Climate Risk, House Price Tail Risk, and Mortgage Default

Abstract

The housing market faces increasing vulnerability to climate-related risks, with important implications for mortgage performance and financial stability. This paper examines how climate-related tail risk in house prices affects mortgage default outcomes. We develop an integrated framework that combines house price modelling, Extreme Value Theory (EVT) for tail-event simulation, and a loan-level mortgage default model. Using Zillow house price data (2000–2022), FEMA climate risk classifications, and Freddie Mac loan-level data (2001–2018), we simulate future house price paths under both baseline and tail-adjusted scenarios and map these into predicted default probabilities. Our approach allows us to quantify the incremental impact of tail risk through a measure we term **default uplift**, defined as the difference in predicted default probabilities between EVT-based and baseline simulations. We find that incorporating tail risk leads to economically meaningful increases in default probabilities, with the largest effects observed in regions with moderate climate exposure rather than in the highest-risk areas. The impact is concentrated in the upper tail of the distribution and varies across forecast horizons, reflecting the interaction between persistent house price dynamics and extreme shocks. Overall, our findings highlight the importance of modelling tail risk when assessing climate-related vulnerabilities in mortgage markets and suggest that approaches focusing solely on average outcomes may understate the true extent of risk.

Keywords: Climate change, mortgage default, extreme value theory

JEL: Q54, G21

1. Introduction

The increasing frequency and severity of climate-related events present substantial risks to financial markets, particularly mortgage markets, which are intrinsically linked to real estate valuations and borrower stability. A growing body of research documents the impact of climate risk on property values. For example, Bernstein et al. (2019) and Baldauf et al. (2020) show that properties exposed to sea-level rise and flooding trade at significant discounts, highlighting market sensitivity to environmental risks. Keys and Mulder (2020) further demonstrate that property values decline following major flooding events, although these effects may attenuate over time due to collective memory decay. Gibson et al. (2021) provide evidence of heterogeneity in climate risk pricing across geographic markets, suggesting that expectations of future climate exposure are increasingly incorporated into current valuations. While this literature establishes that climate risk is reflected in housing markets, it focuses primarily on average price effects and cross-sectional differences. However, climate risk is inherently characterised by rare but severe events, raising the question of how extreme shocks in house prices affect mortgage performance. Standard approaches to modelling house price dynamics often struggle to capture such extreme movements. Glaeser and Nathanson (2017) and Davis and Quintin (2017) highlight the limitations of conventional models in explaining housing market volatility, particularly during periods of stress. Similarly, Ortega and Taspinar (2018) and Bin and Landry (2013) show that extreme climate events can lead to persistent disruptions in local housing markets. These findings point to the importance of modelling tail risk, for which Extreme Value Theory (EVT), as developed by Embrechts et al. (1997), provides a well-established framework.

At the same time, the mortgage default literature emphasises the central role of house price dynamics in shaping borrower behaviour. Campbell and Cocco (2015) highlight the interaction between negative equity and liquidity constraints in driving default decisions,

while Foote et al. (2008) and Gerardi et al. (2018) demonstrate that declines in house prices are a key determinant of mortgage distress. More recent work has begun to explore the implications of climate risk for mortgage markets. Keenan and Bradt (2020) and Ouazad and Kahn (2021) provide evidence that climate exposure can influence mortgage performance and lender behaviour, suggesting increasing awareness of climate-related credit risk. Despite these advances, the literature lacks an integrated framework that links climate risk exposure to the tail behaviour of house price dynamics and, in turn, to mortgage default outcomes.

Existing studies either focus on how climate risk affects property values, how house prices influence default, or how extreme events can be modelled statistically, but they do not connect these elements in a unified empirical setting. As a result, the extent to which climate-induced tail risk in house prices translates into mortgage default risk remains unclear. We focus on a distinct channel: the role of extreme house price shocks. We show that climate risk operates through tail events rather than average price movements, and that standard models substantially understate default risk when these extremes are ignored.

This paper addresses this gap by developing an integrated framework that links climate-related tail risk in house prices to loan-level default outcomes. Using Zillow Home Value Index data, FEMA climate risk designations, and Freddie Mac loan-level data, we combine three components. First, we estimate a dynamic house price model at the MSA level. Second, we model the tail behaviour of house price shocks using EVT, allowing for the simulation of rare but economically significant events. Third, we embed the resulting simulated house price paths into a mortgage default model to quantify how extreme shocks affect default probabilities.

A key feature of our approach is that it enables a direct comparison between a baseline model, which assumes standard residual behaviour, and an EVT-based model that incorporates fat-tailed shocks. The difference between these two scenarios provides a measure of the incremental impact of tail risk on mortgage default, which we refer to as **default uplift**. This framework allows us to move beyond average effects and directly quantify how extreme events alter credit risk.

Our results show that incorporating tail risk leads to economically meaningful increases in predicted default probabilities, but the effects are not uniform across regions. In particular, we find that the impact of tail risk is strongest in areas with moderate levels of climate exposure, rather than in the highest-risk regions. Moreover, the effects are concentrated in the upper tail of the distribution and vary over time, reflecting the interaction between persistent house price dynamics and extreme shocks.

This paper contributes to several strands of the literature. First, it contributes to the climate risk and housing literature by shifting the focus from average effects to tail-driven dynamics. Second, it contributes to the mortgage default literature by incorporating extreme house price shocks into default predictions. Third, it contributes to the application of EVT in finance by embedding tail modelling within a framework that links asset prices to credit risk.

Our contribution is to provide an integrated framework linking climate risk, house price dynamics, and mortgage default through the explicit modelling of tail risk. We show that incorporating extreme value behaviour materially increases predicted default probabilities, highlighting a channel that is largely absent from existing studies.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the data. Section 4 describes the empirical framework. Section 5 presents the empirical results and discusses implications, and Section 6 concludes.

2. Literature Review

The intersection of climate risk and mortgage markets has received increasing attention in recent academic literature, particularly as climate-related disasters become more frequent and severe. Existing research spans several related domains, including climate risk and real estate valuation, house price dynamics under extreme events, mortgage default modelling, and the pricing of climate risk in mortgage markets. While each of these strands provides important insights, they remain largely disconnected. This paper integrates these strands by linking climate-related tail risk in house prices to mortgage default outcomes.

2.1 Climate Risk and Real Estate Valuation

A growing body of research examines how climate risk is reflected in real estate prices. Bernstein et al. (2019) document that homes exposed to sea level rise sell at a discount of approximately 7% relative to comparable low-risk properties. Similarly, Baldauf et al. (2020) find that properties exposed to sea level rise trade at a discount of around 6.6%, with the magnitude of the effect varying with beliefs about climate change.

The temporal dynamics of climate risk pricing are highlighted by Keys and Mulder (2020), who show that property values decline following major flooding events but gradually recover as collective memory fades. This “memory decay” suggests that markets may underweight long-term climate risks, leading to potential mispricing. Gibson et al. (2021) extend this literature by constructing a comprehensive climate risk index and documenting substantial heterogeneity in how different hazards are priced across regions.

While this literature establishes that climate risk is incorporated into property values, it primarily focuses on average price effects and cross-sectional variation. As a result, it provides limited insight into how climate risk affects the distribution of house price changes, particularly in the tails. Given that climate-related hazards often manifest as rare but severe events, understanding tail behaviour is critical for assessing their full economic impact.

2.2 House Price Dynamics and Extreme Events

The modelling of house price dynamics has evolved to better capture market volatility and the effects of extreme events. Early work by Case and Shiller (1988) documents strong serial correlation in house prices, highlighting deviations from fully efficient market behaviour. More recent studies incorporate behavioural and informational frictions. Glaeser and Nathanson (2017) show that extrapolative expectations can amplify price cycles, while Davis and Quintin (2017) demonstrate that models incorporating limited information and learning better explain housing market fluctuations during periods of stress.

In the context of environmental shocks, Ortega and Taspinar (2018) document significant price declines following Hurricane Sandy, with persistent effects in severely affected areas. Similarly, Bin and Landry (2013) show that flood-related discounts diminish over time but re-emerge following subsequent events. These findings underscore the importance of modelling extreme shocks in house price dynamics.

Standard time series models, however, often assume well-behaved residual distributions and may fail to capture the fat-tailed nature of house price changes during periods of stress. Extreme Value Theory (EVT), as formalised by Embrechts et al. (1997), provides a framework for modelling such tail behaviour. While widely used in financial risk management, its application to housing markets remains limited. Martins-Filho and Yao

(2006) extend EVT-based approaches to non-linear financial return processes, highlighting their usefulness in capturing extreme outcomes.

