Learning, Catastrophic Risk, and Ambiguity in the Climate Change Era

Frances C. Moore

April 2024

Abstract

Property insurance markets in multiple US states are experiencing volatility in the form of increasing premiums, insurer bankruptcies, and complete withdrawal from some areas. This volatility appears linked to growing costs of natural disasters, with some of the most serious effects in states heavily exposed to hurricanes and wildfires. This paper illustrates how uncertainty introduced by a changing climate can produce abrupt changes in insurance market conditions as insurers (and reinsurers) integrate the possibility of changing climate conditions into their pricing and risk management.

The mechanism operates via changes to the updating process by which agents infer current climate conditions and, by extension, the distribution of weather risk. Key methodologies used for managing weather risks in both engineering and finance applications have, implicitly or explicitly, relied on the assumption that the climate (the probability distribution over weather) is not changing over time, and therefore the historic weather record is representative of current risks. Anthropogenic climate change upends this assumption by introducing the possibility, or even the likelihood that the climate distribution today is different from past experience. This effectively reduces the information available to actors and increases ambiguity in the estimated climate distribution, with associated costs for weather risk management and risk-averse decision-makers. These costs arise purely from the knowledge that the climate could be changing, may arise abruptly, are additional to any direct costs or benefits from actual climate change, and are, to date, entirely unquantified.

Using a case study of extreme rainfall-related flood damages in New York City, this paper illustrates how these ambiguity-related costs arise. Greater uncertainty over the current climate distribution interacts with a steeply non-linear damage function to greatly increase the mean and variance of the loss distribution. I show how this uncertainty can ripple through insurance markets in the form of higher and more volatile premiums and higher reinsurance costs, with limited potential for diversification within the insurance sector, impacts consistent with observed changes in the U.S. property insurance market in recent years.

1 Introduction

Recent decades have seen rapid increases in the frequency and severity of extreme weather events with significant economic losses. In the U.S., the number of events with losses of more than $1 billion (in inflation-adjusted terms) increased from an average of 3.3 events per year in the 1980s to 20.4 events per year in the last 5 years [25]. These growing losses are likely the result of interactions between anthropogenic climate change altering the frequency and intensity of extreme weather events and growth in population density and capital stocks in high-risk areas (e.g. [27]).

While the relative importance of risky development patterns versus anthropogenic climate change in driving extreme event losses may be debated, it is clear that growing losses are posing challenges to private insurance markets. Major insurers largely exited Florida and Louisiana following large hurricane-related losses years since 2005. Those markets are now dominated by small firms with highly-concentrated risk, heavily reliant on the reinsurnace market. As of 2018, over 50% of value underwritten in Florida is from firms without a credit rating from the major ratings agencies and nine Florida insurers became insolvent between 2021 and 2023 [30, 13]. Unprecedented wildfires have driven record losses in California and led major insurers to limit underwriting in the state, leading to massive growth in the state's public "last resort" insurance program [19, 16]. Price volatility or unavailability of property insurance can quickly spillover to the mortgage market because of the requirement from lenders that properties that secure the loan be insured.

Natural hazards are challenging for private insurers to cover because losses are highly concentrated in space and time. Unlike other insurance lines where claims are stable from year-to-year and premiums can be set to closely match, natural hazards losses exhibit substantial interannual variability, even when aggregated across all perils at the global level [34]. Losses from California wildfires in 2017 and 2018 was more than double the industry profit from all property insurance in the state for the last 30 years [19]. The nature of these losses require insurers underwriting these risks to maintain access to large amounts of liquid capital to pay claims in the event of a major disaster. This is expensive as it requires paying fees to reinsurers or premiums to investors in insurance-linked securities (ILS). These costs are passed on to consumers, potentially raising premiums above expected losses, depressing demand.

This paper highlights how climate change can interact with pre-existing catastrophic risks to raise costs of both insurance and reinsurance. Using recent observed instability in the insurance market as a motivation, it highlights an under-appreciated pathway by which climate change impacts society. Simply the knowledge that past experience of weather may no longer be representative of current risks decreases the information available to market actors and increases uncertainty. The cost of this

added uncertainty may be small for some types of risk but could be substantial for natural hazard risks, where expected losses are driven by very rare (and therefore highly uncertain) events. The paper walks through a stylized model of catastrophic risk and Bayesian updating, using a case study illustration based on extreme rainfall-related flood damages in New York City. I show how simply the knowledge the climate might be changing alters the updating problem to add uncertainty over current weather risks. This propagates through the damage distribution to substantially raise expected damages, even though neither the damage function nor historic evidence on extreme events has changed. I trace the implications of these altered damage distributions through insurance markets, addressing 1) expected losses and actuarily fair premiums; 2) premium volatility; 3) reinsurance costs; and 4) the potential for diversification.

2 Background

2.1 The Climate Distribution

"Climate is what you expect, weather is what you get" - Andrew John Herbertson, 1901

Climate, particularly in a period of relatively rapid climate change as we are now in, is best understood as a probability distribution over weather [18]. Because of the nonlinear dynamics that govern the atmosphere, particular weather outcomes are unpredictable beyond a lead-time of somewhere between a couple weeks to about 6 months for seasonal forecasts. An irreducible exists in weather outcomes, meaning any economically-relevant, weather-dependent outcome will have associated risk. A climate can be usefully understood as the probability density over weather outcomes, useful for quantifying the distribution of weather-related risks. A climate may be time- and space-specific and may be defined jointly over multiple relevant weather metrics (e.g. maximum temperature, wind speed, absolute humidity, precipitation etc).

Understanding climate as a probability density makes clear the inherent challenge in attempting to manage (or insure) weather risks. Weather risks are determined by the full climate distribution, but this is inherently unobservable. At any place and time we observe only a single draw from the climate distribution (i.e., the weather). Actors tasked with managing weather risks are therefore faced with a fundamental inference problem: how to estimate the full climate distribution given only a single history of weather observations?[1]

If the climate distribution is known to be unchanging over time (i.e., stationary), then a long enough

[1]The question of how scientific model information can used to help estimate the climate distribution is discussed more fully in Section 5. While potentially an important complement to observational data, there is evidence from both financial markets and the scientific literature that current models do not substantively help constrain near-term, location-specific extreme weather risks.

weather record can constrain the current climate. In the limit, an infinitely long record will perfectly characterize the climate distribution and therefore resolve any weather-related ambiguity over losses (ambiguity is used here to refer to uncertainty over a probability distribution). In reality, however, weather observations are not infinitely long. A standard definition used by the World Meteorological Organization to define a climate distribution (the so-called "climate normal"), is 30 years [9]. Observational weather records date back somewhere between 70 and 100 years in most locations, up to a few hundred years in some places. Paleoclimate records from coral, tree rings, ice cores and sediments, can push records back much further, but for a limited set of weather variables, in limited locations, and with substantial measurement error and uncertainty.

While 30-100 years may be more than enough information to constrain the expected value of thin-tailed weather variables, some weather variables exhibit long-tails, where their expected value depends sensitively on rare events. Even 100 years of observations contain, in expectation, only five 1-in-20 year events and only one 1-in-100 year event. Even under a stationary climate therefore, limits in the observational record could leave substantial ambiguity in the tail of the climate distribution and therefore, for heavy-tailed weather variables, meaningful ambiguity in their expected value.

Anthropogenic climate change complicates this setting further by undermining the stationarity assumption in the interpretation of weather observations. Greenhouse gas emissions have altered the energy-balance of the planet, producing clearly detectable alterations in global and regional temperatures, rainfall patterns and intensity, and river flows, among other variables [28, 1, 37, 23]. The fact that humans are influencing the climate system renders older records *potentially* uninformative of current probabilities, effectively decreasing the information available to estimate the climate distribution. The magnitude of these effects will be most pronounced for extreme events in the tail of the distribution, where the observational record is already limiting. Although absolute probabilities of historically-unusual events may remain small, increasing ambiguity in the climate distribution could produce large *relative* changes in probability.

2.2 Catastrophic Risk

Weather risk in a particular setting depends both on the distribution of a weather variable (or combination of weather variables) and a damage function that maps realizations of weather onto losses. The distribution of losses arises from convolving the distribution of the weather variable (i.e. the climate) with the damage function. Damage functions with thresholds and / or steep non-linearities can amplify the importance of the tail of the weather distribution in determining expected losses: if losses increase non-linearly with the weather variable, then expected *losses* (even more so than expected weather)

will be driven by very rare but extremely damaging events.

Catastrophic risk occurs when the loss distribution is heavy-tailed so that expected losses are heavily driven by very rare events [7]. Any setting where a long-tailed physical driver (for instance, rainfall intensity or earthquake magnitude) interacts with a damage function that is steeply non-linear in the physical driver could produce heavy-tailed catastrophic risks. Non-linear damage functions are more common than not in the literature, with thresholds and non-linear responses documented in a range of settings from agricultural yields, to human mortality [31, 6, 5]. These are associated with excedances of either natural, engineered, or social tolerances (for instance, over-topping river banks, excedance of building design codes, crop physiological limits).

