Mapping the Protection Gap: How Risk Perception and Social Vulnerability Drive Flood Memory Patterns in the U.S.

Eric Contreras, Nadja Veigel[a,b,c], Jacek Suda[d,e], Heidi Kreibich[b] and Andrea Cominola[a,c,*]

[a]Chair of Smart Water Networks, Technische Universität Berlin, Straße des 17. Juni 135 Berlin, 10623, Germany
[b]Section 4.4 Hydrology, GFZ Helmholtz Centre for Geosciences, Telegrafenberg Potsdam, 14473, Germany
[c]Einstein Center Digital Future, Wilhelmstraße 67 Berlin, 10117, Germany
[d]Department of Quantitative Economics, SGH Warsaw School of Economics, al. Niepodległości 162 Warsaw, 02-554, Poland
[e]Directorate General Monetary Policy, European Central Bank, Sonnemannstrasse 20 Frankfurt am Main, 60314, Germany

ARTICLE INFO

Keywords:
Floods
Flood Insurance
Change Point Detection
National Risk Index
Social Memory

ABSTRACT

This study investigates the temporal dynamics of flood insurance adoption patterns across U.S. counties by analyzing National Flood Insurance Program data through the lens of social memory and risk perception. Using change point detection methodologies, we identify significant shifts in insurance purchasing behavior following flood events and quantify both the magnitude (*salience*) and memory (*time-to-forget*) of these post-flood responses. We show that social memory of flood events, as measured through insurance participation rates, may be considerably shorter and more heterogeneous than previously suggested in the literature. By including demographic, environmental, and institutional variables, we disentangle the multi-modal distribution of salience and memory. We identify three components driving insurance adoption: social vulnerability, risk perception, and flood damage patterns. Subsequent cluster analysis of these components reveals five distinct county profiles, with particularly notable findings regarding areas where social vulnerability coincides with low risk perception. The results demonstrate spatiotemporal variations in community responses to flood risks and suggest that only diverse approaches can successfully maintain equitable sustained insurance coverage.

1. Introduction

Natural disasters, particularly floods, pose significant challenges globally, causing loss of life, displacement, and widespread infrastructure damage (Wing et al., 2018, 2020). Flooding stands out as the most frequently occurring natural hazard, with rising risks and unequal exposure expected due to climate change (Wing et al., 2022). Even as the threat and financial impact of flooding grow, recent estimates indicate that between two-thirds and 92% of flood-exposed properties are currently uninsured (Choi et al., 2024; CBO, 2024).

The United States National Flood Insurance Program (NFIP) is a federal government program that provides flood insurance to homeowners, renters, and businesses in participating communities. The program was created by the United States Congress in 1968 to address the lack of flood insurance options in the private market. The primary goal was to reduce the financial burden on the government from post-flood disaster relief and to internalize the costs of floodplain occupation through flood insurance with risk-based premiums. This approach aimed to promote the economically

*Corresponding author
andrea.cominola@tu-berlin.de (A. Cominola)
ORCID(s): 0009-0004-0263-5410 (E. Contreras); 0000-0002-8044-2751 (N. Veigel); 0000-0002-8259-533X (J. Suda); 0000-0001-6274-3625 (H. Kreibich); 0000-0002-4031-4704 (A. Cominola)

efficient use of floodplains while encouraging floodplain management measures for communities (Bin and Landry, 2013).

The NFIP program faced initial challenges with low community enrollment levels and low insurance purchase rates by individual property owners in participating communities (Michel-Kerjan and Kousky, 2010). Subsequent legislation was enacted, establishing that property owners residing in 100-year floodplains (also called Special Flood Hazard Areas, or SFHAs) are bound by federal law to obtain flood insurance if they have a mortgage from a federally regulated or backed lender. Regulations also mandate community enrollment as a pre-condition for qualifying for federal disaster assistance on a communal and individual levels. These measures increased the program's impact, with an estimated 22,000 communities benefiting from NFIP coverage nationwide (Kousky, 2017). Despite these efforts, the current level of flood insurance coverage of the NFIP is insufficient compared to the theoretical social optimum Gallagher (2014); Kunreuther et al. (2009); Kriesel and Landry (2004). The insurance program's low effectiveness has motivated a growing number of researchers to analyze the drivers behind flood insurance demand, finding hazard proximity (Kousky, 2010; Zahran et al., 2009; Bin et al., 2008), risk perception (Cannon et al., 2020; Bin and Landry, 2013; Lindell and Hwang, 2008) and demographic characteristics (Lucas et al., 2021; Wang et al., 2017; Landry and Jahan-Parvar, 2011; Michel-Kerjan and Kousky, 2010; Lindell and Hwang, 2008; Browne and Hoyt, 2000) as significant factors. There is also general agreement on the price inelasticity of insurance (Atreya et al., 2015; Landry and Jahan-Parvar, 2011; Kriesel and Landry, 2004).

Several papers show that in the US insurance is purchased reactively after a flood (Kousky, 2017; Veigel et al., 2023), which is linked to increased risk perception and an overestimation of the probability of a flood shortly after its occurrence (Dumm et al., 2020). Studies characterize both the magnitude (*salience*), measured by the increase in the take-up rate, and the duration (*memory* or *time-to-forget*), which refers to how long the elevated take-up rates persist following the flood. Kousky (2017) finds that communities that suffered a flood in the previous year increase their net flood insurance purchases by 6.7% and 7.2%, depending on whether a Presidential Disaster Declaration was issued. According to her research, this effect disappears three years after the storm. Atreya et al. (2015) finds similar results, stating that the temporary increase in flood insurance purchases after a flood event fades after three years. Studies observe an immediate increase in the number of homeowners who own flood insurance in communities affected by floods. According to Gallagher (2014) and Choi et al. (2024), the effect ranges from 7%-9% and then declines steadily, becoming undistinguishable from zero nine years after the event. Similarly, while studying the housing market with hedonistic models that explicitly incorporate linear and nonlinear temporal flood-zone effects, Atreya et al. (2013) finds that the flood risk discount disappeared between four and nine years after the flood.

While these findings, derived from regression analyses, provide point estimates with confidence intervals, in many cases they rely on the assumption of normally distributed salience and memory effects. Unlike regression-based approaches that assume parametric distributions, CPD can be implemented with non-parametric methods that detect structural breaks in time series data without imposing distributional assumptions on the underlying process. To advance our understanding of these temporal dynamics, we implement a change point detection algorithm (CPD) that allows for an independent quantification of changes in flood insurance coverage for each county. Earlier regression-based studies have highlighted the need to better understand how baseline risk perception and household adaptive capacity influence post-disaster risk perception and subsequent insurance demand. By characterizing these heterogeneous distributions of temporal changes across different risk and capacity profiles, we establish county-specific baselines that capture the varying risk perceptions that influence post-disaster insurance behavior. Our methodological approach addresses this gap by coupling the CPD results with unsupervised clustering of National Risk Index variables, enabling us to characterize the heterogeneous distribution of temporal changes in insurance adoption patterns.

