

Risk Management

Sensitivity of Complex Reinsurance Pricing Metrics to Model Assumptions

Shree Khare¹ , Soadad Farhan¹ , Keven Roy¹ 

¹ Moody's Analytics, Ltd., UK

Keywords: Reinsurance, Technical pricing metrics, catXL, Reinstatements, Aggregate losses, Model sensitivity, Annual loss volatility

<https://doi.org/10.66573/001c.165180>

Variance

Vol. 19, 2026

Using a representative multiperil catastrophe model applicable to a US nationwide industry portfolio, this paper quantifies the sensitivity of reinsurance technical pricing metrics to model assumptions. We show that shifting risk profiles from low to high catastrophe excess of loss layers (or vice versa) involves clear trade-offs. These trade-offs are useful in formulating strategy under climate change. Our findings are applicable to reinsurers whose portfolios mimic a proportional slice of the US nationwide industry portfolio, but the framework developed here can also be applied to general cases. Additionally, this work yields insights into recent property catastrophe market dynamics.

1. INTRODUCTION

The earth's climate system gives rise to events such as hurricanes, mid-latitude windstorms, wildfires, and severe convective storms. The insurance industry creates value by enabling transfer of risk arising from such events. For example, homeowners pay annual premiums to insurers for the right to claim recoveries after damaging events. Insurers purchase reinsurance to diversify risk and improve profitability. Pricing of reinsurance contracts is dependent on economic factors, market dynamics, and technical metrics derived from risk models. This paper focuses on technical pricing metrics from the point of view of the reinsurer.

The motivation behind this paper was to increase understanding of how climate change and reinsurance pricing may interact. A major premise of climate change science is that human-induced increases in greenhouse gas concentrations, such as carbon dioxide and methane, have a direct causal link to large-scale changes in the earth's climate

system. This claim is supported by the peer-reviewed scientific literature (see Lynas et al. 2021, which describes the authors' analysis of 88,125 climate-related papers written between 2012 and 2020). Climate change has recently been front and center in the insurance industry, which is clear from articles written by a spectrum of stakeholders, including rating agencies and analytics firms (Moody's 2024), reinsurance brokers (WTW 2024), regulatory bodies (Hielkema 2023), consulting firms (Fantini et al. 2023), accounting firms (PWC 2024), actuarial organizations (Climate Change Working Party 2020), reinsurers (Zaitsev 2021), and popular science publications (Frank and E&E News 2023).

The idea that climate change impacts the frequency and severity of natural catastrophes also has scientific consensus (Seneviratne et al. 2021; Calvin et al. 2023)—suggesting that natural catastrophe risk models that fail to account for climate change must be “perturbed” in some fashion, raising questions, such as: What impact do these model per-

a Shree Khare currently leads three model development teams at Moody's RMS responsible for climate change, global modeling, and real-time event response. Shree has 20+ years experience in risk modelling and reinsurance. His academic training includes physics (Honors BSc, UBC), quantitative finance (MSc, University of York), and atmospheric modeling (PhD, Princeton University). Shree has contributed to the development of numerous catastrophe risk models used widely across the global financial sector, alongside novel research contributions to data assimilation, catastrophe modeling, and mathematical aspects of reinsurance pricing. Address for correspondence: shree.khare@moodys.com.

b Soadad Farhan is a senior technology and innovation leader with more than 30 years of international experience in software engineering, digital transformation, and insurance technology. A former chief technology officer and head of innovation at Hiscox, he has also held senior technology roles with the Commonwealth Bank of Australia and Moody's RMS. His career spans re/insurance, financial services, and enterprise technology, with a strong focus on delivering complex digital solutions and driving innovation.

c Keven Roy leads the development of the climate change conditioned versions of the Moody's RMS catastrophe models, covering a wide range of natural hazards. He has a background in geophysics (PhD, University of Toronto) with extensive experience in climate research, as well as postdoctoral experience focusing on sea level studies (Nanyang Technological University, Singapore). Prior to joining Moody's, he worked in physical climate risk assessment and emissions modeling for a large asset manager and in catastrophe model evaluation in the (re)insurance industry.

turbations have on reinsurance pricing metrics? What fundamental understanding and/or framework can reinsurers use to handle perturbed risk profiles arising from climate change? This is a complex and important problem that few studies have investigated (e.g., Khare 2024; Khare and Roy 2021; note that we use the notation provided in Khare 2024). To advance understanding in this area, we used a top-down approach that represents the natural catastrophe risk model setup using a simplified frequency and severity model appropriate to a US-wide industry portfolio covering four climate perils (hurricane, windstorm, wildfire, and severe convective storm). The risk model employed here is the same as that in Khare and Roy (2021), with the earthquake peril excluded. Anecdotally, we note that private sector catastrophe risk models used in practice are based on complex bottom-up models using cutting-edge climate science, engineering, financial concepts, and large-scale cloud computing (see Golnaraghi et al. 2018 for a discussion of these topics). Our simplified top-down approach allows us to run a comprehensive sensitivity analysis enabling fundamental insights (traded for a degree of realism).

Using a representative US nationwide portfolio, this paper makes the following contributions:

- Explores the sensitivity of reinsurance price metrics to model parameter perturbations and contract assumptions for a wide range of catastrophe excess of loss (catXL) layers
- Presents a framework for how reinsurers can consider strategy when dealing with model perturbations
- Clarifies various trade-offs made by reinsurers when increasing/decreasing the risk profile of underwritten reinsurance layers
- Provides insights into how reinsurers can manage climate change from a technical perspective
- Presents a framework for understanding the implications of recent property catastrophe market dynamics (as of 2024)

We emphasize that our work is applicable to reinsurers whose portfolios are well approximated by proportional slices of the US nationwide industry portfolio; to some degree this limits the generalizability of this work. However, the method of analysis and framework presented in this paper can be applied to more general cases, as will be discussed. The paper is formatted as follows: Section 2 reviews mathematical definitions and assumptions, provides details on the base (unperturbed) natural catastrophe model, and discusses both analytical and numerical approaches to formulating catastrophe excess of loss (catXL) reinsurance pricing metrics. Section 3 briefly discusses the risk profile of the base natural catastrophe model, describes our numerical experiments and numerical accuracy, presents core results, and ends by discussing the fundamental insights derived from our results. Section 4 provides a summary, discusses the implications of our work for reinsurer strategy under climate change, sheds light on recent market dynamics, and discusses future research directions. Appendices A, B, and C provide mathematical details that underpin the analytical and numerical methods we employed.

2. PRELIMINARIES

This section discusses various mathematical definitions and assumptions, the multiperil frequency and severity model used as our base natural catastrophe model, analytical formulations for the catXL reinsurance contract central to this work, and various numerical schemes we employed to compute contract pricing metrics.

2.1. MATHEMATICAL DEFINITIONS AND ASSUMPTIONS

The following provides elementary definitions for completeness and begins by defining the frequency component of our representative natural catastrophe model.

Definition 2.1. Frequency distribution: The random number of events over a specified time interval (one year in this work) is denoted by N and has a discrete frequency distribution $P_N(k)$, where $k \in [0, 1, 2, \dots, \infty)$, and $\sum_{k=0}^{\infty} P_N(k) = 1$.

In this work all frequency distributions are Poisson, which is defined as follows.

Definition 2.2. Poisson frequency distribution: The Poisson frequency is

$$P_N(k) = \frac{e^{-\lambda} \lambda^k}{k!},$$

where the expectation $\mathbb{E}[N] = \lambda$ is a positive real number representing the annual average rate of events and the variance $\mathbb{V}[N] = \mathbb{E}[N]$.