3 Case Study Illustration

The remainder of this paper develops an extended case study based on rainfall-induced flooding in New York City (NYC) to illustrate how shifting learning models to account for climate change could affect insurance markets.

3.1 Weather Data and Damage Function

The motivation used here to develop the stylized illustration used in this paper is urban flooding. Rainfall intensity, a critical driver of flood frequency and magnitude, is known to potentially have a heavy-tailed distribution. Peak rainfall intensities that drive flood events are typically modeled using Generalized Extreme Value or Peaks Over Threshold models, which can produce heavy-tailed distributions such as the Weibull or Frechet [36]. Moreover, aggregate damages from intense rainfall are likely to be characteristic of catastrophic risk. Rainfall events of moderate intensity can be handled by existing drainage and flood-defense infrastructure but larger intensity events can increasingly overwhelm these systems to produce a steeply increasing damage function as more properties are affected and sustain heavier damage due to deeper flood depth [35].

Underlying climate data comes from daily rainfall data from the Central Park, NY rain gauge, which goes back to 1869. Figure 1a shows annual maximum rainfall data for the most recent 30 year climatology, from 1994 to 2023. The record shows substantial variability. For instance, while the first 13 years saw maximum rainfall of just over 5 inches in a day, 2007 saw 7.6 inches of rain in a day, exceeding the previous maximum by over 50%. Figure 1b shows the best-fit Weibull distribution fit to the 30 year record in Figure 1a.

The shape of the damage function is based on annual data on all flood insurance claims paid through

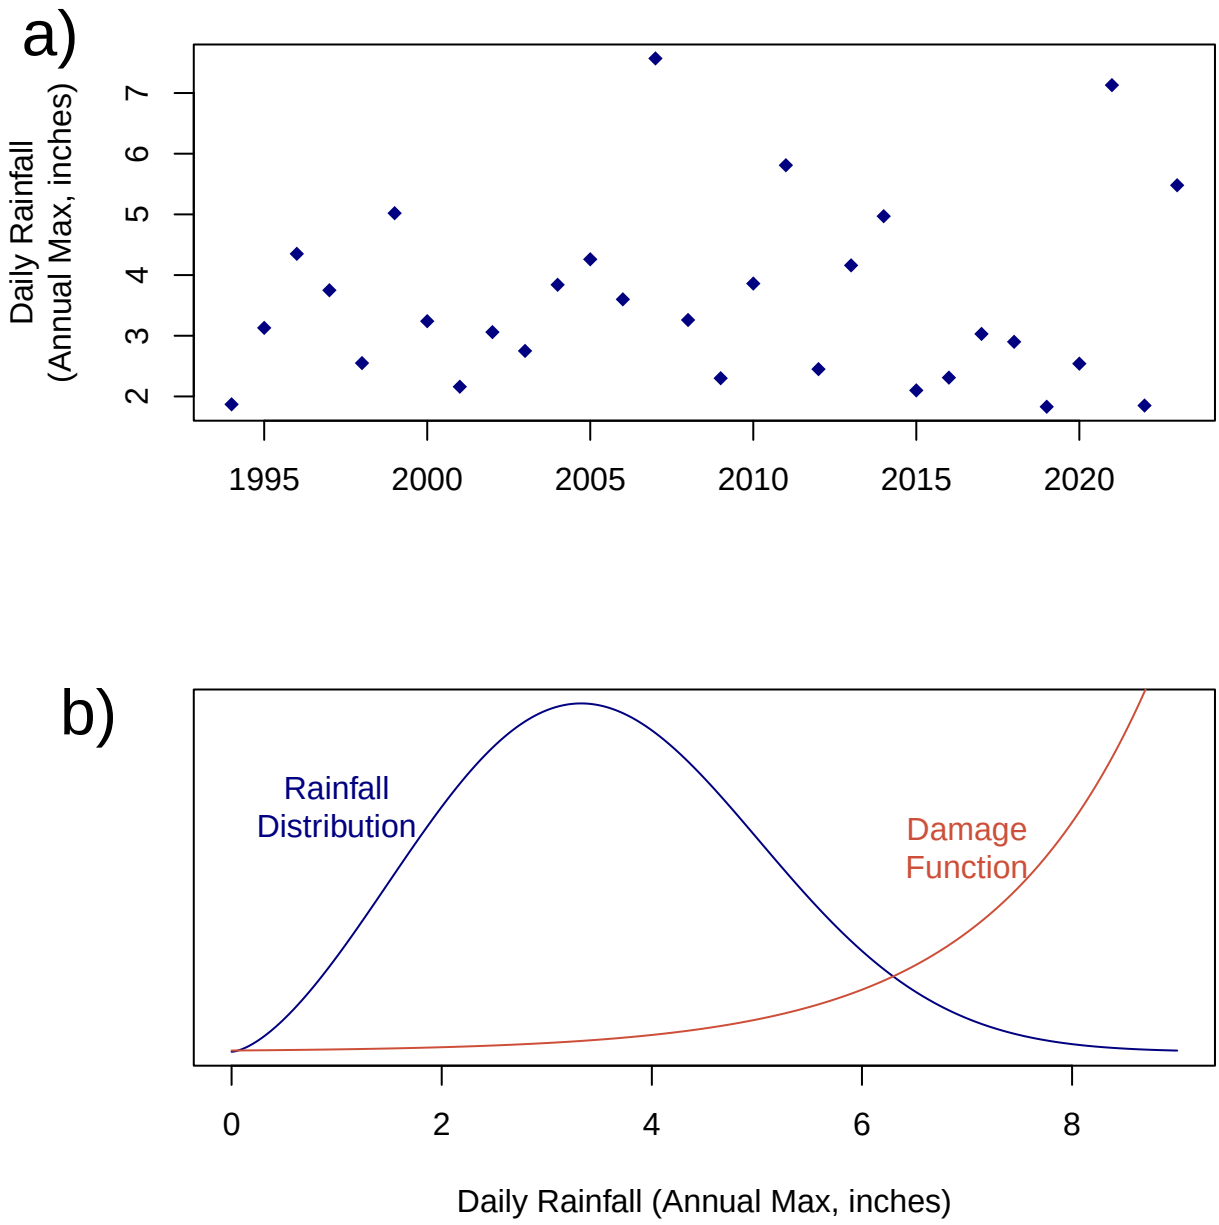


Figure 1: **Rainfall distribution and damage function used for the illustration.** a) Annual maximum daily rainfall from the Central Park, NY rain gauge for 30 years from 1994 to 2023. b) Weibull distribution fitted to the rainfall data with the fitted damage function based on flood insurance claims in New York under the National Flood Insurance Program, controlling for total coverage levels and annual maximum tide heights (additional details in Appendix A.1).

the National Flood Insurance Program (NFIP), 2009-2023 in New York City (NYC). The Federally-run NFIP accounts for more than 90% of flood insurance coverage in the United States. Flood insurance take-up is very low (approximately 3% nationwide) so these damages do not reflect total flood damages, but they do provide an unusually comprehensive view of insurer losses - the most relevant variable for this illustration - and how they vary with rainfall intensity. The shape of the damage function is estimated controlling for annual maximum tide height and total policy coverage, and is robust to the exclusion of 2012 (the year of Hurricane Sandy). Additional details on the damage function estimation and regression model results are given in Appendix A.1.[2]

Figure 1b shows the estimated damage function superimposed on the best-fit Weibull distribution for the 30 year record in Figure 1a. The exponential shape of the damage function is such that most years incur little or no flood-related damages with the bulk of damages concentrated in very intense but unusual events. As one illustration, the 25% of years with lowest maximum rainfall account for just 5% of losses while the top 25% of years account for a disproportionate 65% of losses. Seven percent of damages arise from events not observed in the 30 year climatology and 3% come from events not observed in the full 155 year record at the Central Park station, those with less than a 0.2% annual chance of occurring (under the stationarity assumption).

3.2 Ambiguity and Learning Over the Climate Distribution

Actors seeking to manage (or insure) flooding-related risks in the present (here taken as 2024) face the challenge of inferring the current probability distribution (i.e. the climate distribution) over peak rainfall intensities, given the available history of observations. The climate distribution cannot be known for certain, but instead must be estimated, producing an ambiguity in the current climate distribution. For the set of simulations shown here, I operationalize this learning as a Bayesian updating process over one of the two parameters of the Weibull distribution. The Weibull distribution is commonly used to fit extreme rainfall statistics and is described by two parameters: the shape parameter (α), which describes behaviour of the tail of the distribution ($\alpha < 1$ produces fat-tailed distributions and $\alpha > 1$ gives thin-tailed distributions), and the scale parameter (θ), which describes how "stretched" the distribution is along the x-axis (for a given shape parameter, larger values of the scale parameter will have more probability mass at higher values).