We analyze over 50 million insurance policies across the Contiguous United States (CONUS) from 2010 to 2021, focusing on the following research questions: (1) how big is the change on insurance market dynamics after a flood event, measured as a rise in insurance take-up rates, (2) how long does the effect on the insurance market last, following a flood event until the market returns to initial levels, and (3) which observable characteristics from the counties under study can explain the variability in memory and salience effects after a flood event. We incorporate the National Risk Index to address the differences in the flood memory and salience resulting from the CPD analysis to uncover patterns in memory and salience effects. By correlating risk index components with different aspects of flood memory and insurance adoption patterns, we show how social vulnerability, risk perceptions, and flood damage correlate with insurance coverage. This approach helps identify how strategies to counteract the insurance protection gap can be tailored to the specific characteristics of underprotected communities.

2. Materials and Methods

In this study, we develop a data-driven approach that, in its first phase, analyzes insurance purchase data from the U.S. National Flood Insurance Program (FEMA, 2023c) in combination with household information from the U.S. Census Bureau (U.S. Census Bureau, 2022) to identify behavioral patterns of flood insurance purchase through CPD. CPD is an ensemble of unsupervised learning techniques that identify critical points in a time series or data sequence where the underlying data distribution changes (van den Burg and Williams, 2022). Those are called change points. We compare the timing of the identified change points with the timing of occurrences of storms and flood events, recorded in data from the U.S. National Oceanic and Atmospheric Administration (NOAA, 2023). This comparison allows validating the existence of after-flood salience effects and analyzing their impact, in terms of flood insurance take-up change, and duration. A flowchart summarizing the steps included in our analysis is shown in Figure 1. After building our input dataset (*Panel building* in the figure) and performing change point detection (*Identifying change*

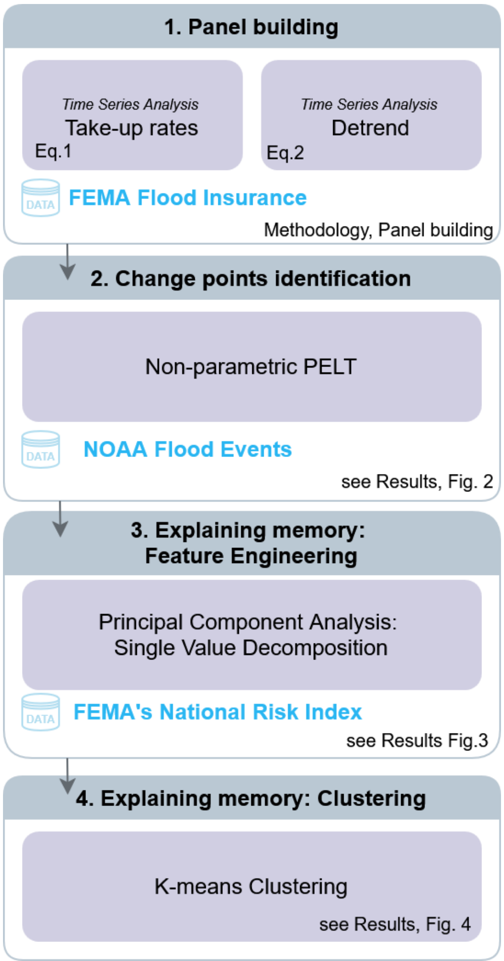

Figure 1: Flowchart summarizing the sequential analysis steps to identify and characterize flood insurance purchasing patterns. Our analysis relies on data from the U.S. National Flood Insurance Program and U.S. Census Bureau. The first two steps (*Panel building* and *Identifying change points*) investigate temporal changes in insurance adoption through CPD analysis, validated against NOAA storm and flood event data to identify post-flood salience effects. The bottom panels (*Explaining memory: Features* and *Explaining memory: Model*) include feature engineering and cluster analysis incorporating demographic data, NFIP Community Rating System participation, and FEMA's National Risk Index to identify key predictors of insurance purchasing behavior. Principal Component Analysis and K-means cluster analysis characterize distinct patterns in insurance adoption and social memory effects across different US counties. The data source of each step is reported in the figure, along with a reference to the results relative to each step.

points), we perform feature engineering via dimensionality reduction (*Explaining memory: Features*) to recognize a relevant subset of potential predictors from a combination of variables including demographic information from the U.S. Census Bureau (U.S. Census Bureau, 2018), NFIP participation information from the Community Rating System (CRS) (FEMA, 2021), and natural and social risk information from FEMA's National Risk Index (FEMA, 2023a). We identify the principal components (PCs) affecting salience and memory effects, and finally group the counties under analysis into clusters with similar insurance purchase and social memory behavior (*Explaining memory: Clustering*. The following sections describe each step of our analysis (Boxes 1-4 in Figure 1) in detail.

2.1. Panel building

We analyze household flood insurance purchase data from the United States NFIP, a federal government initiative tasked with providing affordable flood insurance to homeowners, renters, and businesses in participating communities. The NFIP program maintains a public data source containing over 80 million records of household flood insurance purchases since 1970 (FEMA, 2023c). For this study, we use data comprising over *50* millions of records, covering the period 2010-2021.

We pre-process the data as follows, before feeding them to step 2 of our analysis. First, we spatially and temporally aggregate the insurance purchase data by adding them up at a county level and with a monthly time-step. This decision strikes a balance between managing the computational complexity of a country-wide high-resolution analysis and capturing enough data to have a detailed view of the salience and duration effects. Second, we rescale the insurance data by the number of households per county from the Federal Census Bureau, and subsequently define our take-up rate variable $y_{c,t}$ as the net proportion of households that purchased or renewed their flood insurance policies in a given month t, for a given county c.

We thus build a panel of insurance take-up rate signals defined as

$$y_{c,t} = \frac{p_{c,t}}{h_{c,t}}, \tag{1}$$

where $p_{c,t}$ is the number of flood insurance policies purchased or renewed in county c during month t, and $h_{c,t}$ the estimated number of households in county c in month t. As the household information is provided by the U.S. Census Bureau every 10 years, we fit a linear interpolation to estimate the monthly number of households per county.