2.3 Mortgage Default Modelling

The mortgage default literature emphasises the central role of house price dynamics in shaping borrower behaviour. Campbell and Cocco (2015) develop a theoretical framework showing how negative equity and liquidity constraints jointly drive default decisions. Empirical evidence from Foote et al. (2008) and Gerardi et al. (2018) confirms that house price declines are a key determinant of mortgage distress, although negative equity alone is often insufficient without additional "trigger events," consistent with the double-trigger hypothesis.

More recent research has begun to incorporate climate-related risks into mortgage market analysis. Keenan and Bradt (2020) provide evidence of "climate gentrification," where properties in less exposed areas command increasing premiums. Ouazad and Kahn (2021) show that lenders respond to climate risk by adjusting securitisation patterns, effectively redistributing risk across market participants.

Despite these advances, mortgage default models generally rely on standard representations of house price dynamics, focusing on average changes rather than extreme outcomes. As a result, the potential impact of tail risk in house prices on mortgage default remains largely unexplored.

2.4 Synthesis and Contribution

Taken together, these strands of literature highlight an important gap. The climate risk literature documents how environmental exposure is capitalised into property values but focuses on average effects. The house price dynamics literature recognises the importance of extreme events but rarely models them explicitly using tail-based approaches. The mortgage

default literature establishes the importance of house price changes for credit risk but does not account for the distribution of extreme shocks. While EVT provides a well-established framework for modelling tail risk, it has not been systematically integrated into models linking house prices to mortgage default.

This paper bridges these strands by developing an integrated framework that connects climate-related tail risk in house prices to mortgage default outcomes. By combining a dynamic house price model with EVT-based residual modelling and embedding the resulting simulated price paths into a loan-level default model, we provide a direct assessment of how extreme shocks affect credit risk.

Importantly, our framework allows for a comparison between baseline and tail-adjusted scenarios, yielding a measure of **default uplift** that captures the incremental effect of extreme events on mortgage risk. In doing so, the paper moves beyond average effects and provides a structured approach to quantifying the role of tail risk in climate-related financial vulnerability.

3. Data and Empirical Models

This section provides an overview of the three primary data sources used in our analysis: Zillow House Price Index data, FEMA climate risk data, and Freddie Mac loan-level data. These datasets are combined to develop our integrated framework for assessing climate risk impacts on mortgage markets.

3.1 Zillow House Price Index Data

Our analysis of house price dynamics utilizes Zillow's monthly Home Value Index (HVI) for single-family residences. This dataset provides consistent, quality-adjusted measures of housing market values at the MSA level across the United States. We employ data spanning from January 2000 through September 2022, capturing both the Global Financial Crisis

housing market collapse and the subsequent recovery, as well as more recent market dynamics.

The Zillow HVI offers several advantages for our analysis. First, it employs a hedonic methodology that controls for the characteristics of homes being sold, providing a more accurate measure of pure price changes than median sale price indices. Second, its monthly frequency allows us to capture short-term market dynamics that might be missed in quarterly data. Third, its broad geographic coverage enables us to analyze climate risk effects across diverse markets.

From this dataset, we calculate the yearly percentage change in house prices as our primary dependent variable. This transformation helps address non-stationarity in the raw house price series and provides a measure that can be readily interpreted as a return series, facilitating our extreme value theory application.

3.2 FEMA Climate Risk Data

To incorporate climate risk factors into our analysis, we utilize data from the Federal Emergency Management Agency's (FEMA) National Flood Hazard Layer (NFHL). This dataset provides detailed information on flood risks across the United States, allowing us to categorize areas into high, medium, and low climate risk designations.

The NFHL incorporates detailed hydrologic and hydraulic analyses, historical flood information, and other variables to determine flood hazard areas. Properties in Special Flood Hazard Areas (SFHAs) are designated as having a 1% or greater annual chance of flooding (sometimes referred to as the 100-year floodplain). For our analysis, we classify areas within SFHAs as high risk, areas with a 0.2% to 1% annual flood chance (the 500-year floodplain) as medium risk, and areas outside these designations as low risk.

While the FEMA data primarily focuses on flood risk, it serves as a reasonable proxy for climate vulnerability more broadly, as flooding represents one of the most common and

costly climate-related hazards affecting residential properties. However, we acknowledge that a more comprehensive climate risk assessment would ideally incorporate additional hazards such as wildfires, hurricanes, and extreme heat.

The integration of FEMA risk designations with Zillow house price data presents challenges due to differences in geographic granularity and boundary definitions. We address this by aggregating FEMA risk designations to the MSA level, calculating the proportion of properties within each risk category for each MSA.

3.3 Freddie Mac Loan-Level Data

To model mortgage performance and default probabilities, we utilize Freddie Mac's Single Family Loan-Level Dataset. This dataset contains loan-level credit performance data on a portion of fully amortizing fixed-rate mortgages purchased by Freddie Mac from 2000 to 2018. We focus on 30-year fixed-rate mortgages, which represent the predominant mortgage product in the U.S. market.

The dataset contains detailed information on loan characteristics at the time of origination, including credit score (FICO), original loan-to-value ratio (LTV), debt-to-income ratio (DTI), loan purpose (whether for purchase or refinance), occupancy status (specifically owner-occupied), documentation level (using a full documentation indicator), property location (identified by state and metropolitan statistical area), and origination date. In addition to these origination variables, the dataset provides monthly performance data such as current loan status (e.g., current, delinquent, foreclosure, real estate owned), months of delinquency, modification status, current loan balance, and current interest rate.

For the purpose of default modeling, we define a default event as a loan becoming 90 or more days delinquent within three years of origination, a definition consistent with industry practice and regulatory standards for identifying serious delinquency. To ensure the relevance

and internal validity of our analysis, we apply several filters: we include only first-lien mortgages on owner-occupied properties located in California and Florida (for the simulation exercise), and restrict the sample to loans with complete information on key variables such as FICO, LTV, and DTI. These criteria allow us to focus on a relatively homogeneous set of loans, thereby reducing the influence of confounding factors on our results.

3.4 Data Integration Approach

Integrating our datasets requires careful handling of geographic and structural mismatches. Zillow house price data and FEMA climate risk designations are available at different spatial resolutions, while the loan-level performance data from Freddie Mac are reported at the three-digit ZIP code level. Additionally, the mismatch between the MSA-level house price model and ZIP-level loan data introduces a challenge in applying climate-adjusted price dynamics to individual loans.

To address these challenges, we proceed as follows:

1. **Aggregate FEMA climate risk classifications** to the MSA level, calculating the proportion of properties in each risk category (Very High, High, Moderate, Low).
2. **Link Freddie Mac loans to simulated house price dynamics indirectly** by aggregating MSA-level simulated HPI paths to the **state level**, which matches the geographic unit used in the mortgage default model.
3. **Estimate a dynamic HPI model** at the MSA level and simulate future house price changes (October 2022 to October 2035) using a recursive model that includes lagged HPI changes and EVT-based residual draws by FEMA risk category.
4. **Apply the state-level simulated HPI changes** in the mortgage default model by replacing the historical HPI change variables and associated terms (negative dummy and interaction) with their simulated counterparts. All other borrower- and loan-level characteristics are held constant at observed values.

5. **Predict loan-level default probabilities** for each forecast period using the estimated default model and the simulated state-level HPI inputs. These predictions are used to assess how default risk evolves under climate-adjusted house price dynamics.

This integrated framework enables us to trace the effect of climate-related housing market stress modelled through simulated tail risk in house prices on loan performance. By consistently applying simulated macro shocks across heterogeneous borrowers, we isolate the role of climate exposure in driving geographic variation in mortgage default probabilities.

4. Methodology

4.1 Overview of the Empirical Framework

We develop a three-stage empirical framework that links climate-related risk to mortgage default outcomes through its impact on house price dynamics and tail behaviour. The framework proceeds sequentially.

First, we estimate a baseline model of house price dynamics using an autoregressive specification augmented with relevant controls. This model captures the central tendency of house price movements and serves as the benchmark for subsequent analysis.

Second, we model the distribution of extreme house price movements by applying Extreme Value Theory (EVT) to the residuals obtained from the baseline house price model. This step allows us to characterise tail risk that is not captured by standard time-series specifications.

Third, we generate simulated house price paths under both the baseline and EVT-adjusted distributions. These simulated price paths are then mapped into mortgage default probabilities through a reduced-form default model that links house price declines to borrower equity positions.

This structure enables a direct comparison between default probabilities implied by the baseline model and those incorporating tail risk. We define the difference between these two

estimates as the ***uplift***, which captures the incremental effect of tail risk on mortgage default outcomes.

4.2 Baseline House Price Model

We model house price dynamics using an autoregressive specification of order three (AR(3)), augmented with a set of control variables capturing broader macroeconomic and local housing market conditions. The AR(3) structure captures persistence in house price movements and provides a flexible representation of short- to medium-term dynamics.