In the interests of simplicity and to remain conservative in describing learning model impacts, both learning models described here assume that 1) the shape parameter of the distribution remains con-

[2]For the purposes of this paper, I abstract from the institutional fact that flooding is covered almost entirely by the Federal government in the US, and discuss insurance market implications *as though* losses accrued to a private insurer. Flooding is a useful case study, precisely because of the ready availability of insured loss data from the NFIP to support estimation of the damage function. The essential intuition developed using this case study should extend readily to other climate-related natural disasters such as windstorms and wildfires that are still covered by private insurers in the US.

stant, actors know the value, and that it doesn't change[3]; 2) agents know the damage function precisely; and 3) perform optimal Bayesian updating over the scale parameter of the Weibull distribution. These are clearly conservative assumptions in many ways. In particular, the assumption of a fixed shape parameter substantially limits the potential ambiguity introduced by climate change, by fixing the asymptotic behavior of the right tail of the distribution. Adding uncertainty over the shape parameter would introduce the possibility of much heavier tails into the agent's prior, and therefore would likely produce similar effects but of much larger magnitude than those described here.

To highlight the pure ambiguity costs of climate change (i.e the costs arising from being unable to assume a stationary weather distribution), I contrast two sets of results throughout the remainder of the paper, both with agents using the same 30 year record (1994-2023) and the same damage function, just varying whether or not the agent assumes the rainfall distribution is unchanging (the "Assumed Stationarity" model) over the period, or allows for non-stationarity (the "Potential Non-Stationarity" model). In both models, the agent's problem is to infer the probability distribution over extreme rainfall for the current year (i.e. 2024).

1) Assumed stationarity: Agents assume the climate distribution over the 30 year period is stationary and representative of the present. They know the climate distribution over annual maximum rainfall intensities, x, is distributed Weibull with likelihood:

$$L(x|\alpha,\theta) = \frac{\alpha}{\theta}x^{\alpha-1}e^{-\frac{x^\alpha}{\theta}}$$

where shape parameter, α, is known and the scale parameter, θ must be estimated.

The agent holds a prior over θ distributed inverse gamma (the conjugate prior of the Weibull scale parameter, used to limit computational complexity[4]) with density:

$$p(\theta|a,b) = \frac{b^a e^{-\frac{b}{\theta}}}{\Gamma(a)\theta^{a+1}}$$

The parameters of the prior are set so that the prior is broad but partially informative, with $a = 1.5$ to give a diffuse, heavy-tailed prior distribution and b chosen so that the mean of the distribution ($\frac{b}{a-1}$) is equal to the estimated shape parameter from the prior 30-year climatology (i.e. using data from 1962-1993).

Agents use the 30-year climatology in Figure 1a to update their posterior to a new inverse gamma

[3]The known shape parameter is based on the best-fit Weibull distribution to the 30 year record (here taken to be 1994-2023) and has a value of 2.57, producing a right-skewed but thin-tailed distribution.

[4]In Bayesian learning, using prior distributions from the conjugate of the likelihood distribution provides a closed-form solution for the posterior, allowing the posterior distribution to be calculated directly from the data and the parameters of the prior, rather than deriving it computationally

distribution with parameters $a' = a + 1 + n$ and $b' = b + \Sigma_t x_t^\alpha$ where $n = 30$ is the length of the climatology and x_t is the observation from year t.

This posterior defines the agent's beliefs over possible values of the scale parameter of the rainfall distribution. Each draw from the posterior, when combined with the fixed scale parameter ($\alpha = 2.56$) defines a probability distribution over rainfall outcomes, each of which defines a particular distribution over damages given the fixed damage function. The agent's beliefs over *damages* is calculated by:

1. Drawing 10,000 samples θ_i from the posterior distribution
2. For each draw, drawing 10,000 samples from the Weibull rainfall distribution defined by θ_i and the shape parameter α, producing 10,000 * 10,000 = 100 million samples from the posterior rainfall distribution [5]
3. Passing all 100 million samples through the damage function to give the posterior damage distribution

The posterior distribution over *damages* is estimated by taking 10,000 draws from the posterior θ distribution and then, for each draw, propagating 10,000 samples from the Weibull distribution given by that draw and the known shape parameter α. through the damage function shown in Figure 1b. This gives 10 million samples from the posterior damage distribution.

2) Potential Non-Stationarity: Agents know simply that the climate may be changing, but receive no additional information on exactly how for the particular hazard and location of interest. They are forced to drop the stationarity assumption and allow the unknown scale parameter to vary over time (i.e. the parameter becomes time specific, θ_t). The inference problem is now to estimate the 2024 distribution, i.e. θ_{2024}.

In the interests of limiting computational complexity, possible time variation is limited to the set of linear trends over time t:

$$\theta_t = \theta_0 + \beta t$$

Both the initial scale parameter, θ_0 and the rate of change, β are unknown. The prior over θ_0 is distributed identically to the stationary case (i.e. a broad inverse gamma distribution partly informed by the prior 30-year period). The prior over β is normally distributed around zero, allowing for the scale parameter (and, equivalently, the probability of extreme rainfall events) to be constant ($\beta = 0$),

[5] The potential importance of catastrophic events is of primary interest in this paper. Since these are rare by definition, accurate characterization of the tails of the relevant probability distributions is essential. If computational approximation of distributions is too coarse (i.e. dos not contain enough samples) the tails of the distributions will be poorly sampled and risk estimates will be downward bias. That is why I use what may seem to be excessively large sample sizes (though computational requirements are not excessively burdensome - all code for the paper can run in less than an hour in parallel over 12 cores on a modern laptop computer).

increasing ($\beta > 0$), or decreasing ($\beta < 0$) over time. The width of the posterior is set arbitrarily such that the central 95% of the distribution allows for a change of ± 1 by the end of the n (i.e. 30) year period (from a prior mean starting value of 3.1) giving the prior distribution over β:

$$\beta \sim N(0, \frac{0.5}{n})$$

To estimate the current climate, the agent must now use the same 30 year record to estimate the joint posterior probability distribution over both θ_0 and β. Since the agent must now estimate two parameters instead of one from the same record, they have effectively lost information and the posterior distribution *must* be wider than in the stationary case. This can also be seen by noting that the stationarity case assumed in the first learning model is nested as one possibility in this model ($\beta = 0$). Since this new model admits a broader set of possibilities ($\beta \neq 0$) the priors are broader and, given the same set of data for updating, the posterior must also be wider. The question is just how much wider, and what are the potential implications for the loss distribution given interactions with the non-linear damage function.

Since simple conjugacy no longer applies, the posterior is calculated computationally using Bayes Rule. For a given draw of θ_0 and β, posterior probabilities given the set of observations, $\mathbf{x}$, is given by:

$$p(\ \theta_0, \beta|\mathbf{x}) \propto \Pi_t L(x_t|\alpha, \theta_0, \beta) p(\theta_0) p(\beta)$$

Where $p(\theta_0)$ and $p(\beta)$ are prior probabilities and the likelihood of the data point x_t is given by the Weibull distribution with the time varying scale parameter:

$$L(x_t|\alpha, \theta_0, \beta) = \frac{\alpha}{\theta_0 + \beta t} x_t^{\alpha-1} e^{-\frac{x^\alpha}{\theta_0 + \beta t}}$$

The joint posterior distribution over θ_0 and β is sampled using 16 million draws from the joint prior density (4000 draws from the β prior and, for each draw, 4000 independent draws from the θ_0 prior. The posterior distribution over the current climatology (i.e. θ_{2024}) comes from 10,000 samples of the joint posterior:

$$\theta_{2024} = \theta_0 + 30\beta$$

The posterior *damage* distribution in turn is estimated similarly to the stationary case by, for each 10,000 samples of θ_{2024} from the posterior, taking 10,000 samples from the Weibull rainfall distribution implied by that draw and the known shape parameter, α and propagating those through the damage function. This again gives 100 million draws from the posterior damage distribution.

4 Results

The impacts of being forced to relax the stationarity assumption because of the existence of climate change are illustrated throughout by contrasting results for the two updating processes described in Section 3. I first describe effects on the posterior climate distribution, then discuss how this propagates through to the damage distribution, before describing how uncertainty could propagate through to disrupt functioning of insurance markets.

4.1 Posterior Climate Distribution

Figure 2a shows the posterior distribution over the scale parameter for the two updating processes. Simply relaxing the assumption of stationarity to allow a linear trend in the scale parameter substantially widens the posterior density and shifts it towards higher values. Larger values of the scale parameter give a more "stretched" distribution, with a longer right tail and more probability mass at historically extreme values. While the prior over the trend parameter puts equal probability on increases or decreases in the scale parameter over time, integrating evidence from the historical record decisively shifts the posterior in favor of increases over time (posterior probability of $\beta > 0$ is 83%).