For a given event occurrence, the random financial loss X (for a particular portfolio of physical assets) is drawn from a severity distribution defined as follows.

Definition 2.3. Severity distribution: The continuous severity density is

$$f_X(x),$$

where $x \in [0, \infty)$ (and $\int_0^{\infty} f_X(x) dx = 1$).

Individual peril severities are modeled using continuous gamma distributions defined as follows.

Definition 2.4. Gamma distribution: The gamma distribution is

$$g_X(x) = \frac{x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha) \beta^{-\alpha}},$$

where $x \in (0, \infty)$, the shape parameter $\alpha > 0$, the scale parameter $\beta > 0$, and $\mathbb{E}[X] = \alpha\beta$ and $\mathbb{V}[X] = \alpha\beta^2$.

For simplicity we impose the assumption that samples from $P_N(k)$ and $f_X(x)$ are drawn independently (this is common in many applications, see Shi et al. 2015 for generalizations). The combination of the frequency $P_N(k)$ and severity $f_X(x)$ represents a mathematical abstraction of a risk model applied to a specific portfolio of assets. Our top-down approach allows us to rapidly explore the effects of alternative model assumptions, which is the central focus of this work.

Timeline simulation is a central concept in risk modeling, so reviewing the concept is helpful for setting up the reinsurance pricing problem. It involves sampling both $P_N(k)$ and $f_X(x)$ in a sequence that generates losses for a specified set of simulation years and is defined as follows.

Definition 2.5. Timeline simulation: A timeline simulation for S years is defined by the following process:




Figure 1. Sensitivity of reinsurance pricing metrics. The figure depicts a timeline simulation for a given frequency $P_N(k)$ and severity $f_X(x)$. Each simulation period is one year in length (x-axis), and the y-axis shows the primary insurer portfolio loss in dollars before any reinsurance is applied. The annual maximum loss events are depicted by the 1s with red circles; the second annual maximum losses are indicated by the 2s with yellow circles, and so on. Event losses that are less than or equal to $\$X < \A have values in the closed interval from $\$0$ to $\$X$. The dashed red and blue lines indicated by the $\$A$ and $\$E$ values correspond to the attachment and exhaustion point of a catXL reinsurance contract that will be applied in what follows. Events with losses below $\$X$ are not depicted.

1. For all years indexed by $i = 1, \dots, S$,
2. Draw a random number from the frequency distribution $P_N(k)$,
3. For each year i let the random realization be $N = k_i$, which implies that we must take k_i random (loss) samples from the severity $f_X(x)$ (all samples are tagged with year i).

For clarification, see the visualization in [Figure 1](#).

[Figure 1](#) depicts a timeline simulation generated from a given $P_N(k)$ and $f_X(x)$. The y-axis represents loss to a primary insurer before the application of reinsurance; “pre-catXL” on the y-axis makes clear the losses are depicted before applying a catXL contract. The catXL contract is represented by the attachment point $\$A$ and exhaustion point $\$E$. Losses within the layer (up to a prespecified annual limit) are the responsibility of the reinsurer. [Figure 1](#) also explicitly denotes the order statistics (Barakat and El-Shandidy 2004; Buhrman 1973; Consul 1984; David and Nagaraja 2003; Gupta and Gupta 1984; Young 1970; Khare 2024; Khare and Roy 2021). The red circled 1s depict the annual maxima, the yellow circled 2s the second annual maxima, and so on. In practical applications the order statistics are typically referred to as occurrence losses. We now discuss our parameterization of $P_N(k)$ and $f_X(x)$ for the

base natural catastrophe model, followed by a discussion of technical catXL pricing metrics.

2.2. MULTIPERIL BASE CLIMATE NATURAL CATASTROPHE MODEL FOR A US NATIONWIDE PORTFOLIO

Our overarching aim is to shed light on the impact of model uncertainty (arising from sources such as climate change) in reinsurance pricing using a representative multiperil base natural catastrophe model. Several factors were considered in formulating the base model:

- Maintaining consistency with previously published actuarial and risk literature
- Deriving model parameter settings from publicly available information
- Using commonly applied frequency and severity distributions
- Providing a straightforward way to perform sensitivity analysis (relevant to climate change/variability)
- Representing a large industry level portfolio in a major reinsurance market like the United States

The model setup employed in Khare and Roy (2021) meets these various considerations. The authors used publicly available information to parameterize the frequency and severity of several perils applicable to a large US nationwide portfolio of insured assets (see Khare and Roy 2021 for detailed justification). The model consists of four climate perils: hurricane (HU), winter storm (WS), wildfire (WF), and severe convective storm (SCS). The base model parameters are extracted from Khare and Roy (2021, Table 1) and repeated here in [Table 1](#).

The parameters in [Table 1](#) represent a range of important perils comprising both high frequency, low severity perils (SCS with $\mathbb{E}[N_{SCS}] = 100$, where N_{SCS} denotes the discrete random variable) and low frequency, high severity perils (HU with $\mathbb{E}[N_{HU}] = 2$). WF is high frequency but has much higher severity volatility ($CV = 8$) compared with SCS ($CV = 4$), consistent with the nature of the perils. The multiperil climate natural catastrophe model is formulated as a mixture. For the frequency distribution, as each individual peril is Poisson and independent of all other perils, the multiperil frequency random variable is taken to be $N = N_{HU} + N_{WS} + N_W + N_{SCS}$, and is Poisson distributed with $\lambda = \mathbb{E}[N] = \mathbb{E}[N_{HU}] + \mathbb{E}[N_{WS}] + \mathbb{E}[N_{WF}] + \mathbb{E}[N_{SCS}] = 178$ (e.g., Krishnamoorthy 2015). The severity distribution for each peril is gamma distributed. A unique set of α and β parameters are chosen for each peril given the target average annual loss (AAL), expected frequency, and severity distribution coefficient of variation (CV), where $CV = \frac{\sqrt{\mathbb{V}[X]}}{\mathbb{E}[X]}$. For example, the expectation of the severity for HU is equal to $\$12.5$ billion divided by 2 ($\mathbb{E}_{HU}[N] = 2$), which is $\$6.25$ billion; combined with the target CV of 5 we can infer the α and β parameters listed in the table. The multiperil severity distribution $f_X(x)$ is formulated as a mixture model where the weights are in proportion to the rates given by

$$f_X(x) = w_{HU}f_{HU,X}(x) + w_{WS}f_{WS,X}(x) + w_{WF}f_{WF,X}(x) + w_{SCS}f_{SCS,X}(x),$$

Table 1. Assumed parameters for the individual peril frequency and severity distributions used to formulate the base climate natural catastrophe model.

Peril	AAL (\$ B)	$\mathbb{E}[N]$	CV	α	β	$\mathbb{E}[X]$	$\mathbb{V}[X]$
HU	12.5	2	5	0.04	1.56e11	6.25e9	9.73e20
WS	2.5	6	3	0.11	3.75e9	0.5e9	1.54e18
WF	2.5	70	8	0.016	2.19e9	0.36e8	7.67e16
SCS	10.0	100	4	0.063	1.6e9	0.1e9	1.61e17

Notes: AAL = average annual loss; CV = coefficient of variation; the frequency distribution is assumed to be Poisson for each peril (hence $\mathbb{V}[N] = \mathbb{E}[N]$).