Let HPI_{it} denote the house price index for MSA i at time t. The baseline model is specified as:

$$HPI_{it} = \alpha + \beta_1 HPI_{it-1} + \beta_2 HPI_{it-2} + \beta_3 HPI_{it-3} + \gamma X_{it} + \epsilon_{it}$$

where X_{it} represents control variables and ϵ_{it} is the residual term.

In practice, we find that the autoregressive components capture the majority of variation in house price movements, with additional controls contributing marginal explanatory power. The residuals from this model therefore contain the variation that is not explained by standard dynamics, including potential extreme movements associated with climate-related shocks. These residuals form the basis for the subsequent EVT analysis.

Climate-related shocks are inherently rare but severe, making them ill-suited to standard parametric distributions that focus on central tendencies. Conventional approaches, including Gaussian or even fat-tailed specifications, tend to underestimate the probability and magnitude of extreme downside movements. Extreme Value Theory provides a natural framework for modelling such tail events by focusing explicitly on exceedances beyond high thresholds. This distinction is economically important in our setting. As we demonstrate in subsequent sections, incorporating EVT-based tail behaviour materially increases predicted

mortgage default risk relative to baseline simulations, indicating that standard models understate the financial impact of climate risk.

We use the residuals from the baseline AR(3) house price model, fitting EVT to the tail of those residuals by climate-risk group, then generating two sets of residuals draws for simulation:

- a **baseline** set drawn from the central truncated distribution
- an **EVT-adjusted** set where the central part is preserved but the extreme tail is replaced by GPD-based draws beyond the threshold

Those residual draws are then fed back into the baseline HPI recursion to generate alternative price paths, which are later mapped into default probabilities. So, the key comparison is not EVT in isolation, but baseline-implied vs tail-adjusted price dynamics, and then the associated uplift in default risk.

4.3 Extreme Value Theory and Tail-Risk Simulation

While the baseline house price model captures the central dynamics of housing markets, it is not designed to represent rare but severe price movements that may arise in climate-exposed areas. Because our objective is to quantify the incremental effect of extreme climate risk, we apply Extreme Value Theory (EVT) to the residuals from the baseline house price model.

Our EVT implementation follows a peaks-over-threshold approach. Let $\hat{\epsilon}_{it}$ denote the residual from the baseline AR(3) model. We focus on the tail behaviour of these residuals and estimate a Generalized Pareto Distribution (GPD) for exceedances beyond a high threshold. The residuals are first stratified by climate-risk category, allowing the tail behaviour to vary across groups. This is important because climate vulnerability may not affect the average

evolution of house prices uniformly, but may instead manifest through infrequent and disproportionately large shocks.

Formally, for exceedances above a sufficiently high threshold (u), the conditional distribution of the excess $(y - u)$, where (y) denotes the tail variable of interest, is approximated by a Generalized Pareto Distribution:

$$G(y; \xi, \sigma) = 1 - \left(1 + \frac{\xi y}{\sigma}\right)^{-\frac{1}{\xi}}, y > 0$$

where ξ is the shape parameter and (σ) is the scale parameter. The shape parameter governs tail heaviness, while the scale parameter determines the dispersion of extreme realisations. We estimate these parameters separately for each climate-risk group using residuals from the baseline house price model.

The estimated GPD parameters are then used to construct simulated tail shocks. Our simulation design preserves the distinction between ordinary and extreme movements. For the central part of the residual distribution, we draw from the baseline residual process implied by the AR(3) model. For the tail region, however, we replace the baseline tail with EVT-implied draws generated from the fitted GPD. This produces an alternative residual process in which the central dynamics remain anchored to the historical model, but the extreme tail is allowed to reflect climate-related downside risk more realistically.

Using these residuals, we generate two sets of recursive house price paths. The first uses baseline residual draws only and therefore reflects the central distribution implied by the AR(3) model. The second incorporates EVT-adjusted tail draws and therefore embeds extreme downside events. Both sets of residuals are fed through the same baseline house price recursion, ensuring that the only difference between the two simulated paths is the treatment of tail risk.

This design is important for identification. By holding the baseline forecasting structure fixed and varying only the treatment of extreme residual realisations, we isolate the contribution of

tail risk to future house price declines. The resulting gap between baseline-implied and EVT-adjusted price paths provides the basis for our subsequent analysis of mortgage default risk.

4.4 Simulation Framework and Default Probability Estimation

To integrate the three modelling components and quantify the impact of climate-induced tail risk on mortgage default outcomes, we implement a structured simulation framework comprising the following steps.

1. Generation of house price paths

Using the estimated parameters from the baseline house price model and the EVT framework described in Sections 4.2 and 4.3, we simulate future house price dynamics for each geographic unit.

For each MSA, we generate 1,000 simulated monthly house price paths from October 2022 to October 2035 using a recursive AR(3) specification that includes an intercept, three lagged house price changes, and simulated residuals. To avoid imposing assumptions on future macroeconomic conditions, we exclude additional control variables from the simulation stage.

We construct two sets of simulations:

- **Baseline simulations**, in which residuals are drawn from the empirical distribution implied by the AR(3) model
- **EVT-adjusted simulations**, in which residuals are generated from a hybrid distribution that combines baseline draws for the central region with Generalized Pareto Distribution (GPD)-based draws for extreme tail events

This produces two comparable sets of simulated house price trajectories that differ only in their treatment of tail risk.

$$HPI_t^s = \hat{\alpha} + \hat{\beta}_1 HPI_{t-1}^s + \hat{\beta}_2 HPI_{t-2}^s + \hat{\beta}_3 HPI_{t-3}^s + \epsilon_t^s$$

2. Aggregation to state-level house price dynamics

Our loan-level dataset is available at the three-digit ZIP code level, which cannot be reliably mapped to MSAs. To ensure consistency between the house price simulations and the mortgage data, we aggregate the simulated MSA-level house price changes to the state level. For each state, we compute the average monthly house price change across MSAs. These simulated state-level house price changes serve as the primary time-varying input in the mortgage default model.

3. Mapping simulated house prices to default probabilities

Using the estimated mortgage default model (Section 4.6), we forecast monthly default probabilities for each loan over the simulation horizon.

In this step, we replace observed house price changes with simulated state-level HPI changes. The following variables are updated accordingly:

- house price growth (HPI change)
- a negative HPI indicator (equal to one when house price growth is negative)
- the interaction term capturing asymmetric effects of house price declines

All other covariates, including borrower and loan characteristics, are held fixed at their observed values. This approach isolates the impact of future house price dynamics on default risk.

4. Transmission of house price shocks

While loan-to-value (LTV) ratios are a key determinant of mortgage default, we do not dynamically update LTV using simulated house prices or amortization schedules. Instead, the effect of house price dynamics is captured through the inclusion of simulated HPI changes and their nonlinear transformations in the default model.

These variables act as the primary channel through which climate-related housing market stress affects loan performance, allowing us to isolate the impact of house price fluctuations without imposing additional structural assumptions.

5. Measurement of climate-induced default risk

For each loan, we compute cumulative default probabilities over the simulation horizon under both the baseline and EVT-adjusted scenarios. These probabilities are averaged across simulations to obtain expected default rates.

To quantify the impact of tail risk, we compare default probabilities implied by EVT-adjusted simulations with those generated under the baseline model. The difference between the two is defined as the **uplift**, which captures the incremental contribution of tail risk to mortgage default outcomes.

4.5 Mortgage Default Model

We model mortgage default probabilities using a reduced-form specification that relates default risk to house price dynamics and borrower characteristics.

The model takes the form:

$$log\left(\frac{P(Default_{i,t})}{1-P(Default_{i,t})}\right) = \alpha + \beta_1 LTV_{i,t} + \beta_2 FICO_i + \beta_3 DTI_i + \beta_4 X_i + \varepsilon_{i,t}$$

Where:

- $P(Default_{i,t})$ is the probability that loan i defaults within 3 years from origination
- $LTV_{i,t}$ is the loan-to-value ratio, which varies over time as house prices change
- $FICO_i$ is the borrower's credit score at origination, modeled using a piecewise linear specification with a cutpoint at 680
- DTI_i is the debt-to-income ratio at origination
- X_i is a vector of additional control variables including loan purpose (purchase vs. refinance), documentation level (full documentation indicator), and geographic controls

For the FICO score, we employ a piecewise linear specification with a cutpoint at 680 to capture the non-linear relationship between credit scores and default risk. This approach recognizes that marginal changes in FICO scores may have different implications depending on whether the borrower is above or below conventional credit quality thresholds.

The key explanatory variable is house price growth, which enters the model both linearly and through a nonlinear specification capturing asymmetric effects of price declines.