Note that this effect is almost entirely driven by a shift in the learning model, rather than clear evidence of changing extreme rainfall conditions in the historical data used for updating[6]. This is illustrated by the dotted distributions in Figure 2a which show posterior distributions under identical learning procedures, but updated using a 30 year simulated record that is stationary by construction (drawn from the best-fit Weibull distribution based on the 30 year climatology). Although these are both shifted to the left relative to the posteriors based on real-world data (i.e. place slightly less probability on very extreme rainfall events), the key elements of the simulation remain: posterior probabilities under potential non-stationarity are both broader and substantially shifted to the right compared to the case where stationarity could be assumed.

This asymmetric effect arises from exactly how the available evidence - namely, 30 years of rainfall maxima - act to constrain the set of possible models, given the zero lower-bound and right-ward skew of the underlying rainfall distribution. Because significant sampling variation of tail events in a 30 year record is to be expected, agents in the potential non-stationary case are unable to distinguish between a large upward trend in the scale parameter combined with relatively "normal" draws from

[6]The scientific basis for expecting more extreme rainfall events in a hotter climate is well established. Hotter air can hold more moisture, producing both longer and more intense dry spells and more extreme precipitation events. Evidence for shifts in these patterns over long timescales at the global scale has been demonstrated [23, 37]. Therefore, there is good reason to suspect anthropogenic climate change has had an effect on the Central Park station record used here and increased the intensity of major events. The discussion here is not meant to suggest otherwise, but to point out that such an effect is not required to produce shifts in the posterior probability densities I demonstrate. Instead these can arise purely from the interaction of a broader prior distribution with a skewed likelihood distribution under sampling variability.

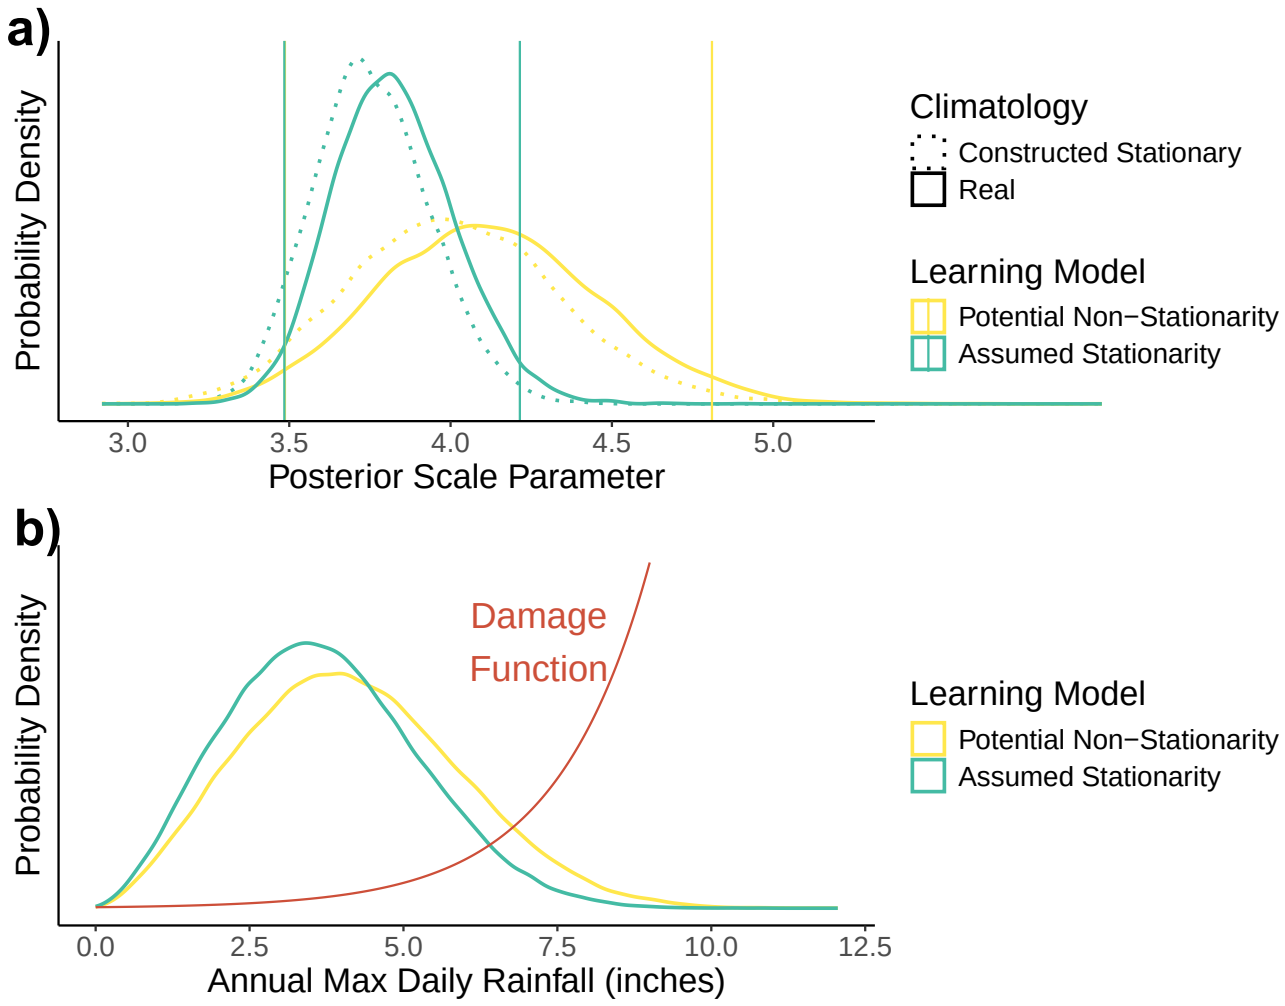


Figure 2: **Posterior Distribution Under Diferent Learning Models.** a) Posterior distribution over the scale parameter of the climate distribution in 2024 after updating using information from the 30-year maximum rainfall climatology shown in Figure 1a under learning models that do and do not assume stationarity. Vertical lines mark the central 95% of the distributions.Dotted distributions posterior densities under two learning models, but based on an artificial time-series of observations that is stationary by construction (i.e. simulated observations are drawn from the Weibull distribution shown in Figure 1b). b) Distribution over maximum daily rainfall given implied by the 95th percentile of the posterior damage function under the two learning models, with the overlaid damage function.

Expected Damages	Variance	Damage Distribution Percentiles							
		25	50	75	90	95	97.5	99	99.5
1.32	2.85	1.11	1.17	1.25	1.33	1.39	1.46	1.54	1.61

Table 1: **Measures of the Damage Distribution** Expected value, variance, and quantiles of the damage distribution, shown as the ratio under the two learning models for each statistic (i.e value under learning allowing for non-stationarity over value under assumption of stationarity).

the underlying climate distribution and little to no trend in the underlying distribution combined with unusually high samples from the distribution. In contrast, just one or two relatively high draws in the dataset can effectively eliminate the possibility of a large downward trend, since the sampling probabilities would be so low. Posteriors in the assumed stationarity case are both narrower and lower because agents have ruled out the possibility of a trend and are therefore better able to use absence of evidence as evidence of absence: if very large rainfall events do not appear in the record, it is probably because of low underlying probabilities and not because sampling variability over 30 years produced a series of "lucky" draws.

Figure 2b maps differences in the posterior θ distribution into difference in rainfall probabilities. The figure shows the rainfall distribution associated with the 95th percentile of both posterior distributions. The larger scale parameter under potential non-stationarity ($\theta_{0.95} = 4.8$ under potential non-stationarity compared with 4.2 under assumed stationarity) stretches the distribution and extends the upper tail. The effect for much of the distribution appears fairly modest, but essential for economic applications is the interaction with the damage function (overlaid for reference). Steeply increasing damages amplify the importance of the tail of the distribution, where *relative* changes in probability are largest. For instance, the probability of an annual maximum rainfall event of 5 inches or more increases by 56% from 19.8% to 30.9% while the probability of an event of 8 inches or more more than triples.

4.2 The Damage Distribution

The implications of changes in the posterior probabilities of extreme rainfall for economic outcomes depends entirely on the impacts of different events, operationalized here through the damage function based on NFIP claims illustrated in Figure 2b. The posterior distributions over θ are propagated through the associated maximum rainfall distributions and then through the damage function to give distributions over damages under both updating models.

Table 1 shows how summary statistics of the damage distribution shift once the possibility of non-stationarity is integrated into the learning process. Even the fairly modest widening of the posterior rainfall distribution (Figure 2b) has a substantial effect on the damage distribution, raising expected

damages by just over 30% but more than doubling the variance. The largest impacts are concentrated in the tails of the distribution, with just a 17% increase in median damages but a 61% increase in the 1 in 200 year event (99.5th percentile).

4.3 Insurance Implications

Increased ambiguity over the climate distribution and, by extension, the nature of weather risks that property owners and insurers face, could have a range of implications for the functioning of property insurance markets. In this section I trace through these implications, taking the perspective of a single insurer underwriting the set of risks represented by the damage function in Figure 1b. In that sense, the damage function can be thought of as expected claims for the insurer conditional on the rainfall realization and its underwriting exposure.

