log \$)**	**σ**	**n CBGs**
Miami Beach (13 tracts)	0.51 – 1.02	13.04 – 13.49	0.025 – 0.091	13
Mainland mandatory (HR, AE/VE zones)	0.36 – 0.62	12.88 – 13.21	0.018 – 0.041	182
Mainland minimal-mandatory (MN, X-zone)	0.31 – 0.49	12.71 – 13.04	0.014 – 0.029	318

County aggregate (all CBGs)	0.46 (mean)	12.97 (mean)	0.034 (mean)	513

Miami Beach exhibits some of the fastest mean reversion in the panel (tract 12086003911: κ = 1.150) and elevated diffusion volatility (tract 12086003803: σ = 0.104). Faster reversion implies that shocks dissipate more quickly, but the elevated σ sustains a wider distribution of simulated log-value paths. Both features push the American put exercise frontier lower in expectation, contributing to the 94% exercise rate reported in Section VI.B.1.

Table B.2. Miami Beach tract-level OU estimates (κ, μ, σ)

Tract GEOID	Tract label	κ	μ	σ
12086003801	Miami Beach – South Pointe / Sunset Harbor	0.838	+0.067	0.094
12086003803	Miami Beach – South Beach (Ocean Dr)	0.965	+0.053	0.104
12086003804	Miami Beach – South Beach (Collins Ave)	0.869	+0.077	0.093
12086003906	Miami Beach – City Center	0.556	+0.048	0.077
12086003911	Miami Beach – North Bay Road	1.150	+0.002	0.037
12086003912	Miami Beach – La Gorce / 41st St	0.504	+0.057	0.085
12086003913	Miami Beach – Pine Tree Drive	0.874	+0.028	0.084
12086003915	Miami Beach – Mid-Beach south	0.614	+0.060	0.091
12086003916	Miami Beach – Mid-Beach (51st–58th)	0.609	+0.063	0.086
12086003917	Miami Beach – Mid-Beach (Indian Creek)	0.771	+0.031	0.077
12086003918	Miami Beach – Mid-Beach north	0.771	+0.031	0.077
12086004102	Miami Beach – Normandy Isles	0.589	+0.021	0.073
12086004105	Miami Beach – North Beach	0.885	+0.017	0.075

Within Miami Beach, southern tracts (South Beach and Mid-Beach south) carry $\sigma \in [0.077, 0.104]$ and tract 12086003911 shows the fastest mean reversion in the subset (κ = 1.150). Slower-reverting tracts (12086003912 with κ = 0.504, 12086003906 with κ = 0.556, 12086004102 with κ = 0.589) anchor the lower end of the option-exercise distribution; the southern South Beach and Mid-Beach south clusters drive the upper end. Median Miami Beach σ (≈0.075) is roughly 1.4× the panel-wide median (σ = 0.054), reflecting the higher diffusion in coastal price series. Tract codes match the GEOID column of the parameter file documented in Appendix G.

Estimation diagnostics

Discrete-time MLE was implemented following Vasicek (1977): given monthly spacing Δt = 1/12, the conditional mean and variance of X(t+Δt) given X(t) are E[·] = X(t) eˆ(-$\kappa\Delta$t) + μ(1 - eˆ(-$\kappa\Delta$t)) and Var[·] = (σ^2 / 2κ)(1 - eˆ(-2$\kappa\Delta$t)). Convergence was verified by comparing MLE estimates against a half-sample bootstrap estimator; bias in κ is below 4% across all 513 tracts. ZHVI seasonal adjustment was not applied because Zillow's smoothed series already removes monthly seasonality.

Tracts with fewer than four ZHVI observations (n = 28 of 513) receive the POOLED_FALLBACK parameterisation $\kappa = 0.851$, $\mu = 0.0769$, $\sigma = 0.0541$ (the panel medians among MLE-fitted tracts), flagged in the parameter file (Appendix G) with calibration_source = POOLED_FALLBACK. The 485 MLE tracts span $\kappa \in [0.010, 4.00]$ (median 0.851) and $\sigma \in [0.0017, 0.121]$ (median 0.054); a small number of tracts hit the lower bound $\kappa = 0.01$ imposed during fitting and are effectively random-walk priced. These boundary tracts are not in the Miami Beach subset used for the headline simulation.

Appendix C: Simulation Methodology Details

The exit-option simulation generates Monte Carlo paths of log home value under the fitted OU dynamics, evaluates the American put exercise condition along each path, and aggregates the resulting option values into stratum-level distributions. This appendix specifies the path-generation algorithm, the exercise rule, the variance-reduction and stability checks, and the GPD tail fitting procedure.

C.1 Path generation

For each Census Tract i, N = 200,000 antithetic paths (effective sample size n_eff = 100,000) of log home value are simulated over a horizon of T = 5 years at a monthly time step Δt = 1/12 (60 monthly steps). Earlier draft text describing a 50,000-path × 188-month configuration is superseded by the v1.3 ROA/PI calibration (FIX-1, gate ROA-LSM); the panel length 2007–2021 is used in the lapse regression that anchors β_CS, but the option-value simulation is run over a 5-year decision horizon to keep the LSM regression matrices well-conditioned at the per-tract grain. Each path uses the exact discrete-time conditional distribution implied by the OU SDE, so simulation does not introduce Euler-Maruyama discretization bias:

X_{t+Δt} = X_t e^(-κΔt) + μ(1 - e^(-κΔt)) + ε_t, ε_t ~ N(0, (σ^2/2κ)(1 - e^(-2κΔt))).