Column 1 provides the peril name abbreviations. Column 2 provides the AAL in billions of dollars for each peril individually against what is presumed to be a large US nationwide portfolio. Column 3 shows the mean number of events per year. Column 4 provides the coefficient of variation for each peril's respective severity distribution. Columns 5 and 6 give the α and β parameters, respectively, of the gamma distribution used to model the severity for each peril. Each peril-specific gamma distribution has expectation $\mathbb{E}[X] = \alpha\beta$, variance $\mathbb{V}[X] = \alpha\beta^2$ and coefficient of variation $CV = \sqrt{\frac{\mathbb{V}[X]}{\mathbb{E}[X]}} = \frac{1}{\sqrt{\alpha}}$. Columns 6 and 7 provide the mean and variance of each peril's severity, respectively. Values in the table differ slightly from actual numbers due to rounding.

with $w_{HU} = \frac{\mathbb{E}[N_{HU}]}{\mathbb{E}[N]}$, $w_{WS} = \frac{\mathbb{E}[N_{WS}]}{\mathbb{E}[N]}$, $w_{WF} = \frac{\mathbb{E}[N_{WF}]}{\mathbb{E}[N]}$, and $w_{SCS} = \frac{\mathbb{E}[N_{SCS}]}{\mathbb{E}[N]}$ (note that $w_{HU} + w_{WS} + w_{WF} + w_{SCS} = 1$, and $\int_0^\infty f_X(x) = 1$), and $f_{HU,X}(x)$ is the hurricane-specific severity (which is similarly depicted in other perils). As stated above, the multiperil frequency and severity distributions are sampled independently.

2.3. ANALYTICAL FORMULATIONS OF CATXL PRICING METRICS

This work studies the catXL contract with reinstatements (Mata 2000; Cummins et al. 1999), providing analysis of loss metrics from the point of view of the reinsurer. For a catXL contract, the reinsurer is responsible for the payout of all event losses, as shown in Figure 1, beginning just above the attachment point $\$A$ up to an event maximum of $\$(E - A)$, subject to an annual aggregate limit $A_{lim} = (n + 1)(E - A)$, where the integer $n \geq 0$. Note that n is called the number of reinstatements because it represents a multiplier on the “width of the layer” given by $\$(E - A)$. To proceed we must define an annual aggregate loss random variable that depends on the contract parameters $\$A$, $\$E$, A_{lim} . As in Khare and Roy (2021), the mean and standard deviation of this annual aggregate loss random variable will be used to compute a “technical price,” which is the focus of this work. We start by defining the annual aggregate loss random variable Y prior to the application of any reinsurance (pre-catXL). Using classical compound notation (Klugman et al. 2008) we have the following definition.

Definition 2.6. Pre-catXL annual aggregate loss random variable: The annual aggregate loss is

$$Y = X_1 + \dots + X_N,$$

where $X_i \sim f_X(x)$ for all $i = 1, \dots, N$, and the number of events is random where $N \sim P_N(k)$.

We now define the annual aggregate loss random variable for the reinsurer. First, for a given layer defined by attachment point $\$A$ and exhaustion $\$E$, we must apply an operator to all samples of $f_X(x)$. This operator results in a transformed random variable: $X_{AE} \equiv \min[\max[0, X - A], E - A]$ (all losses below attachment $\$A$ are set to zero, the payout is $X - A$ up to the maximum $E - A$). We call X_{AE} the “post-catXL” severity

random variable. With X_{AE} in hand, we are ready to define the post-catXL annual aggregate loss random variable.

Definition 2.7. Post-catXL annual aggregate loss random variable: The annual aggregate loss is

$$Y_{AE} = \min[X_{1,AE} + \dots + X_{N,AE}, A_{lim}(n)],$$

where $X_{i,AE} = \min[\max[0, X_i - A], E - A]$ where $X_i \sim f_X(x)$ for $i = 1, \dots, N$, $A_{lim}(n) = (n + 1)(E - A)$ with reinstatements $n \in [0, 1, 2, \dots, \infty)$, and the random number of events $N \sim P_N(k)$.

We emphasize that Y_{AE} is defined for a particular fixed aggregate limit A_{lim} for a specified number of reinstatements n (not made explicit in our notation to keep it succinct). Reinsurers can use the distribution of Y_{AE} to quantify technical metrics that aid in the underwriting and portfolio management. Here we study the expectation $\mathbb{E}[Y_{AE}]$ and standard deviation $\mathbb{S}[Y_{AE}]$ (and their sum). This is motivated by common market practice where a technical pricing metric T is used, and is formulated as

$$T = \frac{\mathbb{E}[Y_{AE}] + \nu \mathbb{S}[Y_{AE}]}{P_{net}},$$

where ν is a volatility “loading factor,” and P_{net} is the net premium, which the reinsurer receives as compensation for taking responsibility for the losses Y_{AE} (Khare and Roy 2021). Other factors, such as diversification debit or credit in the context of the entire reinsurer portfolio, are typically considered in practice but are outside the scope of this work. The aim of this paper is to understand the effect of model parameter changes (around the base model specified in Section 2.2) on the metrics $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ and their sum $\mathbb{E}[Y_{AE}] + \mathbb{S}[Y_{AE}]$ for a variety of practically relevant contract setups.

To simplify the complex sensitivity analysis to follow, we have not included an aggregate deductible in defining Y_{AE} as was done in previous work (Hürlimann 2005; Mata 2000). Our notation is designed to be consistent with the recent paper on catXL pricing that motivated this work, Khare and Roy (2021).

2.4. COMPUTATIONAL METHODS FOR COMPUTING CATXL PRICING METRICS

Our objective is to compute the pricing metrics $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ for a given attachment $\$A$, exhaustion $\$E$, and specified number of reinstatements n , with aggregate limit

$A_{lim} = (n + 1)(E - A)$. All the main sensitivity results for any $n \geq 0$ presented in this work compute $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ using the fast Fourier transform (FFT) method discussed in Appendix C. We believe other methods could be employed to compute $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$, but we used the FFT method because it was found to be suitable for our application. To increase confidence in our results and implementation, we also conducted various accuracy tests, which are discussed in what follows.

2.4.1. ACCURACY CHECK ON FFT IMPLEMENTATION FOR COMPUTING $\mathbb{E}[Y_{AE}]$ AND $\mathbb{S}[Y_{AE}]$

To gauge the numerical accuracy of our FFT implementation, we can compare $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ computed using our FFT method with results derived from numerical integration, *specifically* for the case where $n \rightarrow \infty$. The numerical integral method to obtain $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ in the $n \rightarrow \infty$ limit is discussed in Appendix B. The integral formulas are solved using numerical methods found in the NumPy Python libraries (Harris, Millman, and van der Walt 2020). We assumed that if the FFT method compares well with the results derived from numerical integration for the $n \rightarrow \infty$ case, it is accurate for all cases. The justification for this is discussed in Appendix C.

2.4.2. TECHNICAL PRICING METRICS USING DISCRETE FOURIER TRANSFORMS: ALL CASES ($n \geq 0$)

Appendix C reviews the details behind our FFT based calculations of $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$. A few necessary details are provided here to make our implementation clear. The main idea is to compute a discrete vector of probabilities that accurately represents the post-catXL aggregate loss distribution for any A , E , and n . This is readily used to compute $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$.

A numerical grid is formed as $x_{[0]}, x_{[1]}, \dots, x_{[M-1]}$, where $x_{[0]} = 0$ and $x_{[M-1]} = L_{max} - A$, L_{max} is a fixed “maximum loss,” and M is the number of grid points. The grid spacing is constant for a given catXL layer with $\Delta x = x_{i+1} - x_i = \frac{L_{max}-A}{M}$ for $i = 0, \dots, M-2$. Appendix C formulates the probability vector $\mathbf{p}_{S,AE,n}$, which is the M -dimensional vector where $\mathbf{p}_{S,AE,n,[i]}$ is the probability of the annual loss Y_{AE} we assign to each discrete grid point from $i = 0$ to $i = M-1$ for a given n and $A_{lim} = (n + 1)(E - A)$.