4.3.1 Premium Prices and Volatility

One of the first-order effects of relaxing the stationarity assumption, made clear in Table 1, is a substantial increase in expected damages. Assuming that regulators allow premiums to adjust to reflect new understanding of risks under potential non-stationarity, this would produce a sudden increase in premiums of 32% (in line with the shift in expected losses). This increase occurs despite the fact that neither the weather data, historical loss data, nor the current damage function have changed. It is purely the result of the learner (namely the insurer) adjusting their updating model to integrate the possibility of a shifting climate distribution. The reasoning behind a sudden shift in average premiums may well be opaque to consumers (and potentially regulators), particularly in the absence of structural models of catastrophic risk able to integrate anthropogenic climate change effects (addressed further in the Discussion section).

The additional uncertainty over catastrophic events creates a problem for consumers not just from higher premium prices, but also from price volatility. Insurance contracts are renewed each year, allowing insurers (subject to regulatory approval) to rapidly adjust prices in response to new climatological information. However, volatile and unpredictable insurance prices create challenges for property owners since relevant decisions that impact exposure to insurance price volatility (namely over location, ownership, and mortgages) are long-term, forward looking decisions that can not be easily adjusted in response to changing insurance costs.

Figure 3 shows how greater ambiguity in the climate distribution could lead to more volatility in insurance premiums, particularly very extreme new events. The figure shows expected losses are for both updating models under both the observed 30-year record and a modified record where the

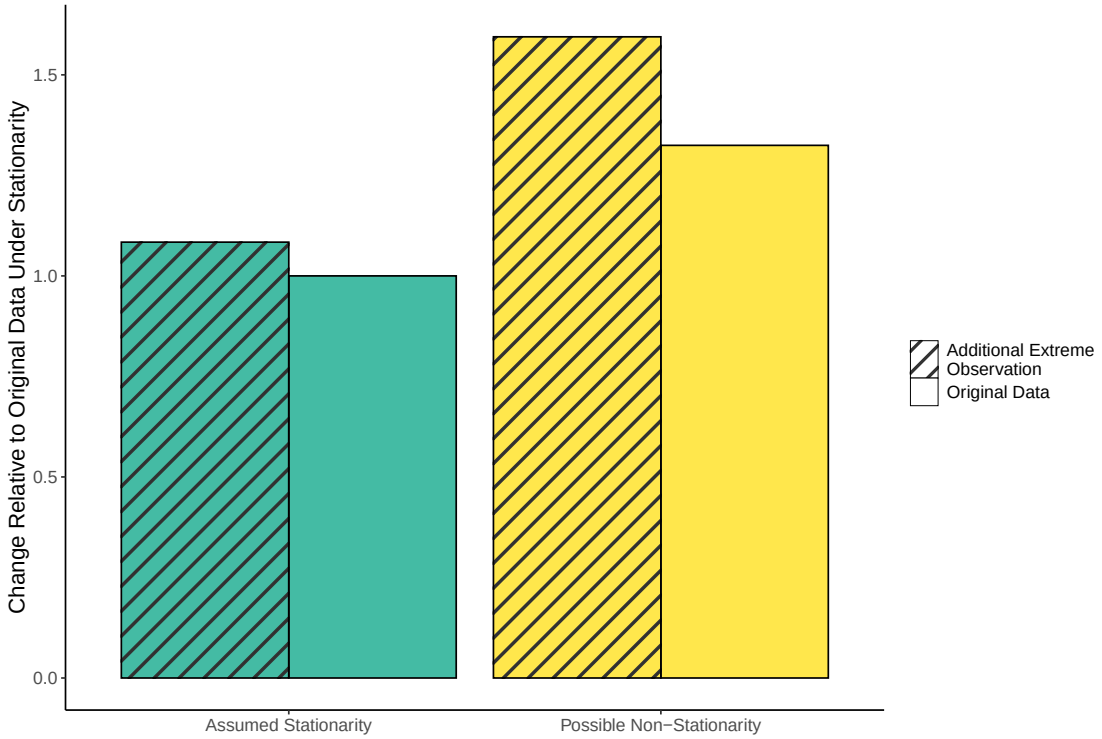


Figure 3: **Effect of Extreme Events on Expected Damages and Premiums.** Shows expected damages for 2024 both assuming stationarity and allowing for possible non-stationarity. Shaded bars show expected damages when the final observational datapoint has been adjusted to an extreme value, slightly larger than the maximum in the 30-year climatology. Values are shown normalized to the level in the stationary updating case using original weather data.

final observation is altered to an extreme value slightly larger than the previous maximum value. The additional extreme observation alters agents' beliefs about the underlying climate distribution, shifting the posterior distribution and raising expected losses. The effect is much larger, however, if agents believe the climate may be changing: expected losses increase 8% under assumed stationarity but 20% with possible non-stationarity in response to the new extreme observation. This arises because the agent is far less confident regarding parameter values under potential non-stationarity and therefore adjusts their beliefs far more in response to new observational evidence.

4.3.2 Loss Variance and Reinsurance Costs

Beyond the higher and more variable insurance premiums faced by consumers, the much larger variance in the damage distribution (Table 1) poses a challenge to insurers. A fundamental challenge of natural hazard risk for insurance is the correlated nature of losses; insurers must maintain access to large amounts of liquid capital in order to pay claims should a large event occur or risk bankruptcy. In the limit, over an infinitely long time horizon, premiums set at expected losses should cover total claims. But insurers need to be able to pay claims not just in the limit, but every time period they

are underwriting risks, including years immediately following a major disaster when any accumulated capital reserves are depleted. Insurers can address these risks by either reducing exposure to catastrophic risks by limiting underwriting (as we observe some firms doing in both the US Gulf Coast and California), attempting to diversify portfolios through exposure to other uncorrelated catastrophes, or passing risks on to global capital markets through reinsurance contracts or insurance-linked securities.

The additional uncertainty from a potentially non-stationary climate adds substantial variance to an insurer's position. In the case study used here, variance in the insurer's net position (i.e. aggregate claims minus total revenues, where revenues are set at expected losses) almost triples. Assuming regulator approval, insurers may be able to charge higher premiums in response to higher expected losses, but the increased variance of losses adds additional costs for the insurer not captured in expected loss. Conditional on a particular underwriting portfolio, insurers will have to pay more to transfer risks to reinsurers or capital markets because of the higher possibility of large losses. I illustrate this effect by simulating returns for a hypothetical, insurance linked security (ILS) that indemnifies the insurer up to losses equivalent to the most extreme event in the observational record for one year[7]. This guarantees the insurer will be able to pay claims for any event up to this threshold, but comes at a cost that compensates the investor for the risk of lost capital[8].

Figure 4 shows the distribution of losses faced by an investor in the ILS. Despite higher premiums under potential non-stationarity (arising from higher expected losses), expected loss for the security increases by almost 40% from 2.2% to 3.5% due to the longer tail of the climate distribution increasing the probability of very large losses. The probability that a large fraction of the collateral is lost increases even more substantially: the probability of a loss of 50% or more almost triples from 0.5% to 1.3%.

This changing loss distribution will affect the return insurers must pay investors to undertake the risk transfer. A number of papers have empirically examined the determinants of ILS pricing and suggest

[7] An insurance-linked security (ILS) is a contract between an investor and an insurer. The investor places collateral into a trust account that provides a base safe-asset return. The insurer pays an additional premium to the investor, essentially the price of the security. If a trigger event occurs then the contracted amount of the collateral is released to the insurer. If the term of the security ends without a trigger event, the collateral is returned to the investor. Triggers can be defined based on insurer losses directly (the case considered in the example), total industry losses, or parametric triggers related to physical variables such as hurricane intensity in a particular geographic region. ILS function similarly to reinsurnace but have far more observable prices compared to largely private reinsurance contracts, which is why I use them as a motivation in this example.

[8] For the time being I abstract from any potential for spatial or temporal smoothing. Spatial smoothing through diversification across independent catastrophic risks is discussed later in the paper. Temporal smoothing is more complicated for insurers since it requires them to amass large capital reserves to pay claims in the event of large but unlikely losses. As discussed in Jaffe and Russell [17], capital market structures make this challenging. Insurers are not able to credibly earmark retained earnings to pay out future claims, and would be liable for tax on both the earnings set aside and any interest earned by that capital. Moreover, accumulation of large reserves could make firms target of hostile takeovers and could attract scrutiny from rate regulators given the appearance of large profits being generated from excessive premiums.