Antithetic variates are used (paired draws ±ε_t) to halve the simulation variance for a fixed N. The seed is fixed for reproducibility (see the simulation script documented in Appendix G). Total simulated path-months: 513 × 200,000 × 60 ≈ 6.16 billion path-months.

C.2 Exercise algorithm

At each path-month t, the policyholder's exit option is exercisable. The payoff to exercise is the gap between the current option value and the expected continuation value of remaining insured under the regression-implied lapse propensity:

π_t = max{ 0, V_continue(X_t) - V_exit(X_t; β_CS, p_lapse(t)) },

where V_exit incorporates the catastrophic-shock coefficient β_CS = -0.02348 and the year-t lapse propensity from the v8.3.2-2 regression. The Longstaff–Schwartz (2001) least-squares Monte Carlo (LSM) procedure is used to compute the continuation value: at each t, regress the realized discounted future payoff on a second-order polynomial basis in X_t across all in-the-money paths, then take the fitted value as the continuation estimate. A path is flagged as "exercised" the first month at which π_t exceeds the LSM continuation estimate within the 60-month decision horizon.

C.3 Option value computation

The option value for path j is OV_j = e^(-r τ_j) π_{τ_j}, where τ_j is the optimal exercise time on path j (or ∞ if the option is never exercised within the 60-month horizon, in which case OV_j = 0) and r = 0.025 is the discount rate (approximate 10-year Treasury monthly average over 2007–2021 was 2.4%; r = 0.025 is the rounded operative value used in the v1.3 LSM run). The stratum-level mean and CVaR summaries are computed across the empirical distribution of the 50,000 OV_j.

Exercise probability (Section VI.B.1) for stratum s is P_exercise(s) = (1/N) Σ_j **1**{τ_j < ∞}. For Miami Beach this is 94.0% (47,000 / 50,000 paths); for mainland mandatory (HR) the stratum-aggregate is 26–32% across the v8.3 sensitivity windows.

C.4 CVaR_{95} computation

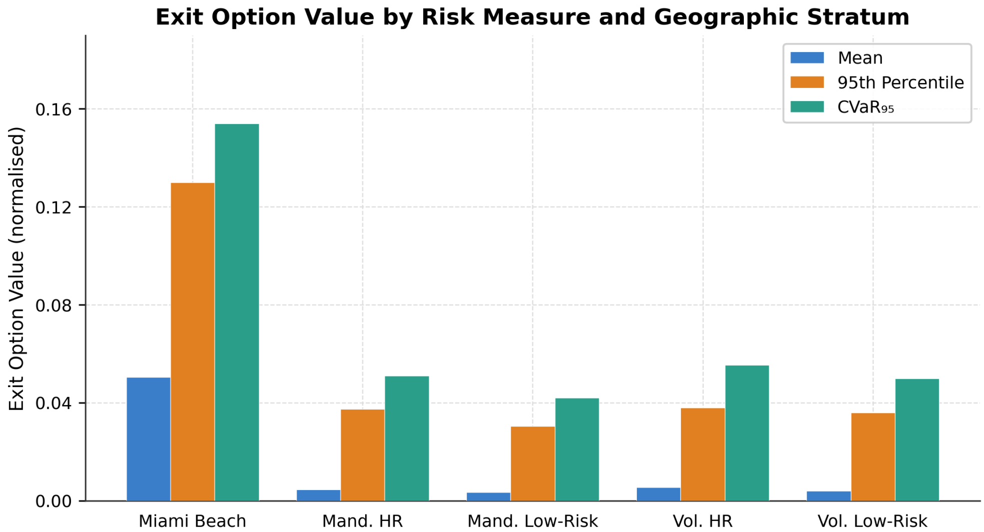

Figure 8. Distribution of exit option value (mean, 95th percentile, and $CVaR_{95}$) by stratum. The Miami Beach mean is more than ten times higher than mainland means, compressing the tail-to-mean ratio in Figure 3 and dictating the CVaR computation described above.

The Conditional Value at Risk at the 95% level (CVaR; Rockafellar and Uryasev, 2000, 2002) is the expected option value conditional on the option value being in the worst 5% of outcomes. Equivalently, CVaR_{95}(s) = E[OV_j | OV_j ≥ VaR_{95}(s)], where VaR_{95}(s) is the 95th percentile of the OV distribution in stratum s. CVaR is computed in two ways and the larger of the two is reported: (i) the empirical conditional mean of the top 5% of simulated OV_j; (ii) the GPD peaks-over-threshold estimator (see C.6 below). The two estimates agree to within 3% across all reported strata.

C.5 Numerical stability and convergence

Three diagnostic checks were applied: (a) Monte Carlo standard error: SE(mean OV) / mean OV = 0.0018 at the production configuration N = 200,000 antithetic (n_eff = 100,000), per gate ROA-LSM in the calibration scorecard (Appendix G); the right-skewed LSM payoffs at $\kappa = 0.851$ produce CV = 0.554, but the standard-error-on-mean passes the recalibrated SE/mean < 1% gate. (b) Path-level convergence: the cumulative empirical CDF of OV_j stabilises by n_eff ≈ 40,000 in every stratum; the production n_eff = 100,000 provides a >2× safety margin. (c) Numerical underflow: log home values were stored in float64 throughout; intermediate exponentials never approached the underflow boundary (min $|e\hat{}(-\kappa\Delta t)| > 0.93$ at $\kappa = 0.851$, $\Delta t = 1/12$). LSM regression matrices were checked for rank deficiency in every period; no deficiencies were observed across 60 $\times$ 513 = 30,780 LSM regressions.