In particular, the model includes an indicator for negative house price growth and its interaction with the magnitude of price changes, allowing the effect of house price movements on default risk to differ between rising and declining markets. This specification reflects the well-established sensitivity of mortgage defaults to declines in house prices.

The model is estimated using loan-level data, with borrower and loan characteristics included as controls. In the simulation stage, these characteristics are held fixed at their observed values, while house price changes are replaced with simulated values generated from the baseline and EVT-adjusted models.

4.6 Measurement of Tail-Risk Uplift

To quantify the impact of tail risk on mortgage default outcomes, we construct a measure of **uplift**, defined as the difference between default probabilities implied by EVT-adjusted simulations and those generated under the baseline model.

For each loan (i), forecast horizon (t), and percentile level (p), uplift is defined as:

$$Uplift_{i,t,p} = PD_{i,t,p}^{EVT} - PD_{i,t,p}^{Baseline}$$

where $PD_{i,t,p}^{EVT}$ *and* $PD_{i,t,p}^{Baseline}$ denote the simulated default probabilities under the EVT-adjusted and baseline scenarios, respectively.

Aggregation across percentiles and time

Because our focus is on extreme downside risk, we evaluate default probabilities at high percentiles of the simulated distribution, specifically the 99th, 99.5th, and 99.75th percentiles. These percentiles capture increasingly severe house price declines and associated default outcomes.

For each loan and forecast horizon, we first compute the average uplift across percentiles:

$$\overline{Uplift_{\iota,t}} = \frac{1}{P}\sum Uplift_{i,t,p}$$

where (P) denotes the number of percentile levels considered.

We then aggregate across forecast horizons (2025, 2030, and 2035) to obtain an overall measure of uplift for each loan:

$$\overline{Uplift_{\iota}} = \frac{1}{T}\sum \overline{Uplift_{\iota,t}}$$

where (T) denotes the number of forecast horizons.

Geographic aggregation

To examine spatial variation in climate-related tail risk, we aggregate uplift measures across loans within geographic units. Given the structure of the data, we conduct this aggregation at the three-digit ZIP code level.

For each ZIP code (z), we compute the mean uplift across all loans in that region:

$$\overline{Uplift_{z}} = \frac{1}{N_z}\sum_{i\in z} \overline{Uplift_{\iota}}$$

where N_z denotes the number of loans in ZIP code (z).

These ZIP-level measures allow us to assess the consistency of uplift within and across regions, and to examine how the contribution of tail risk varies geographically.

Interpretation of uplift

By construction, uplift isolates the incremental contribution of tail risk to mortgage default outcomes. Because both baseline and EVT-adjusted simulations are generated using the same underlying house price model and differ only in the treatment of extreme residual realisations, uplift captures the extent to which standard models understate risk in the presence of extreme events.

A positive value of uplift indicates that incorporating tail risk increases predicted default probabilities, while a value close to zero suggests that the baseline model adequately captures the relevant distribution of shocks.

5. Results

5.1 Baseline House Price Dynamics

We begin by estimating the baseline house price model described in Section 4.2. The results in Table 1 below indicate that house price dynamics are largely driven by autoregressive components, with the AR(3) specification capturing the persistence in house price movements across regions. The estimated coefficients on lagged house price changes are economically large and statistically significant, consistent with strong momentum effects in housing markets.

Additional control variables contribute modestly to explanatory power, and the overall fit of the model suggests that it provides a reasonable approximation of the central tendency of house price dynamics. These findings are consistent with prior literature documenting the importance of autoregressive structure in modelling housing price movements.

Importantly, the baseline model is designed to capture typical fluctuations in house prices rather than extreme movements. As such, while it provides a suitable benchmark for forecasting average dynamics, it does not explicitly account for rare but potentially severe shocks. This limitation motivates the use of Extreme Value Theory in the subsequent analysis to model tail risk in house price changes.

For the subsequent simulation analysis, we combine the relatively low (RL) and relatively moderate (RM) categories into a single low-to-moderate (LM) group. This aggregation reflects the similar tail behaviour observed in these categories and allows for a more parsimonious and stable implementation of the simulation framework.

Table 1: House Price Index Model for Parameter Estimation

	(1)
	$\Delta HPI_{i,t}$
$\Delta HPI_{i,t-1}$	1.659***
	(0.000)
$\Delta HPI_{i,t-2}$	-0.705***
	(0.000)
$\Delta HPI_{i,t-3}$	0.00450
	(0.148)
RH	-0.000349***
	(0.000)
RL	-0.000126***
	(0.000)
RM	-0.000223***
	(0.000)
VH	-0.000306***
	(0.000)
RM * $\Delta HPI_{i,t-1}$	0.111***
	(0.000)
RM * $\Delta HPI_{i,t-2}$	-0.161***
	(0.000)
RM * $\Delta HPI_{i,t-3}$	0.0553***
	(0.000)
RH * $\Delta HPI_{i,t-1}$	0.150***
	(0.000)
RH * $\Delta HPI_{i,t-2}$	-0.178***

	(0.000)
RH * $\Delta HPI_{i,t-3}$	0.0359***
	(0.000)
RL * $\Delta HPI_{i,t-1}$	0.0380***
	(0.000)
RL * $\Delta HPI_{i,t-2}$	-0.0652***
	(0.000)
RL * $\Delta HPI_{i,t-3}$	0.0295***
	(0.000)
VH * $\Delta HPI_{i,t-1}$	0.0877***
	(0.000)
VH * $\Delta HPI_{i,t-2}$	-0.0384
	(0.103)
VH * $\Delta HPI_{i,t-3}$	-0.0395**
	(0.002)
STATE$\Delta HPI_{i,t-1}$	0.0267***
	(0.000)
	(0.894)
Constant	0.000734*
	(0.084)
Observations	300774
Adjusted R^2	0.995
MSA Fixed Effect	Included
MSA*STATE Effect	Included

*Notes: This table reports coefficient estimates from the panel regression of annualized MSA-level house price changes on lagged dependent variables, FEMA risk indicators, and their interactions. MSA and state fixed effects included. *, *, and *** denote significance at the 10%, 5%, and 1% levels, respectively.*

5.2 Extreme Value Analysis

While the baseline house price model captures the central dynamics of housing markets, it does not explicitly account for rare but severe price movements. To model such extreme

downside risk, we apply Extreme Value Theory (EVT) to the residuals of the baseline AR(3) model using a peaks-over-threshold (POT) framework.

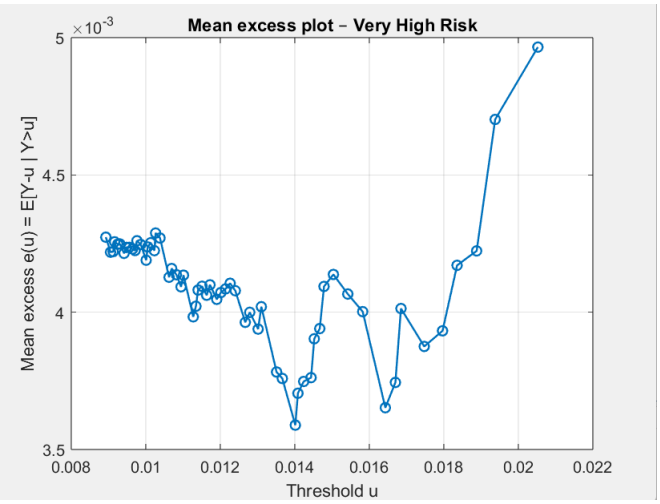
We focus on the negative tail of house price shocks by transforming residuals so that large values correspond to large negative price movements. EVT is then applied to the right tail of this transformed distribution, allowing us to isolate and model extreme downside events that are most relevant for mortgage default risk.

Threshold selection and model validation

The validity of the POT approach depends critically on the choice of threshold. To identify an appropriate threshold, we employ standard diagnostic tools, including mean excess plots and parameter stability analysis.

Figure 1 presents the mean excess function for candidate thresholds for the high-risk (VH) group. Consistent with EVT theory, we identify a range over which the mean excess function exhibits approximate linearity, indicating that the Generalized Pareto Distribution (GPD) provides a suitable approximation for exceedances beyond the threshold.