From $\mathbf{p}_{S,AE,n}$, we then apply standard formulas to compute the technical pricing metrics as

$$\mathbb{E}[Y_{AE}] = \sum_{i=0}^{M-2} \mathbf{p}_{S,AE,n,[i]} x_i,$$

and

$$\mathbb{S}[Y_{AE}] = \sqrt{\sum_{i=0}^{M-2} \frac{\mathbf{p}_{S,AE,n,[i]} (x_i - \mathbb{E}[Y_{AE}])^2}{M-1}}.$$

We emphasize that the above is applicable to all cases for any combination of n , A , E including the pre-catXL case (if $A = 0$, $E \rightarrow \infty$ and $n \rightarrow \infty$). The accuracy of the probabilities (and in turn $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$) is contingent on the number of grid points M , and the choice of the

maximum modeled loss L_{max} . We used fixed values for M and L_{max} , which are chosen to achieve accurate results across the entire set of reinsurance layers and model configurations examined in this work (numerical accuracy results are discussed in Section 3). We use fast FFT methods where M is chosen as a power of 2 (to optimize computational speed), and L_{max} is large enough to provide an appropriate degree of “zero padding” required for accuracy. See Brigham (1988) for a complete discussion of discrete Fourier transform methods.

2.4.3. MEASURING NUMERICAL ACCURACY OF FFT BASED AGGREGATE LOSS METRICS

The accuracy of the FFT implementation is dependent on our choices for the maximum modeled loss L_{max} and grid size M . We now detail our method for comparing FFT and numerical integral results for the $n \rightarrow \infty$ case. As discussed above, we assume that if accurate results are obtained when $n \rightarrow \infty$, the FFT method is accurate for all cases. Our rationale is provided in Appendix C. The integral results for a given catXL layer are denoted $\mathbb{E}^I[Y_{AE}]$ and $\mathbb{S}^I[Y_{AE}]$, respectively (where we use superscript I to indicate integration). The analogous results for the FFT method are denoted $\mathbb{E}^F[Y_{AE}]$ and $\mathbb{S}^F[Y_{AE}]$ (for a given A , E , L_{max} , M , and $n \rightarrow \infty$). We use two methods to quantify accuracy: (1) absolute percentage errors to the targets $\mathbb{E}^I[Y_{AE}]$ and $\mathbb{S}^I[Y_{AE}]$ and (2) the “effective number of years of simulation” using standard errors (Wasserman 2004). Clearly $\mathbb{E}^F[Y_{AE}]$ is not a random variable estimated from independent and identically distributed samples from a well-defined distribution, but for the sake of assessing accuracy, we treat it as though it were, and make use of the central limit theorem (Wasserman 2004) to characterize the absolute error as

$$abs(\mathbb{E}^F[Y_{AE}] - \mathbb{E}^I[Y_{AE}]) = \frac{\mathbb{S}^I[Y_{AE}]}{\sqrt{S_{eq}}},$$

where $\mathbb{E}^I[Y_{AE}]$ and $\mathbb{S}^I[Y_{AE}]$ are deemed as the true values, $\mathbb{E}^F[Y_{AE}]$ is the FFT result (dependent on the numerical setup including M , the number of grid points), and S_{eq} is what we call the “equivalent number of simulation years” analogous to the number of independent and identically distributed samples. Clearly,

$$S_{eq} = \left(\frac{\mathbb{S}^I[Y_{AE}]}{abs(\mathbb{E}^F[Y_{AE}] - \mathbb{E}^I[Y_{AE}])} \right)^2$$

serves as a useful alternative measure of accuracy, since the number of years terminology is commonplace in applied risk modeling (Definition 2.5). The above formulation is used for the pre-catXL case where Y_{AE} is replaced by Y .

3. NUMERICAL EXPERIMENTS AND DISCUSSION

Sections 3.1, 3.2, and 3.3, respectively, cover the base model loss profile, the design of our numerical experiments, and details on the numerical accuracy of our FFT calculations, which sets the stage for the core numerical results presented in Sections 3.4, 3.5, and 3.6. Section 3.4 presents numerical results for two extreme cases made up of a low and high catXL layer. Section 3.5 discusses intermediate

cases lying between our low and high layers. Section 3.6 reviews key findings and insights garnered from our numerical work.

3.1. BASE MODEL LOSS DISTRIBUTION PROFILE

Figure 2 displays results from a $S = 500K$ year timeline simulation for the base model (Section 2.2). Annual aggregate losses are displayed for the four perils (WS, WF, SCS, and HU), including a histogram and bars that indicate the min, max, and interquartile ranges (losses are spread vertically for display purposes only). Figure 2 reveals interesting differences in the annual aggregate losses across the perils. For example, the HU losses are highly volatile and reach far into the tail (with the highest loss of order 10^{12}). SCS annual aggregate losses are much higher than those of WF and WS but also much less volatile, reflecting its high frequency and low volatility (compared with WF). The full group of multiperil losses are labeled “ALL” and depicted in the final set of results in Figure 2. The content in Figure 2 suggests that the contributions of the perils will depend on the particular catXL contract in question (with parameters $\$A$, $\$E$, and n). Figure 2 makes clear the materiality of the annual losses for the case studied in this work with significant probability above US\$100 billion.

Figure 3 displays various exceedance probability (EP, defined as 1 minus the cumulative probability) curves obtained from our $S = 500K$ year simulation (or by direct numerical integration in cases with analytical forms). Panels A, C, and D display return period on the vertical axis, which is the inverse of the EP. Panel A displays the annual aggregate EP for the multiperil model (ALL) and for each individual peril. Panel B displays the EPs of the multiperil severity, as well as the individual perils. Panel C displays the first five order statistic EPs for the multiperil model, as well as a blown up version in Panel D. Panel A makes clear the dominance of the HU peril to tail losses and suggests that contributions for layers with low attachments $\$A$ should come from all perils. Panel B demonstrates the stark differences in the severities for each peril, showing that HU has high loss per event (high area under the curve, noting the log scale on the x-axis), and WF has high volatility (flat severity EP). Panel D shows that the annual maximum loss (Order 1) is expected to dominate higher layers (with high attachment $\$A$), whereas the picture is more complex for lower layers (with low $\$A$ and $\$E$). In results to follow, we chose $\$A$ and $\$E$ by reading off loss thresholds consistent with particular return periods from the occurrence exceedance probability (OEP1), following what we believe to be common market practice.

3.2. DESIGN OF THE NUMERICAL EXPERIMENTS AND MODEL PERTURBATIONS

Broad scientific consensus is that anthropogenic climate change impacts both frequency and severity characteristics of natural catastrophe perils (Seneviratne et al. 2021; Calvin et al. 2023). Table 1 of Section 2.2 displays the various parameter choices associated with our base model. A natural question is how climate change will impact the




Figure 2. Annual aggregate loss distribution. This figure displays results from a $S = 500 K$ year timeline simulation derived from the multiperil frequency and severity model, which shows the annual aggregate losses for winter storm (WS), wildfire (WF), severe convective storm (SCS), and hurricane (HU). The colored dots represent the annual aggregate losses overlaid by a bar indicating the minimum, maximum, median, and interquartile ranges. A histogram is also provided for each peril. The final set of results (labeled ALL) displays the analogous results for the four perils aggregated. Note that this figure is identical to Figure 3 of Khare and Roy (2021), except for earthquake, which is excluded.

parameterization in Table 1 and what effect this has on technical metrics. We do not tackle this question directly; rather, we choose to display results for reasonable ranges of model parameters, hoping to garner interesting insights about the impact climate change may have on various technical reinsurance perspectives.