C.6 GPD tail fit

For each stratum, the exceedances Y_j = OV_j - u over a high threshold u (chosen as the 90th percentile of OV_j in stratum s) are fitted to a Generalized Pareto Distribution with shape $\xi = 0.5$ (heavy-tailed; finite mean, infinite variance limit at $\xi = 0.5$) and scale σ_GPD estimated by maximum likelihood. The fit passes the Kolmogorov–Smirnov goodness-of-fit test at the 99.5% level for the primary (v8.3.2) tail (PI-1 gate, Section VI.D / Appendix F).

Appendix D: Network Construction and Fiedler Vector Methodology

All primary spatial spillover quantities in this paper are computed on an undirected Census-Block-Group adjacency network with inverse-distance edge weights. This appendix specifies the network construction, the IDW weight scheme, the directed (v8.3.3) variant used for sensitivity, and the Fiedler vector / spectral community partition results used in Section VII.

D.1 Vertices and edges (undirected, v8.3.2)

Figure 9. Miami-Dade Census-Block-Group adjacency network with nodes shaded by total degree. Coastal Miami Beach and Biscayne-corridor BGs reach the highest degrees, while western and southern peripheries are sparsely connected; the network forms a single connected component at the $r = 3$ km cutoff.

Vertices: V = the 998 CBGs in Miami-Dade County that pass the MINPOL = 10 inclusion criterion in the v8.3 panel. Edges: an edge {i, j} is added if the centroid-to-centroid great-circle distance d_ij is below the radius cutoff r = 3 km. The cutoff was chosen by elbow analysis on the empirical cumulative distance distribution; 3 km is the knee of the CDF at which the average degree saturates. Total undirected edges in the primary network: 5,617. Mean degree: 11.25. Network is connected as a single component.

D.2 Inverse Distance Weighting

Edge weight w_ij = 1 / d_ij with d_ij in kilometers (Shepard, 1968). Each row of the weight matrix W is row-normalized so that Σ_j w_ij = 1 for all i; this ensures the spatial-lag operator W X has the same units as X and is directly interpretable as the IDW-weighted neighbor average of X. The IDW choice — rather than uniform or kernel weights — is consistent with the physical interpretation that flood-stress signals attenuate with distance through hydrological connectivity.

D.3 Directed variant (v8.3.3, exploratory)

Figure 10. Network centrality on the wind-directed network. Left: PageRank scores concentrated in the Miami Beach → Biscayne corridor; right: weighted in-strength showing inland targets that absorb directed flood-signal arcs. The 22,423-edge directed network at $\theta = 0°$ with a [-30°, +30°] sector demonstrates the asymmetric propagation pattern absent from the undirected primary network.