Finally, to remove trends from the time series signals, we substract the take-up rate at time t from the closest previous element in the series. This results in the detrended take-up rate $d_{c,t}$, formulated as

$$d_{c,t} = y_{c,t} - y_{c,t-1}. \tag{2}$$

We analyse these variables as time series and use change point detection to find relevant changes in their distributions. We define the salience as the spike in the insurance take-up—expected to happen shortly after a flood due to updated risk perception—expressed in percents, and the memory effect as the time it takes for the insurance coverage to drop back to pre-flood levels, expressed in months. We refer to the time to drop back to pre-flood levels as the "memory effect" and the "time to forget" interchangeably.

Since the objective of our study is to analyze the salience effect on the insurance market after a flood event, we set minimum thresholds on both the insurance take-up rate at t_0, i.e., prior to a flood event, at 0.2%) and the damage to infrastructures due to flood events at USD $6,000. These thresholds account for 99% of the infrastructural damage in the NOAA database and 99% of the total insurance policies active during the study period. With those thresholds in place, our study sample consists of a panel of 1.678 counties in the contiguous U.S., corresponding to 53% of all US counties. The detrended time series of insurance take-up rate for each of retained counties comprises 144 data points recorded between January 2010 and December 2021.

2.2. Change points identification

Owing to the high dimensionality of our problem and the need for a tool that is sensitive to short-term variations, we use CPD to analyze the detrended time series of flood insurance take-up rate obtained in the above data gathering and pre-processing step. CPD is the task of (i) statistically determining if there are distribution changes within a signal or time series and (ii) pinpointing the exact locations where changes in distribution occur (Harchaoui et al., 2008). It can be seen as a segmentation exercise, where a full signal y_t is split into segments of similar distribution, delimited by a set of indexes aptly named change points (van den Burg and Williams, 2020). CPD methods are flexible and computationally light, allowing us to test a broad palette of algorithms to find the most suitable one for our sample. The optimization problem associated with CPD can be generalized as

$$\min_{\tau} \sum_{i=1}^{n-1} C(y_{\tau_i,\tau_{i+1}}) + \lambda P(n) \tag{3}$$

where τ represents a segmentation of the signal, defined by a set of change points $\{\tau_1, ..., \tau_n\}$. The function $C(y_{\tau,\tau+1})$ represents the cost for each segment of the signal, a measure of disagreement between the data points and a proposed distribution, with negative maximum log-likelihood being one of the most common. $P(n)$ represents a penalty function for the number of segments, and λ stands for the set of hyperparameters used to calibrate said penalty (different methods define λ as a single value or as a complex set of parameters). The risk of overfitting by segmenting the signal into too many parts is balanced by the penalty value and further calibrated by the set of hyperparameters λ. In this way, C, P, and λ become the three fundamental components of any CPD method.

Taking into account previous literature on the comparative performance of known CPD methods (Truong et al., 2020; van den Burg and Williams, 2020; Aminikhanghahi and Cook, 2017), we test a battery of six CPD methods with twenty combinations of C, P, and λ parameters. We include both parametric approaches - which generally assume normal distribution and independence of the variables in the signal - and non-parametric approaches - which are able to segment a signal with an unknown distribution - in our analysis: Binary Segmentation [BINSEG] (Scott and Knott, 1974) and Pruned Exact Linear Time [PELT] (Killick et al., 2012) as parametric approaches, and Non-parametric PELT (Haynes, Kaylea et al., 2016), Energy Change Point Agglomerative (Matteson and James, 2013), Energy Change Point Divisive (Matteson and James, 2013), and Energy Change Point Divisive on Medians [EDM] (James et al., 2016) as non parametric methods. For more details on the methods refer to Supporting Notes - Section 1.3. We then compare the detected change points with the actual flood events data from NOAA (NOAA, 2023). For this comparison, we define a 15-months window after a flood event, in which a change point is considered a true positive. The true positive window is set in that range to capture short-term effects, but considering that flood insurance policyholders tend to let their year-long policies lapse instead of canceling them in mid-duration (Michel-Kerjan et al., 2012). Detailed explanations and mathematical formulations of the selected metrics are described in the Supporting Notes - Section 1.1.

2.3. Explaining memory: Feature Engineering

Following the identification of change points with CPD, we implement Principal Component Analysis (PCA) to reduce the dimensionality of multiple potential flood-related determinants of flood insurance adoption patterns into interpretable components (Jolliffe, 2002). PCA transforms the original n-dimensional data through eigendecomposition of the covariance matrix Σ:

$$\Sigma = V \Lambda V^T \tag{4}$$

where V contains the eigenvectors (principal components (PCs)) and Λ contains the eigenvalues representing the variance explained by each PC. Each PC is constructed as a linear combination of the standardized original variables:

$$PC_i = \sum_{j=1}^{n} w_{ij} x_j \tag{5}$$

where w_{ij} represents the loading of variable j on principal component i, and x_j is the standardized transformation of the original j-th variable. As an output of this dimensionality reduction step, we obtain the set of numerical weights, representing the relative contribution of each input variable on the resulting principal components. The input variables are CRS participation, risk score, yearly average events, average yearly damage, class rating, initial Take-up, water cover, resilience score, disadvantaged population amd social vulnerability score. They were extracted from the National Risk Index data (FEMA, 2023b).

2.4. Explaining memory: Clustering

In order to identify distinct patterns in county-level flood response characteristics, we implement the k-prototypes clustering algorithm (Huang, 1998). Proposed as an extension of k-means clustering (Lloyd, 1982) k-prototypes clustering is built with capabilities to handle both numerical and categorical variables. We base the clustering process on our objective variables *salience* and *memory*, our dependent variables transformed into *principal component scores*, and a categorical variable *coastal* that indicates whether a county is located on the shoreline. This last element provides key inputs for the clustering process, as coastal and in-land counties are subject to significantly different risks, contexts, and policies when it comes to floods. The k-prototypes algorithm partitions counties into clusters by iteratively minimizing both the within-cluster sum of squared Euclidean distances between each numerical observation and its assigned cluster centroid, and also the matching dissimilarity measure on the categorical attributes. A weight γ is used to avoid favoring either type of attribute. We conducted an iterative random search comparing gamma values in the range from 5 to 5000. We selected the best fit based on how well categorical values (most importantly coastal/inland indicators) were represented in our results, determining a γ value of 1000. The function to minimize can thus be expressed as:

$$\arg\min_{C} \sum_{i=1}^{k} \left(\sum_{x \in C_i} \sum_{j=1}^{p} (x_j - \mu_{i,j})^2 + \gamma \sum_{x \in C_i} \sum_{j=p+1}^{m} \delta(x_j, \mu_{i,j}) \right) \tag{6}$$

where C_i represents cluster i, x is a vector representing a county's position in the three-dimensional PC space (p numerical values) and categorical space ($m - p$ categorical values), and $\mu_{i,j}$ is the centroid of cluster i.