Table 1 suggests an approach to generating perturbed model configurations by varying the AAL, the expected frequency for each peril $\mathbb{E}[N]$, and the peril severity mean and standard deviation (and hence CV). We chose to generate the perturbations by varying one peril at a time and quantifying the impact on catXL pricing metrics, keeping all other perils fixed to the base model configuration:

- Results are displayed across a range of perturbation factors $p_1 \in [.5, 1.5]$ and $p_2 \in [.5, 1.5]$
- Perturbation parameter p_1 simultaneously multiplies the AAL and $\mathbb{E}[N]$ with the same factor
- Perturbation parameter p_2 multiplies the standard deviation of the severity distribution




Figure 3. Exceedance probabilities for various models. Panel A depicts the annual aggregate EP for the individual peril models, as well as the multiperil model denoted by ALL. Panel B depicts the EP of the individual severity distributions, as well as the multiperil model. Panels C and D depict the EP curves for order statistics up to $M = 5$. Panel D is a zoomed in version of Panel C. This figure is identical to Figure 4 of Khare and Roy (2021) except for the earthquake peril, which is excluded.

Our p_1 frequency parameter enables us to explore scenarios where the frequency and overall losses from AAL of a given peril changes in future states, exploring both 50% decreases and increases. By choice, p_1 multiplies both AAL and $\mathbb{E}[N]$ simultaneously. This is motivated by the notion that future states with higher frequency are also likely to incur higher AALs due to economic exposure growth. Our p_2 perturbations reflect future states with 50% increases or decreases in severity volatility (which are intended to capture a range of possible outcomes under climate change). By design, our model perturbations do not change the mean of the peril severities (or the combined model). We generated numerical results for a variety of low and high catXL layers where the attachment $\$A$ and exhaustion $\$E$ were chosen to be consistent with following return period thresholds (from the OEP1 of the multiperil base model): 0–2, 2–5, 5–10, 10–20, 20–50, and 50–100 years return period. Our layer definitions are designed to mimic typical industry practice as much as possible to make our results relevant for practitioners. To quantify the impact of reinstatements, we generated results for $n \in [0, 1, 2, 4, 8, 16]$ and the unlimited case with $n \rightarrow \infty$. Results to follow are displayed in a space where the y-axis is proportional to p_2 and the x-axis is proportional to p_1 using contour plots of $\mathbb{E}[Y_{AE}]$, $\mathbb{S}[Y_{AE}]$ and their sum.

3.3. NUMERICAL ACCURACY OF DISCRETE FOURIER TRANSFORM CALCULATIONS OF TECHNICAL PRICING METRICS

The pricing metrics discussed in this work were computed using the FFT implementation outlined in Section 2.4.2 and detailed in Appendix C. For all results we set the maximum modeled loss to $L_{max} = \$10,000$ billion and used a uniformly spaced grid of dimension $M = 2^{25} > 33e6$. This implies a grid spacing of approximately 0.000298 billion dollars, orders of magnitude smaller than the mean losses per event for the high frequency, low severity perils WF (≈ 0.035 billion dollars) and SCS (0.1 billion dollars) (base model configuration). To assess numerical accuracy we used benchmarks obtained using the integrals in Appendix A (pre-catXL) and Appendix B (post-catXL). Numerical integrals were computed using a standard python NumPy package. We deemed integration errors to be at most $1e-8$. Independent code was developed in R achieving nearly identical results. The numerical integration results are therefore suitable benchmarks for the FFT results. We used the percentage error and “equivalent simulation years” method discussed in Section 2.4.3. Results are displayed in Table 2 (base model).

The results in Table 2 show that our FFT implementation is accurate. For example, in the pre-catXL case, the AAL and standard deviation have errors of less than 0.1% and yield an equivalent number of years of simulation, roughly

Table 2. Results from testing the accuracy of our FFT implementation.

Layer	AALFT	AALI	SDFT	SDI	%-AAL	%-SD	Years
pre-catXL	27.5	27.490	45.436	45.436	< 0.1%	< 0.1%	20e6
0–2 RP	15.370	15.327	5.210	5.210	0.3%	< 0.1%	4000e6
50–100 RP	0.982	0.997	7.698	7.807	1.5%	1.4%	0.3e6

Column 1 indicates the case studied (“pre-catXL” denotes that no reinsurance layer is applied, and 0–2 RP denotes a catXL layer with loss thresholds from the 0–2 year return period of the multiperil OEP1 for the base model configuration, and so on). Columns 2 and 3 display the AAL for the FFT and integral method, respectively. Columns 4 and 5 display results for the annual loss standard deviation. Column 6 displays the absolute percentage error (rounded off to one decimal) for the AAL, and column 7 has the error for the standard deviation. Column 8 displays the “equivalent number of years of simulation.”

20e6 years, much larger than in most practical settings (typically up to 1e5 years). Low (0–2 RP) layer and high (50–100) layer results are accurate enough for our purposes. The higher layer errors are larger, likely from the lower attachment probability (consistent with our intuition). Since the results in Table 2 are for the base model configuration, we also tested the accuracy using various combinations of perturbations to the model as outlined in Section 3.2, yielding qualitatively similar accuracy (omitted for brevity). With assurance that our numerical scheme is accurate, we set the stage to present the core results.

3.4. NUMERICAL RESULTS

First we present the results for a low and high catXL layer in the case of unlimited reinstatements ($n \rightarrow \infty$).

3.4.1. LOW AND HIGH CATXL LAYERS: UNLIMITED REINSTATEMENTS

The low layer loss thresholds are defined by the base model multiperil OEP1 from 0 to 2 years return period, with $\$A = 0$ and $\$E = 3.79204$ (billion), and the high layer is for the 50 to 100 years layer with $\$A = 146.763$ billion and $\$E = 215.837$ billion.

Figure 4 displays the low layer results for HU model perturbations ($n \rightarrow \infty$). The upper panel displays the $\mathbb{E}[Y_{AE}]$ (AAL) with $\mathbb{E}[N_{HU}]$ on the x-axis, and CV for the HU severity on the y-axis (following our design in Section 3.2). The upper panel shows that increasing frequency drives up AAL (not surprisingly), but higher CV results in lower AAL. The middle panel displays results for the standard deviation $\mathbb{S}[Y_{AE}]$ in the same format, showing similar behavior to the AAL, and the lower panel displays the “risk metric” $\mathbb{E}[Y_{AE}] + \mathbb{S}[Y_{AE}]$, which amalgamates the upper and middle panel. Figure 4 demonstrates that the pricing metrics are indeed sensitive to HU model perturbations. The base model configuration is where $\mathbb{E}[N_{HU}] = 2$ and $CV = 5$, and we see that the risk metrics vary between (roughly) -5% and $+20\%$.

Figure 5 displays low layer results for WS perturbations. The results are similar to HU but the range of risk metrics is lower, and while increasing frequency always drives up $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$, the effect of increasing CV for WS differs from that of the HU case. For example, the middle panel shows a convex relationship of $\mathbb{E}[Y_{AE}]$ with CV (opposite of Figure 4 for HU). Hence, no simple rule of thumb

clarifies the effect of changing peril severity CV (for fixed frequency).