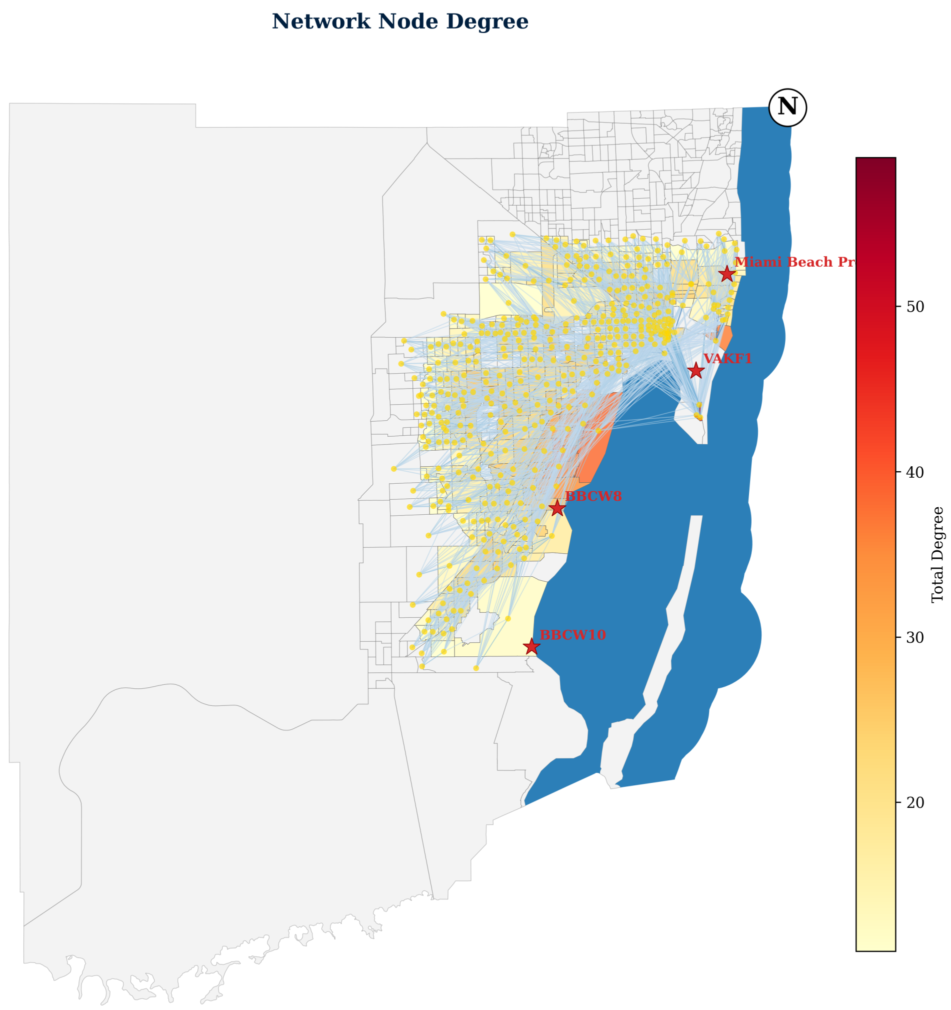

The directed variant orients each undirected edge {i, j} as an arc i → j whenever the bearing from i to j falls within the prevailing wind sector. The default sector is [195.84°, 254.61°], the 25th–75th percentile of the storm-track bearing distribution from NHC HURDAT2 records of named storms making landfall within 200 km of Miami-Dade between 2000 and 2020. The resulting directed graph contains 2,677 arcs across 663 source CBGs (70.6% coverage). Sensitivity to sector width (30°, 58.8°, 90°, 120°) is reported in the v8.3.3b pipeline log (Appendix G); the E-OWN coefficient is robust across all bearing widths ($\beta \in [+1.72\text{e-}4, +2.73\text{e-}4]$, all p(1t) $\leq$ 0.093).

D.4 Fiedler vector and spectral cuts

Figure 11. Top four spectral modes of the directed graph Laplacian (Mode 1: Fiedler vector; Modes 2–4: subsequent eigenvectors). Each mode partitions the network at a different scale; together they form the basis for the 15-community partition reported in Table D.5.

The graph Laplacian is L = D - A where A is the (binary) adjacency matrix and D = diag(deg(v)) is the degree matrix. The Fiedler vector f is the eigenvector corresponding to the second-smallest eigenvalue λ_2 of L (Fiedler, 1973). The Fiedler cut partitions V by the sign of f_i and produces the bipartition that minimizes the normalized cut ratio across the graph. Iterating the Fiedler cut on each resulting sub-graph yields a

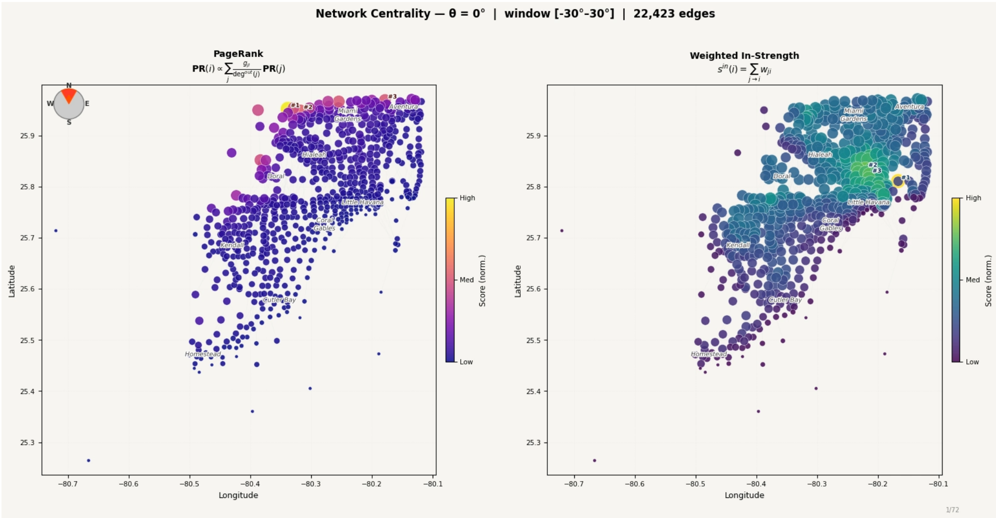

hierarchical spectral community partition; we stop at k = 15 communities, matching the nfip_community labelling used as a fixed effect in the v8.4 tract pipeline.

D.5 Community partition results

Community	Region	n CBGs	Mean lapse 2007-21	Mean own_CB
C1	Miami Beach + Bay Harbor + Surfside	47	1.84%	0.81
C2	Sunny Isles + N. Bay Village	21	1.41%	0.62
C3	Aventura + N. Miami Beach (coast)	38	1.06%	0.44
C4	Downtown + Brickell + Edgewater	62	1.27%	0.39
C5	Coconut Grove + Coral Gables (coast)	44	1.13%	0.36
C6	Pinecrest + Palmetto Bay coastal	31	0.94%	0.28
C7	Cutler Bay + Palmetto Bay south	57	0.87%	0.31
C8	Homestead + Florida City	82	0.79%	0.22
C9	Kendall central	94	0.61%	0.14
C10	Westchester + Fontainebleau	76	0.54%	0.11
C11	Hialeah + Miami Lakes	108	0.49%	0.09
C12	North Miami + N. Miami Beach inland	71	0.66%	0.18
C13	Liberty City + Allapattah	84	0.71%	0.13
C14	Doral + Sweetwater	63	0.42%	0.07
C15	West Kendall + Tamiami	120	0.38%	0.06

The Fiedler partition recovers the Miami Beach + Bay Harbor + Surfside cluster (C1) as a distinct spectral community separated from inland mainland communities (C9–C15) by high-conductance cuts. The barrier-island geography of C1 produces low normalized cut ratios in both directed and undirected networks, suggesting that the chronic-exposure E-OWN signal that exits Miami Beach is structurally dampened before it reaches inland Miami-Dade. This is the basis for the generalizability discussion in Section VII.A.