3. Results

The set of counties shown in Figure 2 (top map) represents the working sample of counties retained after establishing the minimum thresholds for damage and take-up rate. This working sample of 1678 counties covers all 48 states in the Contiguous U.S. It contains representative samples from coastal and inland territories alike, all nine groups of climate in CONUS, and different levels of flood risk as measured by the SFHA risk indicator.

3.1. Overall trend in policies

While the number of housing units in all counties appears reasonably stable, with a median of 1% growth over the ten years of study, the median number of active flood insurance policies dropped by 34% in a decade, dragging the median take-up rate to a similar decrease of 36% during our sample period (see Equation 1 for the definition of the research variable). The existence of this downward trend in the monthly take-up rate is counter-intuitive as the underlying risk of flood in the US is increasing due to the climate change emergency (Wing et al., 2022). This indicates that households in their flood insurance decisions are not considering the long-term increase in the underlying flood hazard or that they may be affected by short-term affordability constraints.

3.2. CPD Results

The metrics and performance evaluation of the CPD method is described in detail in the Supporting Notes - Section 1.1 The Recall metric represents the number of simulated change points that match with a change point in the ground truth data. We defined Non-parametric PELT (Killick et al., 2012) as the best-fit method for our study, based on a

Table 1
Descriptive statistics of CPD results and sample of US counties considered for CPD in our analysis.

Descriptive statistic	Value	Relative Percentage
Total Counties in sample (n)	1,678	100.00%
Counties with positive recall (n)	947	56.44%
Counties with salience effect (n)	756	45.05%
Salience mean (%)	7.92	
Salience median (%)	2.30	
Salience std. dev. (%)	22.54	
Memory mean (months)	28.03	
Memory median (months)	18.00	
Memory std. dev. (months)	25.92	

combination of performance metrics including Recall, Precision, Annotation error, and Hausdorff distance. The full results of the CPD methods performance test can be found in Supporting Information Figure 1. In our specific case, the PELT CPD method's Recall indicates that 24.5% of the change points in the take-up rate time series overlays with damage-inducing flood events (see Supporting Information Table 2). Moreover, 947 of the 1,678 counties in the sample (56.44%) have at least one flood event which correlates with a detected change point in the take-up rate signal (see Supporting information Figure 1 for a comparison of model performance). Although the CPD methods do not allow us to claim causality these results, along with multiple other studies analyzing similar dynamics (Gallagher, 2014; Choi et al., 2024), our analysis shows that such methods are suitable to analyze flood insurance fluctuations and identify substantial changes.

For the counties and flood events with an increase in the insurance market after a flood event, we proceed to analyze if and how long it takes for their memory effect to wear out. Figure 2 presents the results of the analysis, where we find 756 counties where a flood event is followed by an increase in the take-up rate and then a drop back to base levels, and the distance of this drop, the time to forget exhibits a mean of 28.03 months (i.e., 2.3 years) and a median of 18 months (1.5 years).

The top map in Figure 2 displays three levels of flood impact on insurance spikes (low to high) across US counties, showing high salience values along major river systems like the Mississippi river, and low salience in coastal regions. The accompanying histograms show the distribution of the flood memory effect (ranging to 120 months) and flood salience (measured as percentage change). A detailed highlight of the selected case study of Suwannee County, Florida, showcases a typical pattern of temporal dynamics of insurance take-up rates, showing distinct change points following flood events and a characteristic "time to forget" pattern where insurance participation gradually declines. Social memory might be way shorter than proposed by previous studies, with the average time-to-forget standing at 28.03 months. This time is just a little over two years, which is in contrast to the three to nine years previously proposed. Extreme weather events can have an average impact of a +7.92% in insurance take-up rate, consistent with previous

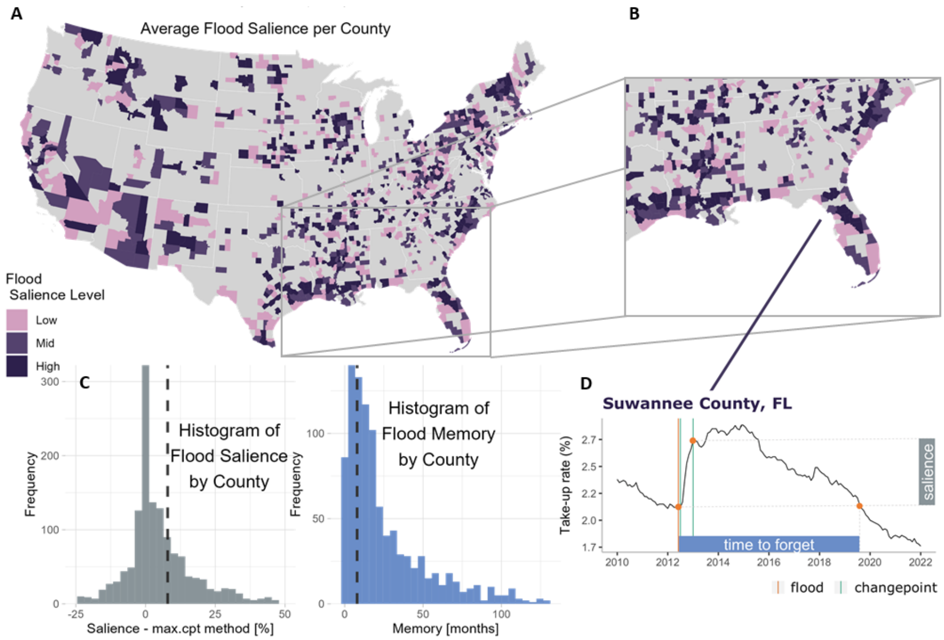

Figure 2: Analysis of flood salience and memory across a sample of 1678 U.S. counties, selected based on threshold criteria of damage and flood insurance take-up rates. The map in panel A, top-left, shows the geographical distribution of three levels of salience effects (low, mid, high), indicated with different colors. The zoom window in panel B shows the geographical distribution of memory and salience in the southeastern coast. A corresponding frequency distribution of the salience and memory variables by county are represented in the two histograms below the map (panel C). As an example, Suwanee County in Florida is selected to illustrate the temporal evolution of flood insurance take-up rates in a high salience county (panel D). The identified change points are marked as orange dots in the timeline and the flood occurrences are indicated as green vertical lines.

researchers' conclusions. However, considering the distribution median, this result decreases to +2.30%, which is indicated in Table 1 as "salience median".