Figures 6 and 7 display analogous low layer results for WF and SCS perturbations, respectively. The results are qualitatively similar to those of the HU (Figure 4) and WS (Figure 5) cases but differ in the details. Figure 7 demonstrates that the results are most sensitive to changes in the specification of the SCS model parameters. The bottom panel shows an expanded range risk metric (-25% to $+20\%$) compared with that of HU (-5% to $+20\%$).

We now move to the high layer ($n \rightarrow \infty$). Figure 8 displays results for HU (analogous to Figure 4). The results are starkly different to the low layer. First, we see that increasing HU CV increases the AAL, $\mathbb{S}[Y_{AE}]$ and the risk metric (opposite to the effect seen in Figure 4). We also see that for the base model configuration, the ratio of $\frac{\mathbb{S}[Y_{AE}]}{\mathbb{E}[Y_{AE}]}$ in high layer case (7) is around 20 times the value for the low layer ($1/3$). The high layer losses are significantly more volatile (not surprisingly). The high layer also exhibits extreme sensitivity to the HU model perturbations, as revealed by the lower panel of Figure 8 (-85% to $+70\%$). This range is much higher than our results for the low layer (-5% to $+20\%$).

Figure 9 displays the risk metrics for the high layer experiments, with WS in the upper panel, WF in the middle panel, and SCS in the lower panel. Figure 9 implies nearly zero sensitivity to model perturbations for WS, WF, and SCS. This starkly contrasts the low layer results with non-trivial sensitivity to all peril perturbations. Apparently the losses from WS, WF, and SCS are simply not material enough to influence the losses to the high layer.

The unlimited reinstatement cases in Figures 4, 5, 6, and 7 (low layers) and Figures 8 and 9 (high layers) lead to the following summary:

- Low layers have the following properties: (1) Low volatility for annual losses (standard deviation per unit AAL); (2) Sensitivity to all four perils—with highest sensitivity from HU and SCS; (3) Increased risk metric (for fixed CV) at higher frequency, while the relationship with CV is peril dependent; and (4) A range of -25% to 20% in the risk metric.
- High layers have the following properties: (1) Annual losses are highly volatile; (2) Results are highly sensitive to HU model perturbations with (nearly) zero sensitivity to WS, WF, and SCS; (3) Risk metric is increased with frequency and CV (for HU); and (4) Sensitivity to HU model perturbations is high, with a range of -85% to 70% in risk metric.




Figure 4. Sensitivity results for the 0-2 RP low layer varying the HU model parameters only (all other parameters kept fixed). Results are for unlimited reinstatements ($n \rightarrow \infty$). The upper panel displays the AAL with annual average frequency on the x-axis and CV of the severity on the y-axis. The size of the perturbations represent 50% upward and downward perturbations of the parameters in the base model configuration. The center panel displays results for the annual loss standard deviation, and the lower panel displays the risk metric (sum of AAL and standard deviation). The center of each panel represents results for the base model configuration.

The above results represent two extremes on a spectrum. There are clear trade-offs. We now explore the effects of adding finite reinstatements.

3.4.2. LOW AND HIGH CATXL LAYERS: FINITE REINSTATEMENTS

Figure 10 displays low layer results with $n = 0$. The upper left panel displays HU, the upper right shows WS, the lower left shows WF, and the lower right show SCS. With $n = 0$, the aggregate limit is $A_{lim} = (E - A)$. This annual loss cut-off has a dramatic effect on the range of risk metric sensitivity. For example, the range of risk metrics for HU in the upper left is very close to the base model configuration with $\lambda_{HU} = 2$ and $CV = 5$. Qualitatively, the results for WS, WF, and SCS are similar. We also note that the risk metric itself is a factor of 6 lower than the case with $n \rightarrow \infty$ (Figure

4), and the risk metric is only slightly larger than $(E - A)$. In fact, for the upper left panel HU results, $\frac{S[Y_{AE}]}{E[Y_{AE}]} \approx 0.005$, compared with $1/3$ in the $n \rightarrow \infty$ limit (similar results for WS, WF, and SCS). For the low layer, including the reinstatement, ($n = 0$) nearly eliminates model sensitivity and annual loss volatility.

Figure 11 displays the analogous results to Figure 10, but for the case where $n = 1$ and the aggregate limit $A_{lim} = 2(E - A)$. Comparing Figure 11 with Figure 10, we see a larger range of risk metrics and more annual loss volatility for which $\frac{S[Y_{AE}]}{E[Y_{AE}]} \approx 0.05$ (10 times the $n = 0$ case, and similar for other perils). Adding reinstatements drives up model sensitivity and annual loss sensitivity. We expect this trend to continue as we add reinstatements.

Figures 12 and 13 display low layer results for $n = 4$ and $n = 8$, respectively, and demonstrates the trend dis-




Figure 5. Low layer results for WS perturbations. This figure displays the analogous results to Figure 4, except the WS (rather than HU) parameters are varied. The size of the upward and downward perturbations relative to the base model configuration for WS are the same as in Figure 4 (50%). The center of each panel represents results for the base model configuration.

cussed above that increasing the number of reinstatements enables higher model sensitivity and volatility ($\frac{S[Y_{AE}]}{E[Y_{AE}]}$ for $n = 4$ is slightly less $n = 8$, which is very close to the $n \rightarrow \infty$ case). The annual loss volatility appears to “saturate” once we reach $n = 8$.

Results for the high layer are depicted in Figure 14, which shows HU with $n = 0$ reinstatements in the upper panel, $n = 1$ in middle panel, $n = 8$ in the bottom panel. In stark contrast to the low layer, reinstatements have little to no effect on the volatility and/or model sensitivity of the high layer (we omitted results for other perils as there is no sensitivity even in the case for $n \rightarrow \infty$ as shown in Figure 9). For the high layer, reinstatements are not an effective “control” mechanism.

3.4.3. INTERMEDIATE LAYERS

To this point, we have presented results for (extreme) low and high layers under varying numbers of reinstatements. Our results suggest a progression from low to high layers when we consider model sensitivity, annual loss volatility, which perils are important (peril uncertainty), and the effect of varying the reinstatements n . We therefore expect the progression to evidence itself when we look at “intermediate layers” lying between the low and high layers. Figure 15 displays risk metric sensitivity for the 2 to 5 year layer where $\$A = 3.79204$ and $\$E = 8.5$ (billion). The upper left displays results for HU where $n \rightarrow \infty$ and the upper right for $n = 0$. The lower left displays results for WF with $n \rightarrow \infty$ and lower right with $n = 0$. Our results show the progression discussed above. As for the $n = 0$ case, the risk metric is more sensitive to HU than for WF, and the effect of including a reinstatement is slightly damped. Qual-




Figure 6. Analogous low layer results for WF perturbations. This figure also displays the analogous results to Figure 4, except the WF (rather than HU) parameters are varied. The size of the upward and downward perturbations relative to the base model configuration for WF are the same as those of the HU in Figure 4. The center of each panel represents results for the base model configuration.

itatively similar results to those of WF are achieved for WS and SCS (omitted for brevity).

Figure 16 displays results analogous to Figure 15, except for the 10–20 year layer from $\$A = 27.588$ to $\$E = 68.743$ (billion). Figure 16 gives more evidence for the progression in that we see high sensitivity to HU, little to no sensitivity for WF, and very minimal impact from finite reinstatements.

3.4.4. KEY FINDINGS FROM THE NUMERICAL EXPERIMENTS

Sections 3.4.1, 3.4.2, and 3.4.3 explored various dimensions important to formulating technical pricing metrics of catXL layers: peril model parameter uncertainty, the return periods associated with a given catXL layer, and the number of reinstatements constraining the annual aggregate loss. Our results suggest some general principles useful for reinsurers dealing with model perturbations.