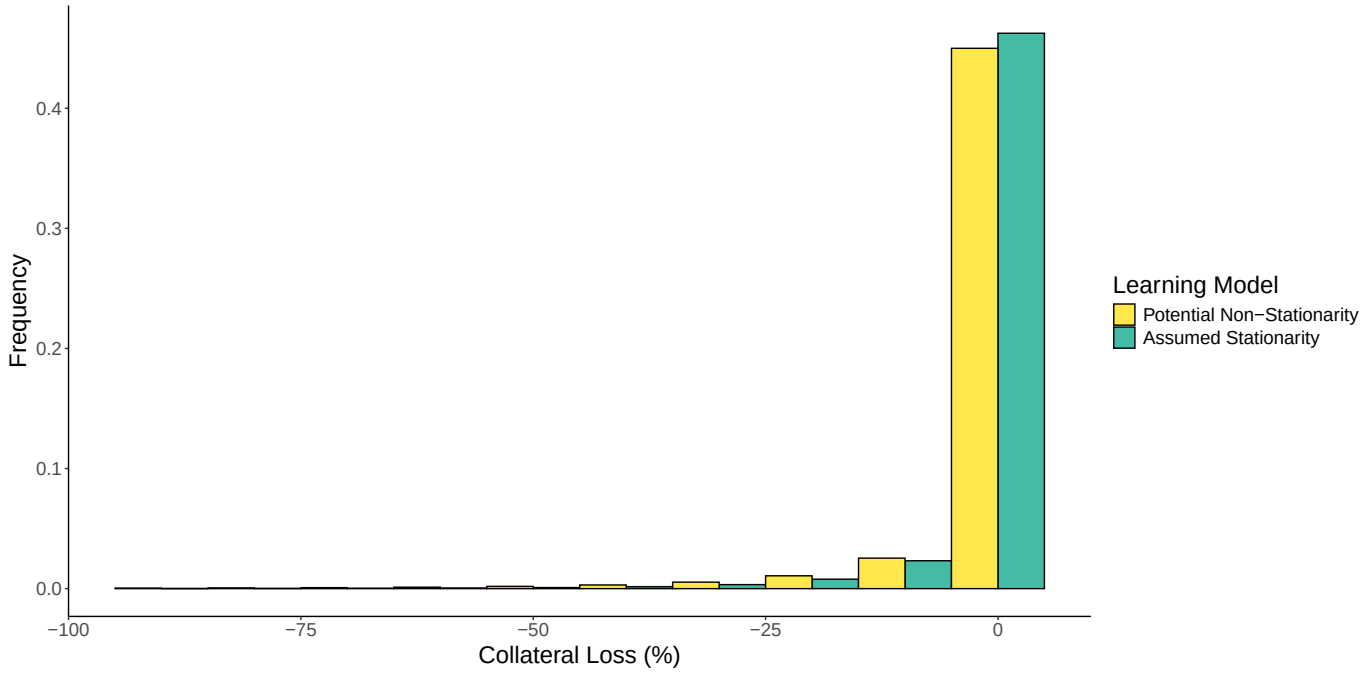


Figure 4: **Loss distribution for a security indemnifying an insurer up to a particular loss level.** Histogram shows the fraction of security collateral lost by the investor under both updating models. Losses to the insurer are defined as aggregate damages minus total premiums, where total premiums are set at expected loss (and are higher in the case of potential non-stationarity relative to assumed stationarity). If aggregate damages are less than total premiums, then the investor incurs no loss.

investors require a substantial premium to hold catastrophic risk. For instance, Braun [4] examines pricing of 437 ILSs issued between 1997 and 2012 and reports a mean spread of 10 times the expected loss, with a median of 4.8 and a minimum of 1.6[9]. Lane and Mahul [21] perform an original analysis that re-models risk statistics for 213 ILSs, enabling them to report how prices vary not just with expected loss but with other moments of the loss distribution. They find evidence that loss variance, including standard deviation and tail value at risk ($TVaR$) are associated with ILS spreads.

I use one of Lane and Mahul's models integrating tail risk metrics to illustrate potential effects on reinsurance costs. They estimate the relationship:

$$PremiumSpread = ExpectedLoss + 0.054 TVaR_{99}$$

where $TVaR_{99}$ is the expected loss conditional on reaching the 99th percentile of the loss distribution. Under this model, the costs of risk transfer for the insurer in terms of premium spread on an ILS increase 43% from a spread of 5.6% over the safe asset return to 8.0%. However, higher risk transfer

[9]The fact that ILSs command *any* premium over the safe asset return and expected loss (let alone the large premium documented in the literature) is perhaps surprising. The standard capital asset pricing model links the risk premium to the covariance between asset returns and broader market volatility. Since natural hazard risk is almost by definition uncorrelated with market returns, one might expect little to no risk premium, but that does not match available evidence on ILS prices.

costs for the insurer are not accompanied by lower risk of insurer bankruptcy. Rather, bankruptcy risk also increases under potential non-stationarity. Probability of an event exceeding the largest event in the full 155 year weather record (and therefore exceeding the indemnity limit for the hypothetical ILS) approximately quadruples from 0.08% to 0.32%.

4.3.3 Diversification

One question is whether sufficient diversification can ameliorate the effect of greater uncertainty in the climate distribution and associated risk profile faced by insurers. By underwriting multiple, uncorrelated risks simultaneously, insurers can lower the variance in their net position. Figure 5 simulates the effect of diversification on insurer positions and the interaction with updating processes. Rather than assume insurers face catastrophic risks exclusively in 1 location, the simulation assumes insurers spread the same exposure equally across n independent markets, all facing the same climatology and damage function.

As Figure 5 shows, diversification across independent risks is an important tool for insurers, with variance dropping steeply as the number of markets grows. However, diversification does not mitigate the increased variance associated with a shift to potential non-stationarity. Variance in the non-stationary case is elevated relative to assumed stationarity, and the *relative* increase in variance remains relatively steady. With exposure concentrated in a single market, variance under non-stationarity is 2.8 times larger than under stationarity. With exposure spread over 10 markets, absolute variance falls by an order of magnitude but still remains 2.8 times larger than the case where stationarity could be assumed. The potential non-stationarity introduced by climate change is a systemic shock, simultaneously raising agents' uncertainty over damages in all markets, creating additional risk that cannot be diversified in property insurance markets alone.

5 Discussion and Conclusions

Climate is a statistical distribution over possible weather states. The climate at a particular place and time is not directly observable, but instead must be estimated using either past observations of weather, structural models of the climate system, or a combination of the two. Anthropogenic climate change, by rendering past weather observations potentially less informative of current risks reduces information available to constrain the current climate distribution and, by necessity, increases uncertainty in the present distribution of weather risks.[10] Like any uncertainty, this is costly to risk averse individuals

[10]Note that this uncertainty also impacts structural climate models such as General Circulation Models (GCMs). GCMs have the challenge of jointly estimating both the effect of greenhouse gas emissions on the climate system (the so-called "forced response") and the distribution of weather conditional on a particular climate state, both of which are uncertain. Historic observations provide only a single draw from the historic climate distribution and must be used

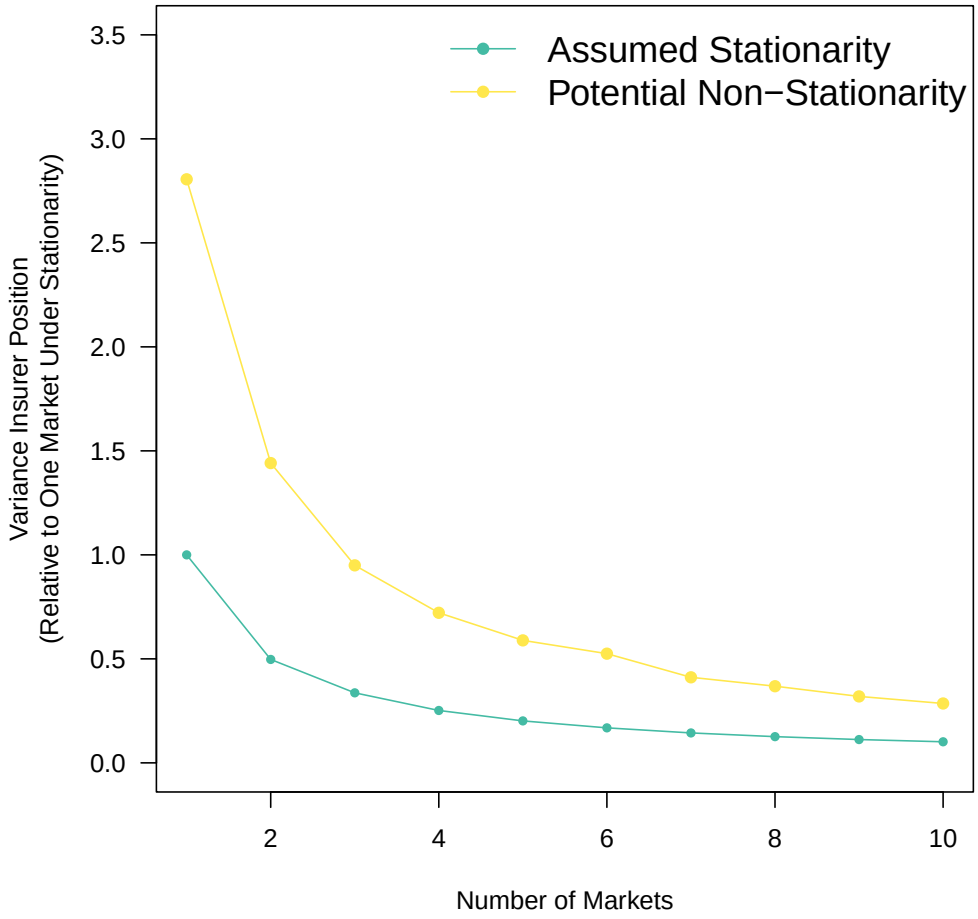


Figure 5: **Effect of diversification on insurer position variance.** Shows the variance in the distribution of premiums (set at expected losses) minus aggregate claims for an insurer with the same total exposure, but split equally across n markets, where n varies from 1 to 10. Shown relative to variance for the stationary case in a single market.

and investors, but the costs of this added uncertainty due to lost information from a non-stationary climate are, as yet, entirely unquantified.