Given the skewed distribution of memory and salience shown in the histograms in Figure 2, the mean values of salience and memory alone do not capture the heterogeneity of flood insurance update and memory across different counties. In the next sections, we disentangle the distribution of memory and salience with additional data, which will allow us to identify and characterize different behavioral groups.

3.3. PCA revealing heterogeneous flood insurance take-up characteristics

We created three combined variables using PCA, which linearly combines and weights individual drivers of flood insurance adoption (shown on y-axis in Figure 3 top panel). The first three PCs accounted for 61.6% of the variance in the data, thus we limit our analysis to such components (see Supporting Figure 2). An analysis of the variables that contribute the most to each of the three main PCs reveal distinct determinants in each PC. PC1 captures *risk perception*, heavily weighted by initial take-up rates, CRS class rating, risk score, and percent water coverage. PC2 primarily captures *social vulnerability* aspects, with strong PC loadings on social vulnerability scores, negative

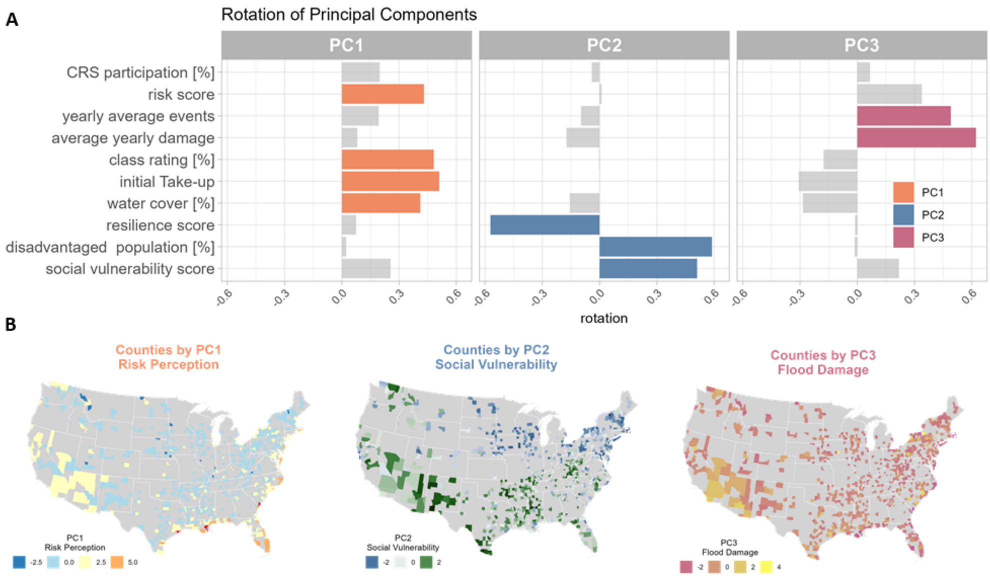

Figure 3: The top panel a shows the rotation of each variable for the PC's 1-3. The first component (PC1) primarily captured social vulnerability metrics, with strong loadings on social vulnerability scores and disadvantaged population percentages. PC2 represented risk perception through variables like initial take-up rates and percent water coverage, while PC3 characterized flood damage patterns through average yearly damage and event frequency measurements. The three maps in the bottom (panel B) show the geographical distribution of the PC values for counties, where change points in insurance uptake match with simulated change points. High risk perception (PC1) is clustered around the southeastern coast, social vulnerability (PC2) is heterogeneously distributed and the highes flood damage values are centered in the southwest.

resilience score, and disadvantaged population percentages from the census data (U.S. Census Bureau, 2022). PC3 clearly identifies *flood damage* patterns, dominated by average yearly damage and event frequency. These results show that risk perception and flood damage appear as independent PCs, both simultaneously contributing to explaining the variance in the dataset individually. This suggests that risk perception and actual risk do not necessarily overlay in the US, meaning that individual choices may be driven by perceptions that can different from observed actual flood risk. The corresponding maps in the bottom part of Figure 3 report a geographical distribution of the three selected PCs, showing which counties are primarily characterized by each of them. The resulting geographical distribution reveals distinct spatial patterns - social vulnerability concentrating in the southern counties, risk perception peaking along coastlines, and flood damage presenting a more dispersed pattern across the eastern United States. These distinct spatial patterns raise the question of how these components interact to form distinct county-level response types, i.e., how they relate to different patterns of flood insurance salience and memory.

3.4. Disentangling the memory and salience distribution with cluster analysis

The clustering analysis presented here synthesizes the previous findings through grouping counties based on our response variables (*salience* and *memory*), our independent variables grouped into principal components (*PC1, PC2,*

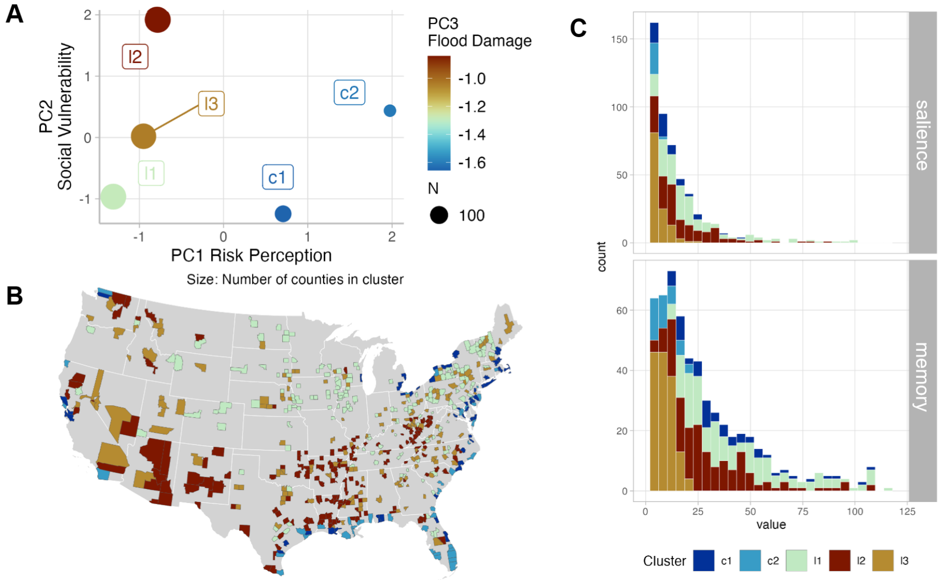

Figure 4: The scatter plot shows clustering results. The clustering shows groupings of counties with similar characteristics, particularly distinguishing between coastal counties with high risk perception (*c1, c2*) and in-land counties with lower flood risk perception (*l1, l2, l3*), counties with both low social vulnerability and low risk perception which respond with a high salience and long memory to flood events (clusters *l1* and *c1*), and highly vulnerable, high risk perception coastal counties with low salience and short flood memory (cluster *c2*). Clusters are shown on the map to verify geographical patterns. In the histograms in the right panel the clustering is mapped onto the salience and memory effects.