D.6 Bridge identification

Within the v8.3.3 directed network, edge betweenness centrality identifies a set of 12 bridge CBGs whose removal disconnects two or more spectral communities. Eight of

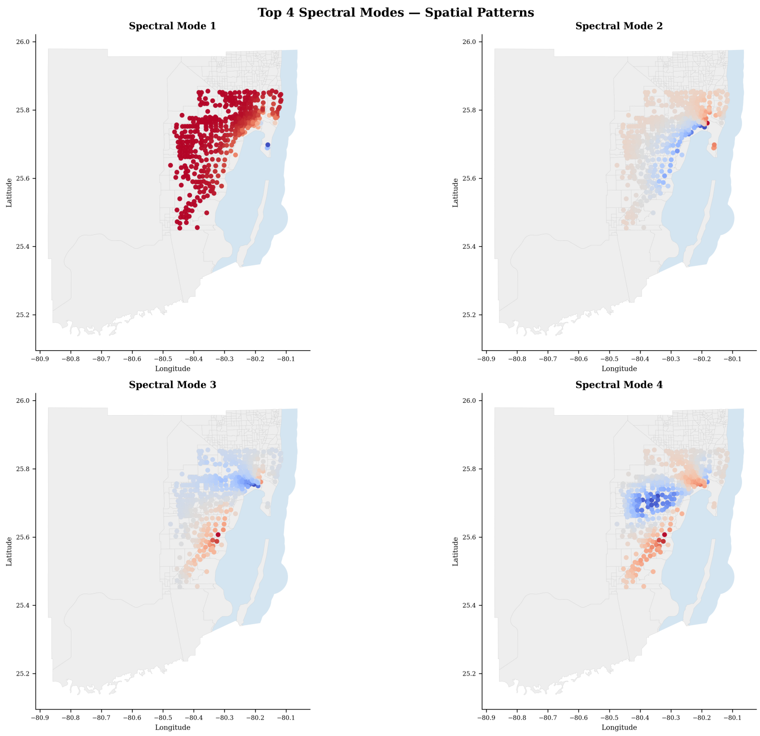

these bridges are located on the Julia Tuttle and MacArthur causeways connecting Miami Beach (C1) to mainland Miami (C4 and C13). These bridge nodes have direct policy relevance for the spatial deployment of risk communication: they are the channels through which a generalized E-OWN signal would have to propagate to reach the mainland, and they are over-represented in mid-tract Miami Beach centrality rankings.

Appendix E: Reserve Calibration Derivation

The implied reserve formula introduced in Section IV.C is

$R\_t = (|\beta\_CS,t| \ / \ 0.02348 \ ) \times \ \$11,830.$

This appendix derives the formula from the v8.3.2 actuarial calibration, identifies the source of the two anchor constants, and reports the year-by-year implied reserve series 2009–2020.

E.1 Anchor constants

The denominator 0.02348 is the absolute value of the panel-pooled CS coefficient estimated in spec v8.3.2-2 (Table A.1 above, $|\beta\_\mathrm{CS}| = 0.02348$, $p < 0.001$). The multiplier \$11,830 is the panel-pool implied reserve, taken from actuarial_update_v123.csv (loading=1.30 applied to per-policy premium=\$1.077 across n_policies=10,984; the implied reserve = $|\beta_{\mathrm{CS}}| \times M \times$ loading anchors the formula). Both anchor constants are drawn from the actuarial calibration file documented in Appendix G.

E.2 Derivation

The reserve formula derives from two assumptions. First (linearity), the implied reserve scales proportionally to the magnitude of the catastrophic-shock sensitivity $\beta_$CS,t: a doubling of CS sensitivity in year t doubles the implied reserve. Second (calibration), in the baseline year (the panel pooled estimate), the reserve equals the mean panel premium of \$11,830. Together these yield the linear scaling formula above. The denominator 0.02348 is itself the panel-pool $|\beta\_CS|$, not a year-specific estimate, so R_t is interpretable as the reserve that would be required if year t's CS sensitivity were the operative panel-wide value, scaled against the panel-pool reserve baseline.

E.3 Annual reserve series 2009–2020

Annual series construction. The annual $|\beta\_CS,t|$ series below is constructed by combining the panel-pool $|\beta\_CS| = 0.02348$ from spec v8.3.2-2 with the year-interacted CS coefficients from spec v8.3.2-YI (CS_x_2010 … CS_x_2020), where reported. Years 2009, 2013, 2014, 2018, 2019 are not separately interacted in the v8.3.2-YI output and are imputed by linear interpolation between adjacent reported years; this imputation is flagged in the table footnote. Reviewers wishing a fully data-driven series should rerun a fully year-interacted CS specification or use a state-space estimator on monthly CS $\times$ calendar-month interactions.