Figure 17 provides a visual aid to clarify our results. The x-axis depicts the annual loss volatility $\frac{S[Y_{AE}]}{\mathbb{E}[Y_{AE}]}$, the y-axis shows the number of reinstatements for a given catXL layer, and the z-axis shows model sensitivity (the change in risk metric under model perturbations). Our result for the low layer with large ($n \rightarrow \infty$) number of reinstatements is represented by the red “low-RP” box. This result is characterized by low annual loss volatility, moderate model sensitivity, and high peril uncertainty (the risk metric is sensitive to all four perils). The yellow “low-RP” for small n is characterized by (near zero) low volatility and model sensitivity and moderate peril uncertainty. Higher frequency drives up the risk metric, but the effect of increasing severity CV is peril dependent. The other end of this spectrum shows the high layer results indicated by the two blue “high-RP” boxes, which are characterized by high model sensitivity, high annual loss volatility, but low peril uncertainty (only sensitive to HU changes). The high layer results are not sen-




Figure 7. Analogous low layer results for SCS perturbations. This figure displays the analogous results to [Figure 4](#), except the SCS (rather than HU) parameters are varied. The size of the upward and downward perturbations relative to the base model configuration for SCS are the same as for those of HU in [Figure 4](#). The center of each panel represents results for the base model configuration.

sitive to the inclusion of finite reinstatements. For high layers, higher HU frequency and CV drive up the risk metric. The intermediate layer results in Section 3.4.3 fit nicely between the two extreme cases, as we would expect.

Our results demonstrate clear trade-offs if reinsurers shift from low to high layers (or vice versa). With large numbers of reinstatements, the transition from low to high layers drives up model sensitivity and annual loss volatility. This increase in model sensitivity and annual loss volatility is “traded” for lower peril uncertainty. Including reinstatements can control model sensitivity and annual loss volatility, but only for low layers. We note that the risk metric relationship to changes in severity CV also changes as we move from low to high layers (from peril dependent to positive dependence on only HU CV for the high layer). Section 4 discusses the implications of our technical results for reinsurance strategies under climate change and provides context for recent market dynamics (as of 2024).

4. SUMMARY AND CONCLUSIONS

This paper presents a set of numerical experiments using a representative natural catastrophe risk model consisting of four climate perils (HU, WS, WF, and SCS). The base model (Section 2.2) is designed to mimic a US-wide insurance portfolio with realistic relativities across perils. Section 3.4 presents a set of sensitivity analyses where we vary model parameters (related to frequency and severity) for each peril independently. These sensitivity results are motivated by the notion that climate change represents perturbations to the natural catastrophe models used by reinsurers to formulate technical metrics that inform catXL pricing. We have systematically varied catXL contract parameters: (1) Attachment $\$A$ and exhaustion $\$E$ representing typical low, intermediate, and high layers in the market; and (2) number of reinstatements n . Our chosen technical metrics are the layer AAL and standard deviation (and their




Figure 8. Sensitivity results for the 50 to 100 RP high layer varying the HU model parameters only (all other parameters kept fixed). Results are for unlimited reinstatements ($n \rightarrow \infty$). The upper panel displays the AAL with annual average frequency on the x-axis and CV of the severity on the y-axis. The size of the perturbations represent 50% upward and downward perturbations of the parameters in the base model configuration. The center panel displays results for the annual loss standard deviation, and the lower panel displays the risk metric (sum of AAL and standard deviation). The center of each panel represents results for the base model configuration.

sum) motivated by market practice. Insofar as our top-down catastrophe model is relevant to the market, our results yield a number of new and fundamental insights. We find a stark contrast between reinsurance metrics for low (where $\$A$ and $\$E$ represent low return periods) and high (vice versa) catXL layers. For the specific US nationwide portfolio and four-peril representative model studied in this work, the main technical conclusions are stated below.

Low layer results are characterized as follows:

- Technical metrics exhibit low annual loss volatility (per unit AAL), moderate sensitivity to model parameter perturbations, and high sensitivity to all perils (high peril uncertainty) for unlimited reinstatements ($n \rightarrow \infty$)

- Finite reinstatements drastically reduce annual loss volatility and model sensitivity, and moderate peril uncertainty
- Higher frequency drives up the risk metric (as expected), but the relationship with severity CV is peril dependent

High layer results are characterized as follows:

- For unlimited reinstatements, technical metrics exhibit high (20x low layers) annual loss volatility, high sensitivity to model perturbations, and low peril uncertainty (since only HU model perturbations are meaningful)
- Finite reinstatements have no measurable impact on annual loss volatility and model sensitivity




Figure 9. Sensitivity results for the 50 to 100 RP high layer varying the WS (upper panel), WF (middle panel), and SCS (lower panel). Each panel displays the risk metric (sum of AAL and standard deviation). The sensitivities are flat, showing nearly zero sensitivity to model assumptions for lower severity perils.

- Higher HU frequency and CV drive up the risk metric

Our results for the low and high layers represent two ends of a spectrum. We have demonstrated that “intermediate layers” (with attachment \$ A and exhaustion \$ E lying between our chosen low and high layers) fit logically within this spectrum. Therefore, as we progress from low to high layers we see higher annual loss volatility, higher sensitivity to model parameter perturbations and less peril uncertainty.

The above conclusions yield practical insights for reinsurers whose portfolios are well approximated by a proportional slice of the US nationwide industry portfolio for the specific perils studied here. Logically, our results are most relevant for large reinsurers with significant market share in the US. While our study covers a use case of broad interest, the computational framework we have developed in this paper is of value to a broader set of portfolios and combinations of perils. The sensitivity analysis we used is conceptually simple and could be carried out by skilled actu-

aries and/or catastrophe risk analysts. Ideally, commercial and/or open source software would be available to perform the type of analysis presented in this paper. We also emphasize that our results are specific to the case where there are four covered perils (hurricane, winter storm, wildfire, and severe convective storm). In this work the peril uncertainty is negligible as we shift to higher layers. This is dependent on our specific modeling setup. For example, if we had included earthquake (although not strongly influenced by climate change) we would expect higher layers to be affected by both hurricane and earthquake. Also, under climate change, the relative importance of different perils may change, and could increase/decrease peril sensitivity in higher layers.

Despite the specificity of our results to a US nationwide portfolio subject to HU, WS, WF, and SCS, it is useful to ask how our results influence how we think about reinsurance strategies under climate change. Clearly, formulating an optimal reinsurance strategy under climate change is a




Figure 10. Sensitivity results (for the risk metric) for the low layer (0–2 RP) case where the catXL contract has $n = 0$ reinstatements. In stark contrast to Figures 4, 5, 6, and 7 for unlimited reinstatements, the above results show very low sensitivity to model assumptions.