and *PC3*), and the categorical binary variable *coastal*, which indicates whether a county is on the shoreline. The clustering results are then mapped onto the histograms of memory and salience to evaluate the connection between the identified clusters and the distributions of salience and memory. Our cluster analysis resulted in the identification of k=5 different clusters of counties, which exhibit different salience and memory effects, each characterized by different combinations of PCs and coastal/inland communities. The top scatter plot in Figure 4 shows these five resulting distinct clusters, each further detailed with descriptive statistics in Table 2. Overall, trade-offs between social vulnerabity characteristics, flood damage, and risk perception emerge from the cluster distribution in the scatter plot. Communities with high social vulnerability do not comparatively exhibit high risk perception and communities with high risk perception are not necessarily affected by flood damage. Clusters *c1* and *c2* include coastal counties, characterized by a higher risk perception than their in-land counterparts in clusters *l1, l2*, and *l3*. The highest risk perception appears in *c2*, with flood-prone, comparatively vulnerable counties mostly found in Florida, the Gulf of Mexico, and the southern Atlantic coast. The counties in this cluster present a low salience effect of 2.8%, and a short memory effect of 8.6 months on average, which can be due to these markets having a significantly higher flood insurance penetration of 14.3% initial take-up rate, compared to a 1.08% on average for in-land counties. Meanwhile, *c1* includes low vulnerability counties mostly located in the north-Atlantic coast and the Great Lakes. This group shows a lower risk perception than *c2*, which corresponds to a lower participation in the NFIP, but a relatively high salience effect after flood events with a

Table 2
Descriptive statistics from cluster analysis and resulting insurance take-up rate salience and memory effects organized by cluster.

Cluster	c1	c2	l1	l2	l3
Counties in cluster (n)	56	44	160	165	154
Coastal counties in cluster (n)	56	44	0	0	0
Avg. initial take-up rate (%)	5.26	14.31	0.81	1.39	1.04
PC1: Risk Perception (index)	0.79	2.18	-1.21	-0.82	-0.92
PC2: Social Vulnerability (index)	-1.17	0.25	-0.95	1.99	0.13
PC3: Flood Damage (index)	-1.71	-1.79	-1.12	-0.89	-0.99
Salience mean (%)	10.52	2.83	21.50	18.99	5.11
Salience median (%)	7.97	2.63	16.00	12.70	4.27
Salience std. dev. (%)	8.63	2.22	20.27	21.61	4.01
Memory mean (months)	38.15	8.57	45.58	32.21	8.92
Memory median (months)	31	8	37	25	9
Memory std. dev. (months)	23.04	5.25	27.33	22.12	4.29

median impact of 10.5% and a memory effect of 38.2 months on average, indicating a reactive, yet long-lasting flood insurance take-up behavior (once a flood insurance is purchased, it is retained for longer time than the average for all data points combined). A noteworthy point is that both of the coastal clusters have, on average, a slightly lower yearly infrastructural damage measure in comparison to in-land counties, which could relate to flood resilience built over the years of NFIP intervention. When it comes to in-land counties, the *l1* cluster shows similar characteristics to *c1* in terms of reactive behavior: communities of low social vulnerability with minimal NFIP participation (low risk perception), but a strong salience effect when hit with a flood event. Mostly located in the area of the Midwest and Great Lakes towards the Mid-Atlantic, the *l1* cluster shows the highest average after-flood effects in our sample, with an average impact of 21.5% and an average memory of 45.6 months. Cluster *l2* stands at the opposite end of the social spectrum with highly vulnerable counties mostly located in the deltas of main rivers like the Mississippi, Colorado, and Columbia. And yet, this cluster's behavior resembles *c1* in showing a high 19.0% salience effect and an average memory of 32.2 months. Lastly, cluster *l3* is characterized by a moderate social vulnerability, low risk perception (in between *l1* and *l2*) and does not show a defined geographical pattern. Despite its low level of risk perception (and therefore, low NFIP participation) and moderate flood damage, the counties in *l3* exhibit a reactive response with a 5.1% average impact measure and a short average memory of only 8.9 months.

4. Discussion & Conclusions

In this study, we analyze over 50 million records of household flood insurance purchase across the Contiguous U.S. from 2010 to 2021 and contribute a data-driven framework based on CPD, dimensionality reduction via PCA, and clustering to identify heterogeneous flood insurance purchase patterns for US counties, along with their salience and memory effects. Our findings demonstrate that household responses to flood events through insurance adoption are characterized by both heterogeneous magnitude (salience) and duration (memory). This temporal pattern suggests a significant challenge in maintaining equality in long-term flood resilience through voluntary insurance programs. The magnitude of flood events' impact on insurance adoption is an average increase of 7.92% in take-up rates. However, we found that this result decreases to 2.30% when looking at the distribution median. With these results we confirm the findings of past research efforts in that, for most communities in the U.S., flood insurance is purchased reactively after floods.

Yet, we find that social memory in the context of flood events, i.e., the time it takes for the temporary increase in flood risk salience after an extreme weather event to fade, may not be as enduring as previously believed. With the proposed methodological setup we find an average "time-to-forget" period of 28.03 months, with a median of only 18 months. Our results indicate shorter memory than the three to nine years suggested by earlier research (Kousky, 2017; Atreya et al., 2015; Gallagher, 2014), with a heavy-tailed distribution that accounts for a highly heterogeneous sample.