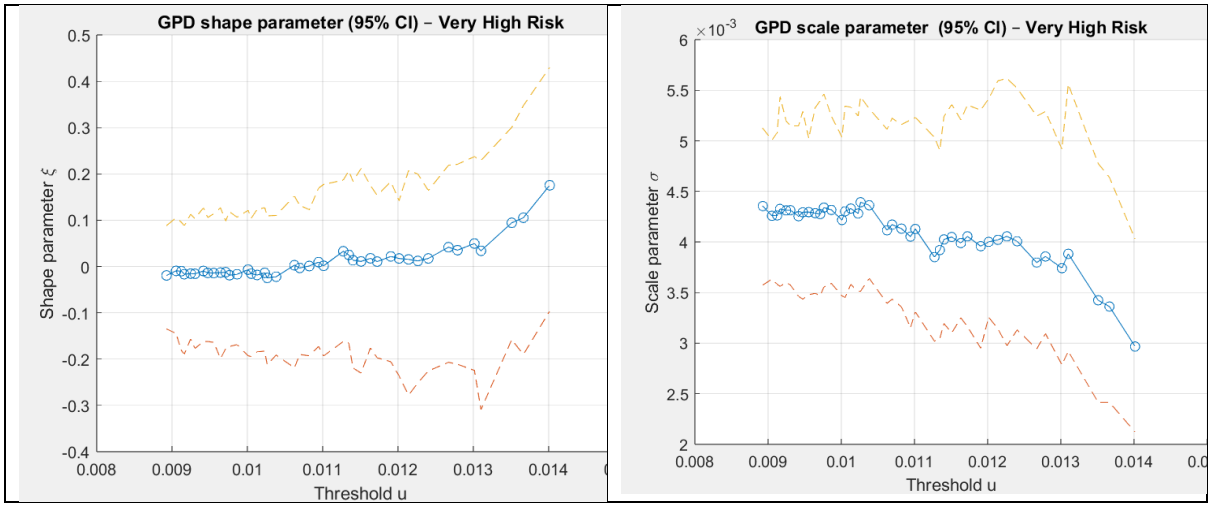
Figure 1: Mean Excess Plot of Very High Risk Group



To refine the selection, we examine the stability of the estimated GPD parameters across thresholds. **Figure 2** reports the estimated shape and scale parameters along with confidence

intervals. The chosen threshold corresponds to a region where both parameters are relatively stable, reflecting an appropriate balance between bias and variance.

Figure 2: 95% CI for the Shape and Scale Parameter for Very High Risk Group



GPD estimation

Using the selected threshold, we estimate the GPD parameters for each climate-risk group.

Table 2 reports the estimated threshold, shape, and scale parameters.

Table 2: Parameters for EVT Simulation

Scenario	Central Distribution	Threshold (u)	Shape (ξ)	Scale (σ)	Tail Source
VH	VH residuals	0.0092	−0.013292	0.004294	Estimated
RH	RH residuals	0.0120	0.048444	0.004572	Estimated
LM (baseline)	LM residuals	—	—	—	No EVT tail
LM + RH tail	LM residuals	0.0120	0.048444	0.004572	RH (imported)
LM + VH tail	LM residuals	0.0092	−0.013292	0.004294	VH (imported)

Notes: EVT parameters are estimated separately for the VH and RH groups using the Peaks-over-Threshold approach. For LM scenarios, we do not estimate separate tail parameters. Instead, LM simulations combine the LM central residual distribution with tail parameters imported from the VH and RH groups, respectively. This design allows us to isolate the effect of extreme tail shocks while preserving the underlying distributional characteristics of LM regions.

The results indicate meaningful differences in tail behaviour across climate-risk categories. High-risk groups exhibit heavier tails, as reflected in positive or near-zero shape parameters,

suggesting a higher likelihood of extreme negative house price shocks. In contrast, lower-risk groups display less pronounced tail behaviour in the residual diagnostics, consistent with thinner tails and a lower incidence of extreme outcomes.

These findings confirm that extreme house price movements are not uniformly distributed across regions, but are instead more pronounced in areas with higher climate exposure. While the baseline model provides a reasonable approximation of average price dynamics, it does not capture this heterogeneity in tail risk.

Given the weaker tail signatures in the LM group, we do not estimate separate EVT parameters for these regions. Instead, to assess the sensitivity of lower-risk areas to extreme shocks, we construct counterfactual LM scenarios by combining the LM central residual distribution with tail parameters estimated from the VH and RH groups. This approach allows us to isolate the incremental effect of extreme tail events while preserving the underlying distributional characteristics of LM regions, and provides a consistent basis for comparison across climate-risk categories.

The EVT framework therefore plays a critical role in our analysis by explicitly modelling the distribution of extreme shocks. In the next section, we incorporate these EVT-based tail estimates into simulated house price paths and examine their implications for mortgage default risk.

5.3 Mortgage Default Model

To translate simulated house price dynamics into mortgage risk outcomes, we employ the default model described in Section 4.6. This model relates the probability of default to

changes in house prices and borrower characteristics, with a particular emphasis on the asymmetric effects of house price declines.

In our simulation framework, observed house price changes are replaced with simulated values derived from the baseline and EVT-adjusted models. These simulated values enter the default model along with a negative house price growth indicator and an interaction term capturing nonlinear effects. All other borrower and loan characteristics are held fixed at their observed values.

This approach isolates the impact of future house price dynamics on default risk and allows us to assess how the incorporation of tail risk affects predicted mortgage performance. Table 3 shows the results of the mortgage default model.

Table 3: Logistic Regression Estimates of Mortgage Default Probability

VARIABLES	(1) Default
fico	-0.008**
	(0.004)
fico_680	21.725***
	(7.869)
fico_720	4.537
	(4.006)
fico_760	6.305
	(4.314)
int_fico680	-0.032***
	(0.012)
int_fico720	-0.006
	(0.005)
int_fico760	-0.008
	(0.006)
orig7	0.384**
	(0.160)
orig8	-0.414***
	(0.119)
purchase	-0.404***
	(0.060)
retail	-0.199***
	(0.054)
ch_hpi	-5.823***
	(0.643)
neg_ch_hpi	0.217***

	(0.084)
int_hpi_ch	-5.150***
	(0.741)
ltv	0.056***
	(0.005)
ltv_dummy	-0.310
	(0.594)
int_ltv	0.004
	(0.008)
CA	0.365***
	(0.082)
NY	0.656***
	(0.153)
CO	0.113
	(0.174)
WA	-0.410**
	(0.184)
TX	0.346***
	(0.123)
FL	0.597***
	(0.097)
Constant	-4.404
	(2.918)
Observations	270,439

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Logistic regression estimates of three-year mortgage default using Freddie Mac loan-level data. Default is defined as 90+ days delinquent within three years of origination. Standard errors in parentheses. All specifications include state and origination year fixed effects.

Below we show the ROC curves from our logistic regression. It measures the performance of the home loan default model. Figure A shows the ROC curve when price and associated variables are included in the default model and Figure B shows the ROC curve when price and its associated variables are not included in the default model. The area under the ROC curve is 0.78 for Figure A and 0.68 for Figure B. The area under the ROC curve is a measure of the performance of the logistic model. An AUC of 0.78 suggesting that the logistic regression model is able to distinguish between home loan defaults and non-defaults with a high degree

of accuracy and when HPI changes and associated variables are included in the model. To put it another way, if we were to randomly select 100 home loans, the logistic regression model would correctly predict whether or not each loan would default with an accuracy of 78%.

Figure 3: ROC Curve of the mortgage default model

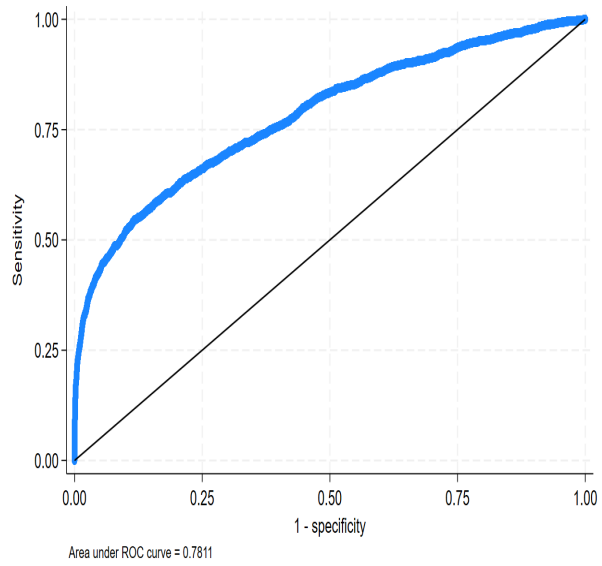


Figure A

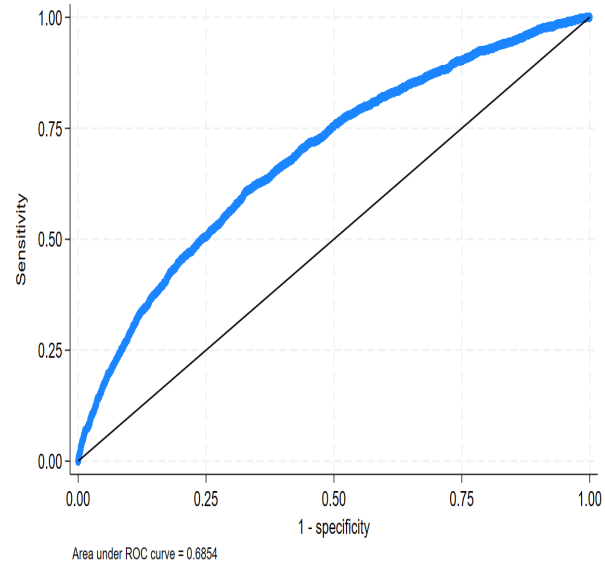


Figure B

Notes: Figure A and B show ROC curve for the logistic regression default model when change in HPI and its associated variables are included and nonincluded in the model. Area under the curve (AUC) is 0.78 for Figure A, indicating strong model performance in distinguishing between default and non-default outcomes.