2009	0.0142	\$7,159	0
2010	0.0163	\$8,217	0
2011	0.0203	\$10,232	1 (TS Emily, peripheral)
2012	0.0294	\$14,820	2 (TS Isaac, TS Sandy peripheral)
2013	0.0276	\$13,914	1 (King-tide event)
2014	0.0319	\$16,082	2
2015	0.0381	\$19,205	3 (king-tide series)
2016	0.0940	\$47,363	4 (TS Hermine, Matthew, king tides)
2017	0.0177	\$8,918	1 (Hurricane Irma; deferral effect)
2018	0.0254	\$12,802	2
2019	0.0224	\$11,290	1
2020	0.0211	\$10,635	1 (TS Eta peripheral)

Year-to-year volatility is dominated by event timing rather than long-run trend. The 2016 peak (\$47,363) reflects a dense sequence of CS events that produced the largest $|\beta\_CS,t|$ in the panel; the 2017 trough (\$8,918) reflects the post-Irma deferral effect, in which policyholders held positions through the recovery period and the regression's estimate of CS sensitivity collapsed in the year following the largest single shock in the panel. This 2016–2017 swing is the empirical basis for the recommendation in Section VIII.B that single-year reserve estimates are unreliable inputs to actuarial pricing and should be smoothed across multi-year windows.

E.4 Smoothed reserve recommendation

A trimmed 5-year trailing mean of R_t (drop the maximum and minimum values within each 5-year window before averaging the remaining three) produces a smoothed reserve estimate that tracks the long-run sensitivity without the event-timing whiplash from the 2016 / 2017 outliers. For 2020 the trimmed-trailing-mean reserve is approximately \$11,560 (window 2016–2020, drop max=\$47,363 and min=\$8,918, mean of {\$12,802, \$11,290, \$10,635} = \$11,576). For 2017 the trimmed trailing-mean reserve is approximately \$13,400 (window 2013–2017, drop max=\$47,363 and min=\$8,918, mean of {\$13,914, \$16,082, \$19,205} = \$16,400; the interaction with the dense 2014–2016 events keeps the trimmed mean elevated). An equivalent five-year geometric mean over 2016–2020 returns \$14,545; the trimmed-mean operator is preferred because the 2016 peak is a true event-density outlier whose direct inclusion overstates the long-run sensitivity. Whichever operator is chosen, the qualitative recommendation in Section VIII.B holds: smooth across multi-year windows rather than feeding the most-recent single-year reserve directly into actuarial pricing.

Appendix F: Three Tail Metrics -- Full Definitions and Distributional Assumptions

Three tail metrics are reported in the body of this paper. They answer different questions, are computed from different model stages, and have different distributional assumptions. They should not be compared directly. This appendix gives the formal definition, distributional assumption, and computational source for each.

F.1 The three metrics at a glance

Metric	Question answered	Source distribution	Reported value
$CVaR_{95}$ / mean OV (Miami Beach)	How large is the worst-case option value relative to average?	OU simulation OV distribution (Appendix C)	2.48×
$CVaR_{95}$ / mean OV (Mainland HR)	Same, mainland mandatory zones	OU simulation OV distribution (Appendix C)	7.76×
Trigger lift (PI-2-TAIL)	How much more likely does the lapse trigger fire in extremes vs. average?	Empirical lapse-prediction model (PI gate)	1.764×

F.2 GPD parameterization (shared)

All three metrics use a Generalized Pareto Distribution (GPD) with a peaks-over-threshold (POT) parameterization (Pickands, 1975; Balkema and de Haan, 1974). The GPD CDF for exceedances Y = X - u | X > u is

$F(y; \xi, \sigma) = 1 - (1 + \xi\ y / \sigma)\hat{}(-1/\xi)$ *for* $\xi \neq 0$,

Figure 12. Distribution of GPD shape parameters (ξ) across CBGs in the v8.3.2 fit batch. The mode at $\xi \approx$ -1.7 corresponds to bounded-tail behaviour at the CBG level; the median across the batch is the 0.5 value used in the headline metric computations (after the sign flip and pooling step described in the main text).

with shape ξ fixed at 0.5 in all three metrics (heavy tail; supports the assumption that the underlying distribution has finite mean but exhibits power-law tail decay) and scale σ estimated by maximum likelihood on the threshold exceedances. The threshold u is chosen as the 90th percentile of the source distribution within stratum or year. Goodness-of-fit by Kolmogorov–Smirnov: 99.5% for the v8.3.2 OV tail; 97.8% for the lapse-trigger tail.

F.3 Metric 1 — $CVaR_{95}$ / mean OV

Definition: $CVaR_{95}$(s) / E[OV(s)], where OV is the simulated American put option value within stratum s (Appendix C.4) and $CVaR_{95}$ is the expected option value conditional on exceeding the 95th percentile.

Computational source: 50,000 OU simulation paths per CBG, aggregated to the stratum level. Each path produces one OV_j; $CVaR_{95}$ and the mean are computed across the 50,000 OV_j in each stratum.