We addressed the gap in identifying drivers of the heterogeneous distribution of flood insurance memory and salience. We show that households in counties with greater risk perception may proactively invest on insurance (clusters *c1* and *c2*), which aligns with the findings of (Landry and Turner, 2020), namely that the risk perception of future damage is a reliable determinant of insurance take-up. We also find that counties with low social vulnerability (clusters *l1* and *c1*) seem to have a lower risk perception and to be less prone to invest a priori in flood insurance, compared to similar communities in other clusters. The counties in these low social vulnerability clusters, however, exhibit high salience and long memory effects, implying an effective process for self-actualization of the risk perception after a flood event. Additionally, we find a group of 154 in-land counties (cluster *l3*) that, despite showing a low level of risk perception, low NFIP participation, and moderate flood damage, exhibit a reactive response to floods with a 5.1% average impact measure and a short average memory of only 8.9 months. In these areas, insurance is usually taken up after the event and drops quickly with a return to the low pre-flood levels, leaving these areas vulnerable to experience uninsured losses. We find no evidence that this reactive behavior is due to affordability issues, since clusters with higher social vulnerability (cluster *l2*) show higher salience and longer memory effects. This group then shows a behavior consistent with previous literature, where the flood damage risk gets discounted based on the relative infrequency of flood events (Dumm et al., 2020).

Overall, the comparison of our clustering results shows large heterogeneity in the flood insurance purchase and trade-offs in the tenure behaviors of different counties and communities, suggesting that a one-size-fits-all approach to flood insurance policy may be suboptimal.

Several methodological constraints remain, posing limitations to interpreting our results. First, the county-level aggregation of flood experience may mask significant within-county variations in direct flood exposure, a limitation also described by Choi et al. (2024), who note that this aggregation could potentially underestimate the insurance response among directly affected households. Additionally, our focus on non-Special Flood Hazard Areas zones may be influenced by pre-existing communication biases about flood risk.

Second, in this study we primarily focus on the effects of direct flood experience while disregarding the potential influence of proximity effects, such as floods occurring in nearby counties or the influence of social networks. Previous research has analyzed social and geographical proximity effects with positive results (Gallagher, 2014; Hu, 2020, 2022). Disregarding these effects might negatively influence our evaluation of change point detection methods, particularly the precision and recall metrics, as some real effects might have been incorrectly considered “false positives.”

Our large-scale data-driven analysis ultimately provides evidence supporting the formulation of strategies to counteract the flood insurance protection gap in the US. For instance, extending policy terms and introducing flexible payment structures could effectively address the identified challenges of short-lived insurance adoption and long-term declining participation rates. Additionally, we contribute to the quantitative basis for equitable accessibility by showing which factors should be considered when adapting premiums and which areas would benefit the most from those policies. Beside addressing the individual limitations discussed above, future research should focus on evaluating the effectiveness of risk communication strategies in fostering and maintaining long-term sustained insurance coverage, encouraging insurance update and moving beyond currently dominating reactive behaviors and short social memory.

Software and Data Availability Statement

The data used in this study is openly available from the cited sources. The R Code used to perform the analysis is available upon request.

CRediT authorship contribution statement

Eric Contreras: Conceptualization, Methodology, Software, Validation, Formal Analysis, Resources, Writing - Review & Editing. **Nadja Veigel:** Conceptualization, Methodology, Software, Validation, Formal Analysis, Resources, Writing - Original Draft, Writing - Review & Editing. **Jacek Suda:** Writing - Review & Editing, Supervision. **Heidi Kreibich:** Writing - Review & Editing, Supervision. **Andrea Cominola:** Conceptualization, Methodology, Writing - Review & Editing, Supervision. **Andrea Cominola:** Conceptualization, Methodology, Software, Validation, Formal Analysis, Data Curation, Writing - Original Draft, Writing - Review & Editing, Visualization.

References

Aminikhanghahi, S., Cook, D.J., 2017. A survey of methods for time series change point detection. Knowledge and Information Systems 51, 339–367. doi:10.1007/s10115-016-0987-z.

Atreya, A., Ferreira, S., Kriesel, W., 2013. Forgetting the Flood? An Analysis of the Flood Risk Discount over Time. Land Economics 89, 577–596. doi:10.3368/le.89.4.577.

Atreya, A., Ferreira, S., Michel-Kerjan, E., 2015. What drives households to buy flood insurance? New evidence from Georgia. Ecological Economics 117, 153–161. doi:10.1016/j.ecolecon.2015.06.024.

Bin, O., Kruse, J.B., Landry, C.E., 2008. Flood Hazards, Insurance Rates, and Amenities: Evidence from the Coastal Housing Market. The Journal of Risk and Insurance 75, 63–82. doi:10.1111/j.1539-6975.2007.00248.x, arXiv:25145263.

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, 361–376. doi:10.1016/j.jeem.2012.12.002.

Browne, M.J., Hoyt, R.E., 2000. The Demand for Flood Insurance: Empirical Evidence. Journal of Risk and Uncertainty 20, 291–306. doi:10.1023/A:1007823631497.

van den Burg, G.J.J., Williams, C.K.I., 2022. An evaluation of change point detection algorithms. URL: https://arxiv.org/abs/2003.06222, arXiv:2003.06222.

Cannon, C., Gotham, KF., Lauve-Moon, K., Powers, B., 2020. The climate change double whammy: Flood damage and the determinants of flood insurance coverage, the case of post-Katrina New Orleans. Climate Risk Management 27. doi:10.1016/j.crm.2019.100210.

CBO, 2024. Flood Insurance in Communities at Risk of Flooding. Reports 60042. Congressional Budget Office. URL: https://www.cbo.gov/publication/60510.

Choi, J., Diffenbaugh, N.S., Burke, M., 2024. The effect of flood exposure on insurance adoption among us households. Earth's Future 12, e2023EF004110.

Dumm, R.E., Eckles, D.L., Nyce, C., Volkman-Wise, J., 2020. The representative heuristic and catastrophe-related risk behaviors. Journal of Risk and Uncertainty 60, 157–185.

FEMA, 2021. Information about the community rating system. https://www.fema.gov/case-study/information-about-community-rating-system. Retrieved 21.06.2022.

FEMA, 2023a. National risk index. https://hazards.fema.gov/nri/data-resources. Retrieved 17.11.2023.

FEMA, 2023b. Natural Hazards National Risk Index.

FEMA, 2023c. Openfema dataset: Fima nfip redacted policies. https://www.fema.gov/about/openfema/data-sets/nfip. Retrieved 17.11.2021.

Gallagher, J., 2014. Learning about an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States. American Economic Journal: Applied Economics 6, 206–233. doi:10.1257/app.6.3.206.

Harchaoui, Z., Moulines, E., Bach, F., 2008. Kernel Change-point Analysis. NIPS 21, 609–616. doi:https://dl.acm.org/doi/10.5555/2981780.2981856.

Haynes, Kaylea, Fearnhead, Paul, Eckley, Idris A., 2016. A computationally efficient nonparametric approach for changepoint detection. Cornell University - arXiv doi:10.48550/arxiv.1602.01254.