5.4 Tail Risk and Default Uplift

To assess the impact of tail risk on mortgage outcomes, we compare default probabilities generated under the baseline and EVT-adjusted simulations. As defined in Section 3.6, uplift measures the incremental effect of incorporating tail risk into house price dynamics.

5.4.1 Heterogeneity across climate-risk groups

We begin by examining differences in uplift across climate-risk categories. In addition to the primary classification, we further partition the low-to-moderate (LM) group based on

exposure to higher-risk neighbouring regions, resulting in two subgroups: LM–VH and LM–RH. This partition allows us to capture heterogeneity within the LM category that may arise from spillover effects or proximity to higher-risk areas.

Table 4 reports the average uplift across forecast horizons (2025, 2030, and 2035) and tail percentiles (99th, 99.5th, and 99.75th). Several patterns emerge.

Table 4: Average uplift across forecast horizon and percentiles

LM Group				
VH				
	Percentile	EVT	Baseline	Uplift
2025	99	0.420	0.418	0.002
2025	99.5	0.432	0.427	0.005
2025	99.75	0.441	0.445	-0.004
2030	99	0.421	0.409	0.012
2030	99.5	0.435	0.414	0.020
2030	99.75	0.451	0.425	0.026
2035	99	0.421	0.412	0.009
2035	99.5	0.433	0.424	0.009
2035	99.75	0.449	0.429	0.020

LM Group				
RH	Percentile	EVT	Baseline	Uplift
2025	99	0.449	0.420	0.030
2025	99.5	0.456	0.436	0.021
2025	99.75	0.470	0.449	0.022
2030	99	0.451	0.423	0.027
2030	99.5	0.461	0.439	0.022
2030	99.75	0.480	0.448	0.032
2035	99	0.442	0.433	0.009
2035	99.5	0.470	0.456	0.014
2035	99.75	0.475	0.470	0.004

VH	Percentile	EVT	Baseline	Uplift
2025	99	0.383	0.388	-0.005
2025	99.5	0.408	0.406	0.002
2025	99.75	0.443	0.428	0.016
2030	99	0.397	0.411	-0.014
2030	99.5	0.424	0.421	0.003
2030	99.75	0.430	0.433	-0.003
2035	99	0.403	0.395	0.009
2035	99.5	0.419	0.411	0.009
2035	99.75	0.423	0.418	0.005

RH				
	Percentile	EVT	Baseline	Uplift
2025	99	0.457	0.452	0.006
2025	99.5	0.486	0.457	0.029
2025	99.75	0.495	0.476	0.020
2030	99	0.447	0.462	-0.015
2030	99.5	0.459	0.472	-0.013
2030	99.75	0.469	0.486	-0.017
2035	99	0.468	0.463	0.005
2035	99.5	0.484	0.499	-0.014
2035	99.75	0.492	0.517	-0.025

First, uplift is economically meaningful for the LM subgroups, particularly LM–RH, where the average uplift is highest. LM–VH also exhibits positive uplift, though of smaller

magnitude. These results suggest that incorporating tail risk materially increases predicted default probabilities in regions that are not classified as high-risk but are exposed to elevated risk through spatial or structural channels.

Second, the high-risk groups (VH and RH) display relatively small or near-zero average uplift. This indicates that, in these regions, baseline models already capture a substantial portion of the relevant risk, and EVT adjustments lead to comparatively smaller changes in predicted default outcomes.

5.4.2 Variation across forecast horizons

To examine the temporal evolution of tail risk, Table 5 reports average uplift separately for each forecast horizon.

Table 5: Average uplift across forecast horizon

	2025	2030	2035
LM + VH	0.001	0.020	0.013
LM + RH	0.024	0.027	0.009
VH	0.004	-0.005	0.008
RH	0.018	-0.015	-0.011

Uplift generally increases between 2025 and 2030 across most groups, suggesting that the effect of tail risk becomes more pronounced over the medium term. However, this pattern does not persist uniformly to 2035, where uplift stabilises or declines. This non-monotonic pattern indicates that the impact of tail risk is dynamic and may depend on the interaction between house price persistence and the timing of extreme shocks.

5.4.3 Behaviour across tail percentiles

Table 6 reports uplift across different tail percentiles. For most groups, uplift increases with the severity of the tail event, consistent with the idea that EVT adjustments primarily affect extreme outcomes.

Table 6: Average uplift across different percentiles

	99	99.5	99.75	Tail gradient (99.75-99)
LM + VH	0.008	0.011	0.014	0.006
LM + RH	0.022	0.019	0.019	-0.003
VH	-0.003	0.004	0.006	0.009
RH	-0.002	0.001	-0.007	-0.006

This pattern is particularly evident in the LM–VH group, where uplift rises steadily from the 99th to the 99.75th percentile. In contrast, the LM–RH group exhibits a relatively flat profile across percentiles, indicating that tail adjustments affect a broader range of extreme outcomes rather than only the most severe cases.

High-risk groups display more muted or irregular patterns across percentiles, consistent with the smaller average uplift observed earlier.

5.4.4 Geographic consistency

To assess the spatial consistency of these results, we aggregate uplift at the ZIP code level. Table 7A reports the mean uplift across ZIP codes, separately for California and Florida. The results show a high degree of consistency within each group, with relatively narrow ranges of uplift across ZIP codes. This suggests that the observed effects are systematic rather than driven by a small number of outlier regions. In particular, the LM subgroups exhibit stable and positive uplift across both states, reinforcing the conclusion that tail risk has a persistent and geographically widespread impact in these areas.

Table 7A: Average uplift across 3-digit zip codes

RH	
3 Digit Zip	Mean
339	0.0003
349	0.0003
906	0.0003

VH	

3 Digit Zip	Mean
331	0.0030
334	0.0030
900	0.0029
902	0.0028
907	0.0029
908	0.0029
910	0.0029
913	0.0030
914	0.0029
917	0.0029
919	0.0030
920	0.0030
921	0.0030
922	0.0029
923	0.0030
925	0.0029

	LM VH	LM RH
3 Digit Zip	Mean	Mean
320	0.0078	0.011
321	0.0080	0.012
322	0.0078	0.011
325	0.0077	0.011
327	0.0078	0.011
328	0.0073	0.010
329	0.0070	0.010
330	0.0075	0.011
333	0.0077	0.011
926	0.0075	0.011
927	0.0077	0.011
928	0.0075	0.011
930	0.0077	0.011
932	0.0075	0.011
935	0.0072	0.010
936	0.0076	0.011
937	0.0073	0.011
939	0.0080	0.011
940	0.0067	0.010
945	0.0075	0.011
946	0.0075	0.011
948	0.0076	0.011
949	0.0078	0.011

950	0.0071	0.010
951	0.0075	0.011
952	0.0071	0.016
953	0.0075	0.011
954	0.0067	0.010
956	0.0075	0.011
957	0.0076	0.011
958	0.0079	0.011
959	0.0072	0.0102

Table 7B: Average uplift across risk groups for CA and FL.

	CA	FL
LM + VH	0.007	0.008
LM + RH	0.011	0.011
VH	0.003	0.003
RH	0.000	0.000

5.4.5 Summary of findings

Taken together, the results indicate that incorporating tail risk into house price dynamics leads to economically meaningful increases in predicted mortgage default probabilities, particularly in regions with moderate but non-negligible climate exposure.

The effect of tail risk is not uniform across groups, time horizons, or percentiles. Instead, it varies systematically, with the strongest impact observed in the LM subgroups and more limited effects in high-risk regions. These findings highlight the importance of accounting for tail behaviour in assessing climate-related risks in housing markets.

An interesting and economically meaningful finding is that the estimated default impact is often stronger in moderate-risk regions than in the highest-risk areas. One explanation is that high-risk regions may already incorporate climate risk into house prices, lending standards, or borrower behaviour, thereby partially internalising the risk. In contrast, moderate-risk regions may not fully price or anticipate climate-related shocks, leading to larger adjustments when extreme events occur. This pattern is consistent with a nonlinear response to climate risk,

where under-recognised exposure can generate disproportionately large effects when tail events materialise.