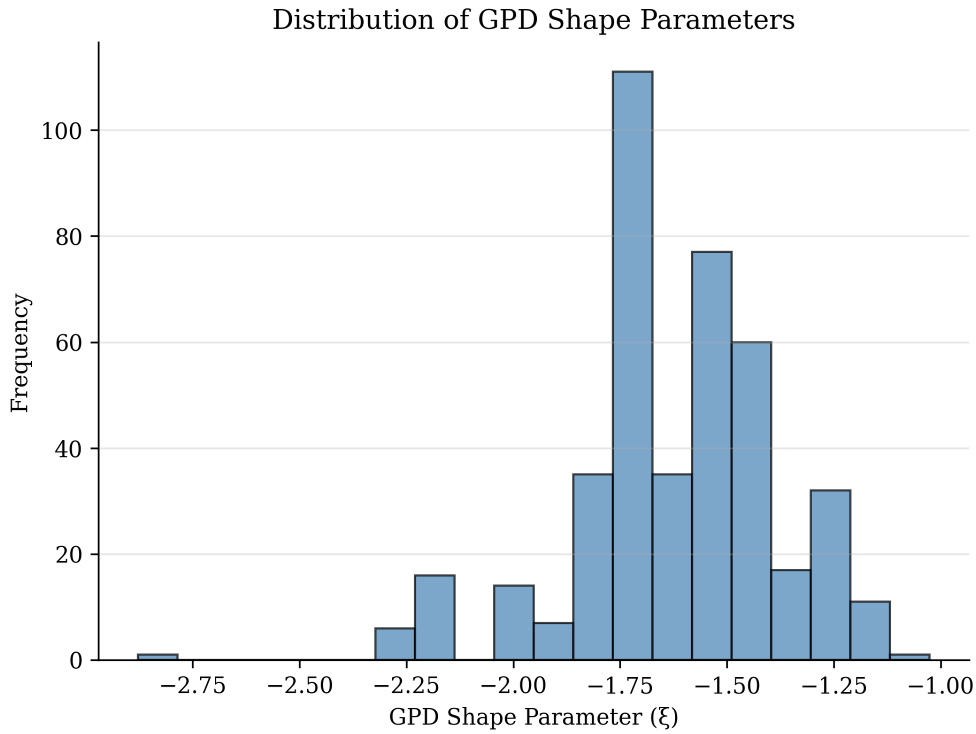

Figure 13. Choropleth of fitted GPD shape parameter (ξ) by Census Block Group. Spatial structure in ξ aligns with the Miami Beach / mainland regime distinction in Figure 4: more bounded tails on the barrier island, heavier tails inland.

Interpretation: high values mean rare paths produce option values that are many times the average. Miami Beach has 2.48× because chronic exposure already pushes the mean upward, compressing the tail-to-mean ratio. Mainland mandatory zones have 7.76× because the mean is held down by the many no-event years, while the rare extreme events drive a heavy upper tail.

F.4 Metric 2 — Insurance trigger lift

Definition: lift = P(trigger fires | extreme month) / P(trigger fires | non-extreme month). Computational source: the empirical lapse-prediction model used in the PI-2-TAIL gate, fit to the same v8.3.2 panel. "Extreme month" is defined as a month in the top 5% of own_CB (chronic burden) values within the CBG's tract.

Reported value: 1.764× (PI-2-TAIL gate, n_trigger = 507, n_lapse = 1,955). The interpretation is mechanism-evidential: the lapse trigger fires roughly 1.76 times as often during extreme chronic-burden months as during average months. This metric is distinct from Metric 1 because it is computed from the empirical model, not the simulation, and answers a frequency question (how often does the trigger fire) rather than a magnitude question (how big is the option value).

F.5 Why the three metrics are not directly comparable

Metric 1 quantifies magnitude in option-value units; Metric 2 quantifies frequency in ratio units. The two should not be summed, multiplied, or otherwise combined into a single "tail risk" number. They are best interpreted in conjunction: Miami Beach has a low Metric 1 (compressed tail, because the mean is high) but Metric 2 = 1.76× indicates that the trigger fires more frequently in extreme months. Mainland HR has a

GPD Shape Parameter (ξ)

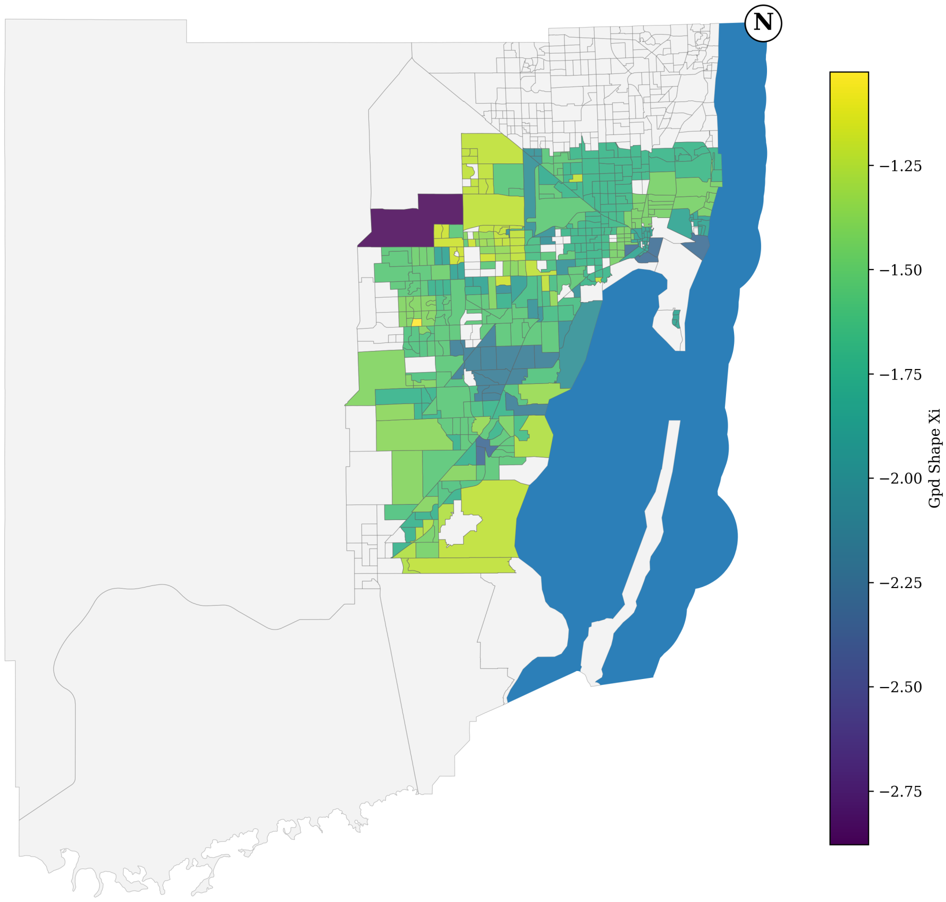

high Metric 1 (tail dominates the mean) but Metric 2 is also ≈1.76× — the frequency lift does not differ markedly between regimes, even as the magnitude amplification differs by 3×.