The implicit assumption of stationarity in the climate distribution has been deeply embedded in how institutions understand and manage weather risk. For instance, methods for designing engineered systems from standards for property construction to the specification requirements for urban drainage systems, rely on the assumption that the envelope of natural weather variability these systems will face can be recovered from the observational record [24]. Catastrophe modeling, used by the insurance industry to estimate and price catastrophe risk, resamples the observational record of weather extremes while overlaying current maps of property locations and vulnerability to estimate losses were those

to evaluate model simulations of both the forced response and internal climate variability. If the forced response were known to be zero, internal climate variability could be better constrained with the same information.

events to occur today. Effectively then, these models assume the distribution of past weather events is representative of today.

While the limitations of the stationarity assumption are increasingly well-recognized, the question of how to adapt risk-management approaches to account for anthropogenic climate change is not resolved. While evidence from GCMs does provide general indications of trends in some extremes (such as increasing heat waves or more intense rainfall and droughts), the ability of GCMs to generate reliable, probabilistic information on extreme distributions at spatial and temporal scales relevant for risk-management and insurance pricing, is not established. GCMs are designed to project long-term, global changes in temperature from elevated greenhouse gas emissions, primarily as a tool to inform global emissions targets. While the models have an excellent track record at this task [14], risk-management applications are very different [12]. Recent papers evaluating performance for these applications cast doubt on models' ability to capture even the direction of change for key variables relevant to both wildfire and hurricane risks [33, 32].

The recognition that the stationarity assumption is inappropriate, with no well-tested methods to replace lost information from the weather record necessarily increases ambiguity in the current and near-future climate distribution. While climate change itself is a relatively gradual, long-term process, this increase in ambiguity can occur abruptly as actors the possibility of a shifting climate into their assessment of current weather risks. The impacts of this added uncertainty is likely to be largest for extreme event risk. Extremes are rare, by definition, in the historical record, meaning long observational records are particularly valuable in constraining current probabilities. Moreover, given non-linear damage functions and long-tailed weather distributions, expected values are heavily influenced by unlikely but highly consequential outcomes. A loss of confidence in tail probabilities could place substantial upward pressure on expected losses.

Simulations presented in this paper illustrate how these effects could ripple through property insurance markets, raising actuarily-fair premiums, premium volatility and reinsurance costs. While higher insurance rates may be important in providing information on changing risks, efficient adaptation will not be accomplished through insurance pricing alone. Decisions over where and how to build, where to live and work, and when and how to purchase property are long-term and forward-looking. Insurance by contrast is sold annually and only covers current risks. Price volatility in insurance poses challenges for consumers, who are not able to rapidly adjust mortgage, location, or property ownership decisions in response.

Further market pressures could arise from higher costs of risk transfer for insurers in response to reinsurers and investors also altering beliefs over the distribution of climate risks (Section 4.3.2). Unlike

standard insurance premiums, reinsurance costs are not regulated. Evidence from ILS prices indicates that catastrophic risk transfer is expensive, with ILS spreads substantially higher than corporate bonds with comparable risk [21, 4], despite the diversification advantage offered by these assets. Insurers will need to pass higher reinsurance costs on to consumers to remain profitable, but this risks raising premiums above expected losses for individual consumers. Uptake of natural hazard insurance, when not required by lenders or regulators, is generally very low [22] and higher rates that are either actually or perceived to be far above expected loss will only exacerbate this market unraveling.

Alternate options for insurers unable or unwilling to purchase risk transfer are to limit exposure to catastrophic risk entirely by limiting policy writing in exposed areas or, if permitted by regulators, to operate at higher risk of bankruptcy. Even more than rapidly increasing premiums, sudden insurer exits from areas rendering insurance unavailable at any price create challenges for property owners, most of whom are committed to 30 year mortgages that require an insurance policy. Insurer bankruptcies, several of which have been seen in recent years in Florida and Louisiana following major hurricanes, risk destabilizing local insurance markets more generally [30]. Losses from bankrupt insurers are assessed on the remaining admitted insurers in the state via State Guarantee Associations, putting additional financial pressure on those firms. Consumers that lose confidence in insurance institutions will be even less willing to pay higher premiums for insurance contracts that may not be paid out.

Climate change poses clear but not insurmountable challenges for U.S. property insurance markets. Insured losses from natural disasters averaged less than $45 billion per year since 2000 (in 2020 dollars) [15], a vanishingly small fraction of an economy of $27 trillion. Natural hazard insurance plays an important role in disaster recovery for those affected [2, 20] and smooths the functioning of property and mortgage markets [29, 3], meaning maintaining access to insurance coverage in most areas will likely be an important part of climate adaptation. At the same time, risk transfer is not risk reduction, and policies to stabilize insurance markets in the face of climate change will not by themselves substantially lower the net costs of climate change. If poorly designed, policies to address insurance availability could end up subsidizing development in the riskiest areas and perversely *increasing* total climate change costs.

A Appendix

A.1 Damage Function Estimation

The damage function used for the case-study illustration in this paper is estimated using the universe of National Flood Insurance Program (NFIP) claims for New York City [11]. Claim amounts for 184

	Full Data	Excluding 2012
$RMax$	0.654*	0.611*
	0.219	0.233
$\log(P)$	1.075	1.202
	1.055	1.096
$TMax$	1.300***	1.615**
	0.182	0.487
Num.Obs.	15	14
R2	0.867	0.766
RMSE	0.95	0.96

* p < 0.05, ** p < 0.01, *** p < 0.001

Table 2: **Damage Function Regression** Regression results for damage function used for case study illustration in paper. Dependent variable is the total NFIP claims paid in New York City, for the period 2009-2023. $RMax$ is daily maximum rainfall, $TMax$ is daily maximum tide height, and P is total flood policy coverage in New York City. Errors are treated as iid.

New York City zip codes are aggregated to the annual level (i.e. total insured flood damages for the city) and merged with data on total flood insurance coverage for the city [10]. The time series runs from 2009 (the first year for which coverage data is available) to 2023. Claim amounts and coverage are converted into real 2020 dollars using the consumer price index from the St Louis Federal Reserve.

The damage function relating annual maximum rainfall with insured flooding damage is estimated using a simple regression controlling for coverage levels and annual maximum tide height (using tide gauge data from The Battery in New York City [26]) [11]. The estimated regression is:

$$log(C_t) = \beta_0 + \beta_1 RMax_t + \beta_2 TMax_t + \beta_3 log(P_t) + \epsilon_t$$

where C_t is total NFIP paid claims in New York City in year t, $RMax_t$ is the maximum daily rainfall at the Central Park station in year t, $TMax_t$ is the maximum daily tide height at The Battery station, and P_t is the total flood policy coverage in New York City in year t.

Regression results are shown in Table 2, showing large and highly statistically significant effects for maximum tide height and substantial effects for maxiumum daily rainfall, significant at the 5% level. Results imply that a 1 inch increase in maximum daily rainfall increases NFIP claims for the year by approximately 65%. Because Hurricane Sandy was such an extreme outlier (in terms of both tide

[11] Flooding can occur either from intense rainfall overwhelming artificial and natural drainage systems (i.e. pluvial flooding) or coastal flooding in which sea water floods onto land, typically during major storms (i.e. storm surge). The period used here includes Hurricane Sandy in 2012, which caused intense storm-surge-related flooding in New York City. The inclusion of maximum tide height controls helps isolate the rainfall-induced flooding, which is the motivation for the example in this paper).

height and flood damage) for New York City in this period, Table 2 also shows results of a regression model dropping the 2012 outlier. The estimated magnitude and direction of extreme rainfall (and maximum tide) on flood claims is robust to dropping this outlier.