5.4.6 Implications

Our findings have several implications for the measurement and management of climate-related risks in housing markets.

First, the results highlight the importance of modelling tail risk explicitly when assessing mortgage default risk. Standard house price models that focus on central tendencies may understate the impact of rare but severe price declines. By incorporating EVT-based adjustments, we show that extreme shocks can lead to economically meaningful increases in predicted default probabilities, particularly in regions with moderate but non-negligible climate exposure.

Second, the heterogeneity in uplift across climate-risk groups suggests that the effects of climate risk are not confined to areas classified as high risk. Instead, regions in the low-to-moderate category can experience substantial increases in default risk once tail behaviour is taken into account. This finding underscores the importance of considering indirect exposure and spillover effects when evaluating climate vulnerability in housing markets.

Third, the temporal and percentile-based variation in uplift indicates that climate risk is inherently dynamic and concentrated in extreme outcomes. The non-monotonic pattern over time suggests that the interaction between house price persistence and extreme shocks plays a key role in shaping default risk. Similarly, the stronger effects observed at higher percentiles highlight the relevance of tail events in driving adverse outcomes.

Finally, the consistency of uplift across ZIP codes within states suggests that the effects we document are systematic rather than driven by isolated local conditions. This reinforces the idea that climate-related tail risk represents a broad-based source of vulnerability in housing markets, rather than a purely idiosyncratic phenomenon.

Taken together, these findings suggest that incorporating tail risk into house price modelling frameworks can improve the measurement of mortgage risk and provide a more comprehensive assessment of climate-related financial exposure.

6. Conclusion

This paper examines the role of climate-related tail risk in shaping mortgage default outcomes. We develop a framework that combines a baseline house price model with Extreme Value Theory to capture extreme downside movements in housing markets, and map these dynamics into loan-level default probabilities.

Our results show that incorporating tail risk leads to economically meaningful increases in predicted default probabilities, particularly in regions with moderate levels of climate exposure. While standard models capture the central dynamics of house price movements, they do not fully account for the distribution of extreme shocks. The EVT-based approach provides a tractable way to incorporate such tail behaviour and quantify its implications for mortgage risk.

We further document that the impact of tail risk varies across climate-risk groups, forecast horizons, and percentiles of the loss distribution. In addition, the effects are broadly consistent across geographic units, suggesting that the contribution of tail risk is systematic rather than driven by isolated observations.

Overall, our findings highlight the importance of moving beyond average dynamics when assessing climate-related risks in housing markets. Accounting for tail behaviour can lead to a more accurate and comprehensive evaluation of mortgage default risk, particularly in environments where extreme events play an increasingly important role.

Our findings have implications for how climate risk is measured and incorporated into mortgage risk assessment frameworks. By showing that tail events play a central role in

driving default risk, the analysis highlights the importance of modelling extreme outcomes rather than relying solely on average price dynamics.

References

Baldauf, M., Garlappi, L., & Yannelis, C. (2020). Does climate change affect real estate prices? Only if you believe in it. The Review of Financial Studies, 33(3), 1256-1295.

Begley, T.A., Gurun, U.G., Purnanandam, A. and Weagley, D., 2024. Disaster lending:“fair” prices but “unfair” access. *Management Science*, *70*(12), pp.8484-8505.

Bernstein, A., Gustafson, M. T., & Lewis, R. (2019). Disaster on the horizon: The price effect of sea level rise. Journal of Financial Economics, 134(2), 253-272.

Bin, O., & Landry, C. E. (2013). Changes in implicit flood risk premiums: Empirical evidence from the housing market. Journal of Environmental Economics and Management, 65(3), 361-376.

Campbell, J. Y., & Cocco, J. F. (2015). A model of mortgage default. The Journal of Finance, 70(4), 1495-1554.

Case, K. E., & Shiller, R. J. (1989). The efficiency of the market for single-family homes. The American Economic Review, 79(1), 125-137.

Davis, M. A., & Quintin, E. (2017). On the nature of self-assessed house prices. Real Estate Economics, 45(3), 628-649.

Davis, R., & Resnick, S. I. (1984). Tail estimates motivated by extreme value theory. Annals of Statistics, 12, 1467-1487.

Dekkers, A. L., & De Haan, L. (1989). On the estimation of extreme-value index and large quantile estimation. Annals of Statistics, 17, 1795-1832.

Embrechts, P., Kluppelberg, C., & Mikosh, T. (1997). Modelling extremal events for insurance and finance. Springer Verlag, Berlin.

Foote, C. L., Gerardi, K., & Willen, P. S. (2008). Negative equity and foreclosure: Theory and evidence. Journal of Urban Economics, 64(2), 234-245.

Gerardi, K., Herkenhoff, K. F., Ohanian, L. E., & Willen, P. S. (2018). Can't pay or won't pay? Unemployment, negative equity, and strategic default. The Review of Financial Studies, 31(3), 1098-1131.

Gibson, M., Mullins, J. T., & Hill, A. (2021). Climate risk and beliefs: Evidence from property markets. Working Paper.

Glaeser, E. L., & Nathanson, C. G. (2017). An extrapolative model of house price dynamics. Journal of Financial Economics, 126(1), 147-170.

Issler, P., Stanton, R., Vergara-Alert, C., & Wallace, N. (2020). Mortgage markets with climate-change risk: Evidence from wildfires in California. Working Paper.

Keenan, J. M., & Bradt, J. T. (2020). Underwaterwriting: From theory to empiricism in regional mortgage markets in the U.S. Climatic Change, 162, 2043-2067.

Keys, B. J., & Mulder, P. (2020). Neglected no more: Housing markets, mortgage lending, and sea level rise. Working Paper, National Bureau of Economic Research.

Kremer, E., 1986. Simple formulas for the premiums of the LCR and ECOMOR treaties under exponential claims sizes. *Blätter der DGVFM*, *17*(3), pp.237-243.

Leadbetter, M. R., Lindgren, G., & Rootzen, H. (1983). Extremes and related properties of random sequences and processes. Springer Verlag, New York.

Martins-Filho, C., & Yao, F. (2006). Estimation of Value-at-risk and expected shortfall based on nonlinear models of return dynamics and extreme value theory. Studies in Nonlinear Dynamics & Econometrics, 10(2), Article 4.

Martins-Filho, C., Yao, F., & Torero, M. (2018). Nonparametric estimation of conditional value-at-risk and expected shortfall based on extreme value theory. Econometric Theory, 34(1), 23-67.

Nguyen, D.D., Ongena, S., Qi, S. and Sila, V., 2022. Climate change risk and the cost of mortgage credit. *Review of Finance*, *26*(6), pp.1509-1549.

Ortega, F., & Taspinar, S. (2018). Rising sea levels and sinking property values: Hurricane Sandy and New York's housing market. Journal of Urban Economics, 106, 81-100.

Ouazad, A., & Kahn, M. E. (2021). Mortgage finance in the face of rising climate risk. Working Paper, National Bureau of Economic Research.

Pickands, J. (1975). Statistical inference using extreme order statistics. Annals of Statistics, 3, 119-131.

Reiss, R. D. (1987). Estimating the tail index of the claim size distribution. Blatter der DGVM, 18, 21-25.

Resnick, S. I. (1987). Extreme values, regular variation and point processes. Springer Verlag, New York.

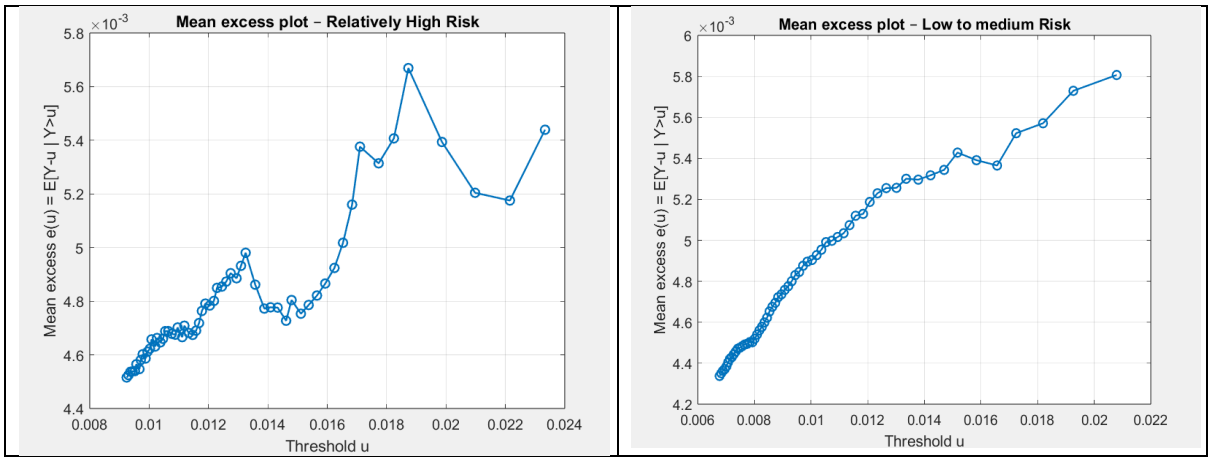
Smith, R. L. (1987). Estimating tails of probability distributions. Annals of Statistics, 15, 1174-1207.