Hu, Z., 2020. Salience and Households' Flood Insurance Decisions. SSRN Electronic Journal doi:10.2139/SSRN.3759016.

Hu, Z., 2022. Social interactions and households' flood insurance decisions. Journal of Financial Economics 144, 414–432.

Huang, Z., 1998. Extensions to the k-Means Algorithm for Clustering Large Data Sets with Categorical Values. 2, 283–304. URL: http://link.springer.com/10.1023/A:1009769707641, doi:10.1023/A:1009769707641.

James, N.A., James, N., Kejariwal, A., Matteson, D.S., 2016. Leveraging cloud data to mitigate user experience from 'breaking bad'. 2016 IEEE International Conference on Big Data (Big Data) , 3499–3508doi:10.1109/bigdata.2016.7841013.

Jolliffe, I.T., 2002. Choosing a subset of principal components or variables. Principal component analysis , 111–149.

Killick, R., Fearnhead, P., Eckley, I.A., 2012. Optimal Detection of Changepoints With a Linear Computational Cost. Journal of the American Statistical Association 107, 1590–1598. doi:10.1080/01621459.2012.737745.

Kousky, C., 2010. Learning from Extreme Events: Risk Perceptions after the Flood. Land Economics 86.

Kousky, C., 2017. Disasters as Learning Experiences or Disasters as Policy Opportunities? Examining Flood Insurance Purchases after Hurricanes. Risk Analysis 37, 517–530. doi:10.1111/risa.12646.

Kriesel, W., Landry, C., 2004. Participation in the National Flood Insurance Program: An Empirical Analysis for Coastal Properties. The Journal of Risk and Insurance Vol. 71, pp. 405–420. arXiv:3520070.

Kunreuther, H., Meyer, R., Michel-Kerjan, E., 2009. Chapter 23. Overcoming Decision Biases to Reduce Losses from Natural Catastrophes, in: Shafir, E. (Ed.), The Behavioral Foundations of Public Policy. Princeton University Press, pp. 398–414. doi:10.1515/9781400845347-027.

Landry, C., Turner, D., 2020. Risk perceptions and flood insurance: Insights from homeowners on the georgia coast. Sustainability 12. URL: https://www.mdpi.com/2071-1050/12/24/10372, doi:10.3390/su122410372.

Landry, C.E., Jahan-Parvar, M.R., 2011. Flood Insurance Coverage in the Coastal Zone. Journal of Risk and Insurance 78, 361–388. doi:10.1111/j.1539-6975.2010.01380.x.

Lindell, M.K., Hwang, S.N., 2008. Households' Perceived Personal Risk and Responses in a Multihazard Environment. Risk Analysis 28, 539–556. doi:10.1111/j.1539-6924.2008.01032.x.

Lloyd, S., 1982. Least squares quantization in pcm. IEEE transactions on information theory 28, 129–137.

Lucas, C.H., Booth, K.I., Garcia, C., 2021. Insuring homes against extreme weather events: A systematic review of the research. Climatic Change 165, 61. doi:10.1007/s10584-021-03093-1.

Matteson, D.S., James, N.A., 2013. A Nonparametric Approach for Multiple Change Point Analysis of Multivariate Data. arXiv: Methodology doi:10.1080/01621459.2013.849605.

Michel-Kerjan, E., de Forges, S.L., Kunreuther, H., 2012. Policy Tenure Under the U.S. National Flood Insurance Program (NFIP). Risk Analysis 32, 644–658. doi:10.1111/j.1539-6924.2011.01671.x.

Michel-Kerjan, E., Kousky, C., 2010. Come Rain or Shine: Evidence on Flood Insurance Purchases in Florida. Journal of Risk and Insurance 77, 369–397. doi:10.1111/j.1539-6975.2009.01349.x.

NOAA, 2023. National centers for environmental information storm events database. https://www.ncdc.noaa.gov/stormevents/. Retrieved 17.11.2022.

Scott, A.J., Knott, M., 1974. A Cluster Analysis Method for Grouping Means in the Analysis of Variance. Biometrics 30, 507. doi:10.2307/2529204, arXiv:2529204.

Truong, C., Oudre, L., Vayatis, N., 2020. Selective review of offline change point detection methods. Signal Processing 167, 107299. doi:10.1016/j.sigpro.2019.107299.

U.S. Census Bureau, 2018. 2018 American Community Survey 5-year Public Use Microdata Samples. URL: api.census.gov/data/2018/acs/.

U.S. Census Bureau, P.D., 2022. Annual Estimates of Housing Units for Counties in the United States: April 1, 2020 to July 1, 2021 (CO-EST2021-HU).

van den Burg, G.J.J., Williams, C.K.I., 2020. An Evaluation of Change Point Detection Algorithms. arXiv: Machine Learning doi:10.48550/arXiv.2003.06222.

Veigel, N., Kreibich, H., Cominola, A., 2023. Interpretable Machine Learning Reveals Potential to Overcome Reactive Flood Adaptation in the Continental US. Earth's Future 11, e2023EF003571. doi:10.1029/2023EF003571.

Wang, D., Davidson, R.A., Trainor, J.E., Nozick, L.K., Kruse, J., 2017. Homeowner purchase of insurance for hurricane-induced wind and flood damage. Natural Hazards 88, 221–245. doi:10.1007/s11069-017-2863-x.

Wing, O.E., Pinter, N., Bates, P.D., Kousky, C., 2020. New insights into US flood vulnerability revealed from flood insurance big data. Nature Communications 11, 1–10. doi:10.1038/s41467-020-15264-2.

Wing, O.E.J., Bates, P.D., Smith, A.M., Sampson, C.C., Johnson, K.A., Fargione, J., Morefield, P., 2018. Estimates of present and future flood risk in the conterminous united states. Environmental Research Letters 13, 034023. URL: https://doi.org/10.1088/1748-9326/aaac65, doi:10.1088/1748-9326/aaac65.

Wing, O.E.J.., Lehman, W., Bates, P.D.., Sampson, C.C.., Quinn, N., Smith, A.M.., Neal, J.C.., Porter, J.R.., Kousky, C., 2022. Inequitable patterns of US flood risk in the Anthropocene. Nature Climate Change 12, 156–162. doi:10.1038/s41558-021-01265-6.

Zahran, S., Weiler, S., Brody, S.D., Lindell, M.K., Highfield, W.E., 2009. Modeling national flood insurance policy holding at the county scale in Florida, 1999–2005. Ecological Economics 68, 2627–2636. doi:10.1016/j.ecolecon.2009.04.021.