Regression results are used to construct a damage function connecting the weather variable of interest, $RMax$, to aggregate flood insurance claims in 2024. The damage function is specified using the average value over $TMax$ for the 2009-2023 period, and the level of policy coverage in 2023, on the assumption that 2024 would be most similar to 2023. Predicted values for $log(C)$ are converted into predicted values for C using the non-parametric Duan smearing estimator [8].

$$C_t = e^{\hat{\gamma_0}+\hat{\beta_1}RMax} * \frac{1}{n}\Sigma e^{\epsilon_t}$$

Where $\hat{\gamma_0} = \hat{\beta_0} + \hat{\beta_2}\overline{TMax} + \hat{\beta_3}P_{2023}$ and $\frac{1}{n}\Sigma e^{\epsilon_t}$ is the Duan smearing term for a log transformed dependent variable. This procedure gives the exponential damage function shown in Figure 1b.

References

[1] Tim P Barnett et al. "Human-induced changes in the hydrology of the western United States." In: *Science* 319 (5866 2008), pp. 1080–1083.

[2] Stephen B. Billings, Emily A. Gallagher, and Lowell Ricketts. "Let the rich be flooded: The distribution of financial aid and distress after hurricane harvey". In: *Journal of Financial Economics* 146 (2 2022). ISSN: 0304405X. DOI: 10.1016/j.jfineco.2021.11.006.

[3] Kristian Blickle and João A. C. Santos. "Unintended Consequences of "Mandatory" Flood Insurance". In: *SSRN Electronic Journal* (2022). DOI: 10.2139/ssrn.4086765.

[4] Alexander Braun. "Pricing in the Primary Market for Cat Bonds: New Empirical Evidence". In: *Journal of Risk and Insurance* 83 (4 2016). ISSN: 15396975. DOI: 10.1111/jori.12067.

[5] R. Daniel Bressler et al. "Estimates of country level temperature-related mortality damage functions". In: *Scientific Reports* 11 (1 Dec. 2021), p. 20282. ISSN: 2045-2322. DOI: 10.1038/s41598-021-99156-5. URL: https://www.nature.com/articles/s41598-021-99156-5.

[6] Tamma A Carleton et al. "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits". In: *Quarterly Journal of Economics* (2022).

[7] Marc N. Conte and David L. Kelly. "An imperfect storm: Fat-tailed tropical cyclone damages, insurance, and climate policy". In: *Journal of Environmental Economics and Management* 92 (2018), pp. 677–706. ISSN: 10960449. DOI: 10.1016/j.jeem.2017.08.010. URL: https://doi.org/10.1016/j.jeem.2017.08.010.

[8] Naihua Duan. "Smearing Estimate: A Nonparametric Retransformation Method". In: *Journal of the American Statistical Association* 78 (383 Sept. 1983), pp. 605–610. ISSN: 0162-1459. DOI: 10.1080/01621459.1983.10478017. URL: http://www.tandfonline.com/doi/abs/10.1080/01621459.1983.10478017.

[9] National Center for Environmental Information. *WMO Climate Normals*. 2022. URL: https://www.ncei.noaa.gov/products/wmo-climate-normals.

[10] FEMA. *OpenFEMA Dataset: FEMA NFIP Redacted Policies -v2*. May 2024. URL: https://www.fema.gov/openfema-data-page/fima-nfip-redacted-policies-v2.

[11] FEMA. *OpenFEMA Dataset: FEMA NFIP Redated Claims-v2*. May 2024. URL: https://www.fema.gov/openfema-data-page/fima-nfip-redacted-claims-v2.

[12] Tanya Fiedler et al. "Business risk and the emergence of climate analytics". In: *Nature Climate Change* 11 (2 Feb. 2021), pp. 87–94. ISSN: 1758-678X. DOI: 10.1038/s41558-020-00984-6. URL: http://www.nature.com/articles/s41558-020-00984-6.

[13] Fliegelman. *Wind, Fire, Water, Hail: What Is Going on In the Property Insurance Market and Why Does It Matter?* Dec. 2023. URL: https://www.financialresearch.gov/the-ofr-blog/2023/12/14/property-insurance-market/.

[14] Zeke Hausfather et al. "Evaluating the Performance of Past Climate Model Projections". In: *Geophysical Research Letters* 47 (1 2020). ISSN: 19448007. DOI: 10.1029/2019GL085378.

[15] Insurance Information Institute. *Natural Catastrophe Losses in the United States by Peril*. 2024. URL: https://www.bea.gov/news/2024/gross-domestic-product-fourth-quarter-and-year-2023-third-estimate-gdp-industry-and.

[16] Insurance Information Institute. *Residual Markets*. 2023. URL: https://www.iii.org/publications/a-firm-foundation-how-insurance-supports-the-economy/a-50-state-commitment/residual-markets.

[17] Dwight M. Jaffee and Thomas Russell. "Catastrophe Insurance, Capital Markets, and Uninsurable Risks". In: *The Journal of Risk and Insurance* 64 (2 1997). ISSN: 00224367. DOI: 10.2307/253729.

[18] Charles D. Kolstad and Frances C. Moore. "Estimating the Economics Impacts of Climate Change Using Weather Observations". In: *Review of Environmental Economics and Policy* 14 (1 2020), pp. 1–24.

[19] C. Kousky. *Understanding Disaster Insurance: New Tools for a More Resilient Future*. Island Press, 2022.

[20] Carolyn Kousky. *The Role of Natural Disaster Insurance in Recovery and Risk Reduction*. 2019. DOI: 10.1146/annurev-resource-100518-094028.

[21] Morton N. Lane and Olivier Mahul. "Catastrophe Risk Pricing An Empirical Analysis". In: *Policy Research Working Paper Series* 4765 (November 2007).

[22] Kendra Marcoux and Katherine R. H. Wagner. "Fifty Years of U.S. Natural Disaster Insurance Policy". In: *SSRN Electronic Journal* (2023). DOI: 10.2139/ssrn.4456363.

[23] K. Marvel and C. Bonfils. "Identifying External Influences on Global Precipitation". In: *Proceedings of the National Academy of Sciences* 110 (48 2013), pp. 19301–19306. ISSN: 0027-8424. DOI: 10.1073/pnas.1314382110.

[24] Author P C D Milly et al. "Stationarity Is Dead : Stationarity Whither Water Management ?" In: *Science* 319 (5863 2008). ISSN: 0036-8075. DOI: 10.1126/science.1151915.

[25] NOAA. *NOAA Tides and Currents*. 2020. URL: https://tidesandcurrents.noaa.gov/.

[26] NOAA. *NOAA Tides and Currents: The Battery, NY - Station ID: 8518750*. 2024. URL: https://tidesandcurrents.noaa.gov/stationhome.html?id=8518750.

[27] Volker C. Radeloff et al. "Rising wildfire risk to houses in the United States, especially in grasslands and shrublands". In: *Science* 382 (6671 2023). ISSN: 10959203. DOI: 10.1126/science.ade9223.

[28] Benjamin D Santer et al. "Identifying human influences on atmospheric temperature." In: *Proceedings of the National Academy of Sciences of the United States of America* 110 (1 Jan. 2013), pp. 26–33. ISSN: 1091-6490. DOI: 10.1073/pnas.1210514109.

[29] Parinitha Sastry. "Who Bears Flood Risk? Evidence from Mortgage Markets in Florida". In: *SSRN Electronic Journal* (2022). DOI: 10.2139/ssrn.4306291.

[30] Parinitha Sastry, Ishita Sen, and Ana-Maria Tenekedjieva. "When Insurers Exit: Climate Losses, Fragile Insurers, and Mortgage Markets". In: *SSRN Electronic Journal* (2024). DOI: 10.2139/ssrn.4674279.

[31] W Schlenker and D L Roberts. "Nonlinear Temperature Effects Indicate Severe Damages to U.S. Corn Yields Under Climate Change". In: *Proceedings of the National Academy of Sciences* 106 (37 2009), pp. 15594–15598.

[32] Isla R. Simpson et al. "Observed humidity trends in dry regions contradict climate models". In: *Proceedings of the National Academy of Sciences of the United States of America* 121 (1 2024). ISSN: 10916490. DOI: 10.1073/pnas.2302480120.

[33] Adam H. Sobel et al. "Near-Term tropical cyclone risk and coupled Earth system model biases". In: *Proceedings of the National Academy of Sciences of the United States of America* 120 (33 2023). ISSN: 10916490. DOI: 10.1073/pnas.2209631120.

[34] SwissRe. *Sigma Explorer*. 2024. URL: https://www.sigma-explorer.com/.

[35] Oliver E. J. Wing et al. "New insights into US flood vulnerability revealed from flood insurance big data". In: *Nature Communications* 11 (1 Dec. 2020), p. 1444. ISSN: 2041-1723. DOI: 10.1038/s41467-020-15264-2. URL: http://www.nature.com/articles/s41467-020-15264-2.

[36] Gordon Woo. *Calculating catastrophe*. 2011. DOI: 10.1142/P786.

[37] Xuebin Zhang et al. "Detection of human influence on twentieth-century precipitation trends." In: *Nature* 448 (7152 July 2007), pp. 461–5. ISSN: 1476-4687. DOI: 10.1038/nature06025.