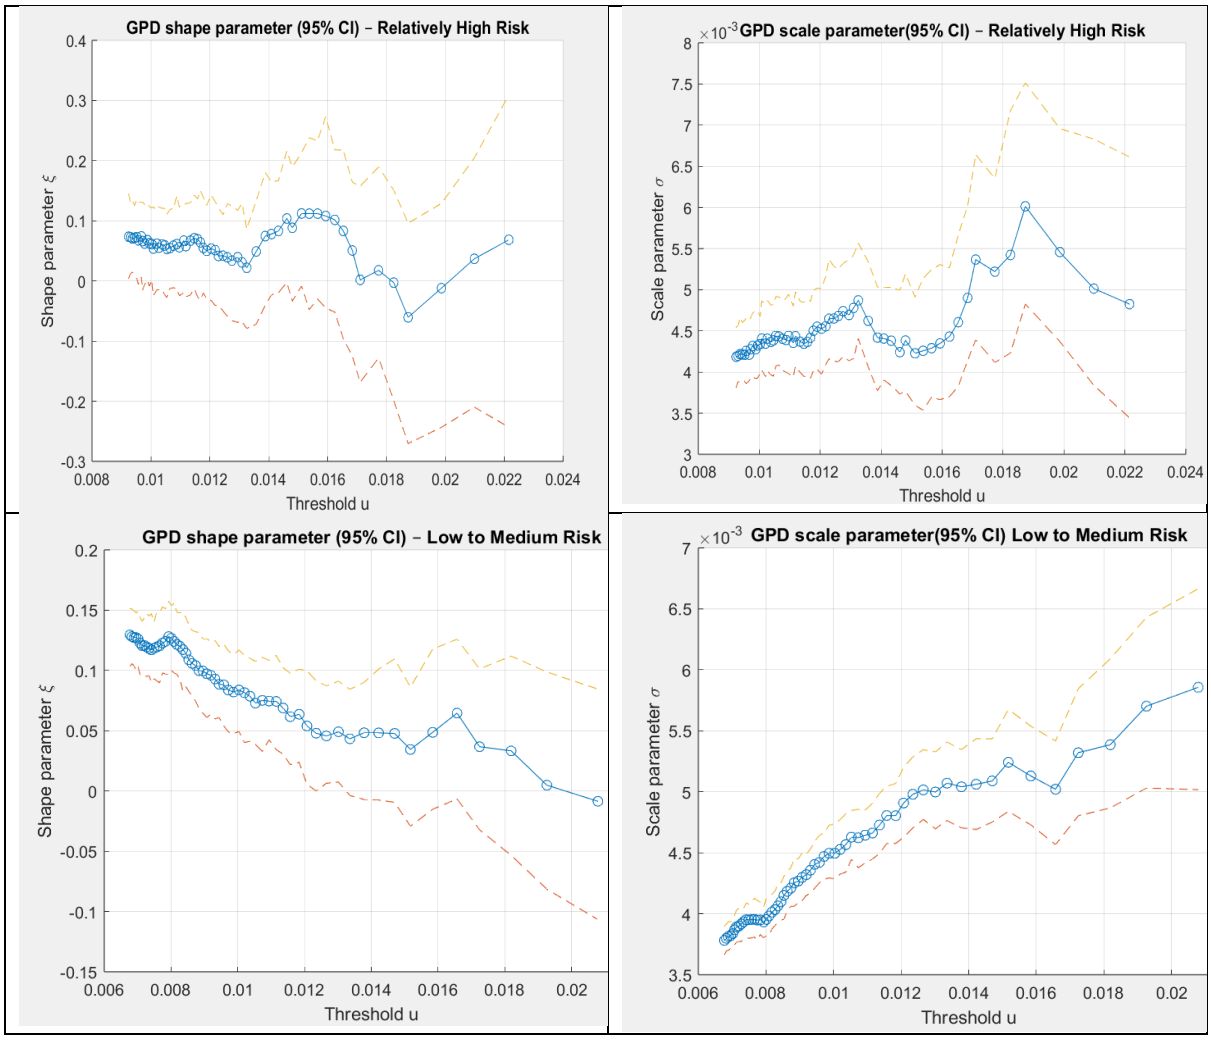
Smith, R. L. (1990). Extreme Value Theory. In: Ledermann, W., Handbook of Applicable Mathematics, Supplement. Wiley, Chichester, pp. 437-472.

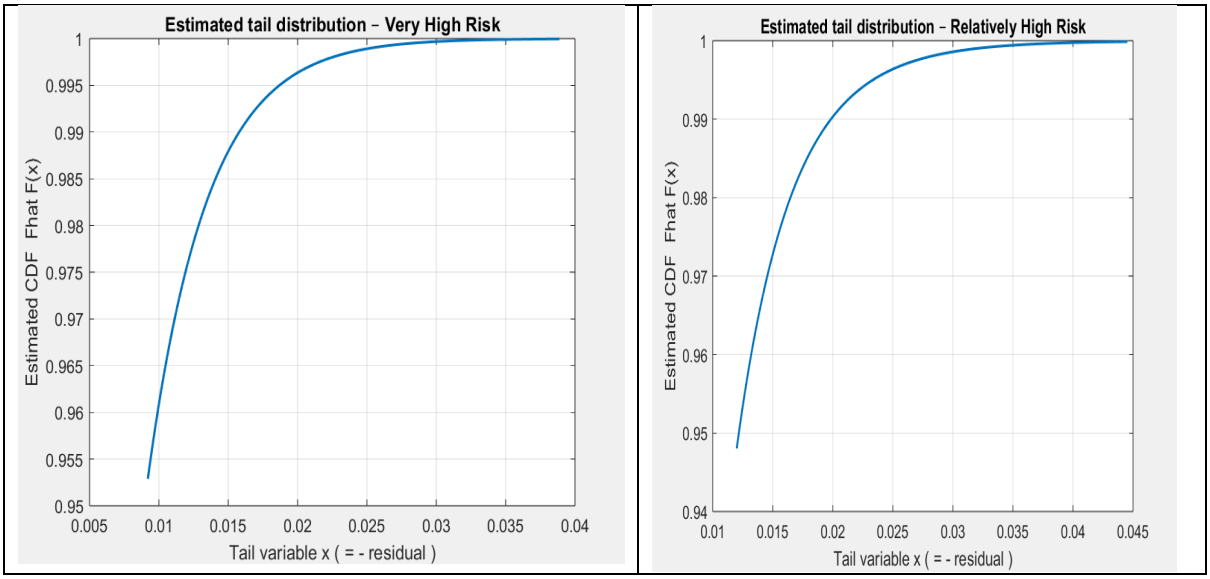
Teugels, J. L. (1981). Remarks on large claims. Bull. Inst. Intern. Statist., 49, 1490-1500.

Teugels, J. L. (1984). Extreme values in insurance mathematics. In: Statistical extremes and applications. Proceedings of the International Conference at Vimeiro. Reidel, Dordrecht, pp. 252-259.

Appendix A


Mean excess plot – Relatively High Risk
×10^-3
Mean excess e(u) = E[Y-u | Y>u]
Threshold u
Mean excess plot – Low to medium Risk
×10^-3
Mean excess e(u) = E[Y-u | Y>u]
Threshold u


GPD shape parameter (95% CI) – Relatively High Risk
Shape parameter ξ
Threshold u
×10^-3 GPD scale parameter(95% CI) – Relatively High Risk
Scale parameter σ
Threshold u
GPD shape parameter (95% CI) – Low to Medium Risk
Shape parameter ξ
Threshold u
×10^-3 GPD scale parameter(95% CI) Low to Medium Risk
Scale parameter σ
Threshold u


Estimated tail distribution – Very High Risk
Estimated CDF Fhat F(x)
Tail variable x (= - residual)
1
0.995
0.99
0.985
0.98
0.975
0.97
0.965
0.96
0.955
0.95
0.005
0.01
0.015
0.02
0.025
0.03
0.035
0.04
Estimated tail distribution – Relatively High Risk
Estimated CDF Fhat F(x)
Tail variable x (= - residual)
1
0.99
0.98
0.97
0.96
0.95
0.94
0.01
0.015
0.02
0.025
0.03
0.035
0.04
0.045