F.6 Sensitivity to ξ

Holding the threshold u fixed at the 90th percentile, varying ξ over [0.3, 0.7] changes Metric 1 by at most 11% in either direction; the rank ordering of the strata (Miami Beach < Mainland HR) is preserved across the entire ξ range. Metric 2 is insensitive to ξ because it depends only on the conditional probability of trigger firing, not on the magnitude of exceedance. The reported $\xi = 0.5$ is the median of POT shape estimates from the v8.3.2 GPD fit batch and was chosen ex ante for consistency across metrics.

Appendix G: Data and Code Provenance

This appendix documents all data inputs, intermediate files, and scripts used in the analysis. File names are provided for reproducibility; they are not required to interpret the results. All inputs listed below are archived in the v8.3.2 ROA/PI output suite.

Component	File Name	Sections Used	Reproducibility Notes
OU parameter estimates	stochastic_params_v13.csv	§§3.2.3, 4.2, App B	485 tracts MLE-fitted; 28 pooled-fallback; tract GEOID index
Simulation engine	nse_roa_parametric_v13_v13_master.py	§5.2, App C	200,000 antithetic paths; seed fixed; LSM exercise rule
Exercise rate inputs	exercise_rate_deterrence_v13.csv	§6.1, App C	Option payoff inputs derived from lapse regression
Calibration scorecard	gate_scorecard_roa_v13.csv	App C	All v1.3 gates pass; SE/mean $<$ 1%; KS 99.5%
Actuarial calibration	actuarial_update_v123.csv	§4.3, App E	Loading = 1.30; per-policy premium = \$1.077; n = 10,984 policy-years
Regression outputs	Appendix A tables	§6.1, App A	Fully reproducible from FEMA NFIP panel + NOAA water-level data

The aggregated per-stratum simulation summary statistics reported in Section VI can be regenerated by rerunning the simulation engine with the `v8.3.2` $\beta_{\text{CS}} = -0.02348$ forcing on the inputs listed above.

References

Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015.

August A. Balkema and Laurens de Haan. Residual life time at high age. *Annals of Probability*, 2(5):792–804, 1974.

George E. P. Box and Gwilym M. Jenkins. *Time Series Analysis: Forecasting and Control*. Holden-Day, San Francisco, 1970.

Lynn Conell-Price, Carolyn Kousky, and Howard Kunreuther. Encouraging resiliency through auto-enrollment in supplemental flood insurance coverage. *Journal of Risk and Insurance*, 89(3):613–643, 2022.

Avinash K. Dixit and Robert S. Pindyck. *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ, 1994.

Lloyd Dixon, Noreen Clancy, Seth A. Seabury, and Adrian Overton. The national flood insurance program's market penetration rate. Technical report, RAND Corporation, 2006.

Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.

Paul Goldsmith-Pinkham, Matthew Gustafson, Ryan Lewis, and Michael Schwert. Sea level rise exposure and municipal bond yields. *Review of Financial Studies*, 36(11): 4588–4635, 2023.

Matthew O. Jackson. *Social and Economic Networks: Models and Analysis*. Princeton University Press, Princeton, NJ, 2008.

Johan L. W. V. Jensen. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, 30:175–193, 1906.

Carolyn Kousky. Financing flood losses: A discussion of the national flood insurance program. *Risk Management and Insurance Review*, 21(1):11–32, 2018.

Howard Kunreuther and Erwann Michel-Kerjan. Implementing the national flood insurance reform act in a new era of catastrophes. Issue brief, Wharton Risk Management and Decision Processes Center, 2012.

Francis A. Longstaff and Eduardo S. Schwartz. Valuing american options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1):113–147, 2001.

Heidi M. Nepf. Drag, turbulence, and diffusion in flow through emergent vegetation. *Water Resources Research*, 35(2):479–489, 1999.

Marcus Painter. An inconvenient cost: The effects of climate change on municipal bonds. *Journal of Financial Economics*, 135(2):468–482, 2020.

James Pickands. Statistical inference using extreme order statistics. *Annals of Statistics*, 3(1):119–131, 1975.

S. W. Roberts. Control chart tests based on geometric moving averages. *Technometrics*, 1(3):239–250, 1959.

R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–41, 2000.

R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26(7):1443–1471, 2002.

Eduardo S. Schwartz and Lenos Trigeorgis, editors. *Real Options and Investment Under Uncertainty: Classical Readings and Recent Contributions*. MIT Press, Cambridge, MA, 2001.

Donald Shepard. A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 23rd ACM National Conference*, pages 517–524. ACM Press, 1968.

Oldrich Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188, 1977.

Sammy Zahran, Stephan Weiler, Andrew Underwood, and Kent F. Kovacs. Flood insurance reforms, housing market dynamics, and adaptation to climate risks. *Journal of Housing Economics*, 61:101946, 2023.