Figure 11. Analogous to Figure 10 but for the case where the number of reinstatements $n = 1$. Compared with the $n = 0$ case in Figure 10, the results demonstrate a higher range of risk metrics and therefore higher sensitivity.

profound and multifactorial problem. Our reductionist approach, which focused on technical metrics and a specific case, is one useful step forward. For example, suppose as a reaction to climate change, a reinsurer decides to tilt their portfolio to lower layers, believing that tail risk is too uncertain under climate change. Our results imply lower annual loss volatility and model sensitivity. We also gain a degree of “control” if we limit the number of reinstatements (lowering both volatility and model sensitivity). However, a clear trade-off is made for higher peril uncertainty. In this scenario, the reinsurer must model all four perils ac-

curately, and this complexity is amplified by uncertainties from climate change. This is a nontrivial consideration, given the nature of perils such as SCS, with small scale features. On the flip side, if the reinsurer decides to underwrite higher layers in reaction to climate change, the result is higher volatility and model uncertainty but minimal peril uncertainty (since only HU is important). The high model uncertainty could be problematic if the reinsurer does not have a good grasp of climate change effects on HU. Steering the portfolio in one direction or the other as strategy there-




Figure 12. Analogous to Figures 10 and 11 but with the number of reinstatements $n = 4$. These results demonstrate that as the aggregate limit is increased we enable a higher sensitivity to model assumptions.




Figure 13. Analogous to Figures 10, 11, and 12 with the number of reinstatements set to $n = 8$. These results demonstrate that increasing numbers of reinstatements lead to higher sensitivity of the risk metric to model assumptions and yield similar sensitivity as that exhibited in the unlimited reinstatements case in Figures 4, 5, 6, and 7.

fore comes with trade-offs. Our results make these trade-offs clear.

Our work also provides insight on recent property catastrophe market dynamics (2024). Rate on line for global property catastrophe is near all-time highs, with considerable increases from 5 to 6 years ago (Guy Carpenter 2024). In this environment, reinsurers have tightened their terms and conditions and shifted their portfolios to higher layers (broadly speaking). This has implications for both insurers (cedants) and reinsurers. Insurers are left with more re-

tained risk in the lower layers; hence, increasing peril uncertainty (and lower volatility and model uncertainty). This raises the importance of complex perils such as WF, WS, and SCS for insurers. Market factors—such as pressure on ability to flex insurance rates due to regulation—exacerbate the complexity. For reinsurers, the move to higher layers introduces more volatility and model uncertainty (driven by HU). Reinsurers must ensure they are adequately compensated for this additional volatility, which may be amplified if the effects of near term climate change on HU are poorly




Figure 14. Results for the high layer (50 – 100 RP) varying HU model parameters for varying numbers of reinstatements ranging from $n = 1$ in the upper left panel to $n = 8$ in the lower right panel. The range of risk metrics in the above four panels is similar to the range seen in Figure 8 for the unlimited reinstatements case. The inclusion of reinstatements has little effect on the range of risk metrics obtained by varying model parameters.

quantified. Any near term market shock that puts downward pressure on reinsurance rate on line would be consequential for reinsurers.

Our work sets the stage for additional studies. To further address climate change, it would be useful to extend our study to the case where we specify a multivariate distribution for all the peril parameters. This requires scientific insight beyond the scope of this paper but would enable a probabilistic interpretation of our results. The present work could be generalized by varying several perils simultaneously (instead of one peril at a time as done here for simplicity). Our analysis can be made more complete and richer by including market factors such as how well reinsurers are compensated in premiums for the various catXL contracts studied in this work. To aid in such analyses,

we note that commercially available climate change conditioned catastrophe risk models are available to allow reinsurers forward looking views of risk. Although we focused on the reinsurer point of view, it would be useful to extend our analysis to consider insurer retained risk. Finally, the framework presented in this work can be applied to a wide range of specific reinsurer portfolios, for different peril combinations, for exposures in the US and around the globe. The insurance industry plays a critical role in mitigating the worst economic effects of climate change, and the authors encourage further work. The relatively recent adoption of natural catastrophe modeling by the banking and investment sectors raises the stakes even further.




Figure 15. Risk metrics for the 2 to 5 year intermediate layer. The upper left panel has HU in the $n \rightarrow \infty$ case. The upper right displays the $n = 0$ case, showing a notable impact on the risk metric but not as extreme as with the low layer. The bottom two panels show the analogous results for WF.




Figure 16. The 10–20 year intermediate layer in the same format as [Figure 15](#). The figure shows increased sensitivity to HU and diminished sensitivity to WF perturbations. In comparison with [Figure 15](#), Figure 16 shows diminishing impact from the inclusion of finite reinstatements.

Submitted: July 06, 2024 EDT. Accepted: May 23, 2026 EDT.

Published: August 11, 2026 EDT.




Figure 17. A qualitative depiction of the numerical results. The x-axis depicts the annual loss volatility $\frac{S[Y_{AE}]}{E[Y_{AE}]}$, the y-axis shows the number of reinstatements n , and the z-axis shows the model uncertainty reflecting the range of risk metrics for a given catXL contract layer. The red “Low-RP” box depicts the case with large ($n \rightarrow \infty$) reinstatements for the low layer with $\$A = 0$ and $\$E = 3.79204$ (billion). The yellow “Low-RP” box depicts low layer results with $n = 0$ reinstatements. Two high layer cases are depicted in the blue “High-RP” boxes. Both model uncertainty and annual loss volatility are (nearly) the same for the two “High-RP” cases, as variation in the number of reinstatements has little quantifiable impact. Peril uncertainty is indicated by the color spectrum in the lower right and quantifies the number of perils of the four perils (HU, WS, WF, and SCS) that influence model uncertainty for a given case (blue indicates low peril uncertainty).

REFERENCES

- Barakat, H. M., and M. A. El-Shandidy. 2004. "Order Statistics with Random Sample Size." *International Journal of Statistics* LXII: 233–46.
- Brigham, E. O. 1988. *Fast Fourier Transform and Its Applications*. Prentice-Hall.
- Buhrman, J. M. 1973. "On Order Statistics When the Sample Size Has a Binomial Distribution." *Statistica Neerlandica* 27: 125–26.
- Calvin, K., D. Dasgupta, G. Krinner, and others. 2023. *Climate Change 2023: Synthesis Report*. IPCC. <https://doi.org/10.59327/ipcc/ar6-9789291691647>.
- Climate Change Working Party. 2020. *Climate Change for Actuaries: An Introduction*. Institute and Faculty of Actuaries. <https://www.actuaries.org.uk/system/files/field/document/Climate-change-report-29072020.pdf>.
- Consul, P. C. 1984. "On the Distributions of Order Statistics for a Random Sample Size." *Statistica Neerlandica* 38: 249–56.
- Cummins, J. D., C. M. Lewis, and R. D. Phillips. 1999. "Pricing Excess-of-Loss Reinsurance Contracts against Catastrophic Loss." In *The Financing of Catastrophe Risk*, edited by K. A. Froot. The University of Chicago Press.
- David, H. A., and H. N. Nagaraja. 2003. *Order Statistics*. 3rd ed. Wiley.
- Fantini, L., B. Blanchard, S. Rath, P. Removille, S. Schwemer, and F. Mayeres. 2023. *An Insurance Risk Framework for Climate Adaptation*. Boston Consulting Group. <https://mkt-bcg-com-public-pdfs.s3.amazonaws.com/prod/an-insurance-risk-framework-for-climate-adaptation.pdf>.
- Frank, E. and E&E News. 2023. "Climate Change Is Destabilizing Insurance Industry." *Scientific American*. <https://www.scientificamerican.com/article/climate-change-is-destabilizing-insurance-industry/>.
- Golnaraghi, M., P. C. Nunn, R. Muir-Wood, et al. 2018. *Managing Physical Climate Risk: Leveraging Innovations in Catastrophe Risk Modelling: Research Brief*. The Geneva Association. https://www.genevaassociation.org/sites/default/files/research-topics-document-type/pdf_public/ga_risk_modelling_18112018.pdf.
- Gupta, D., and R. C. Gupta. 1984. "On the Distribution of Order Statistics for a Random Sample Size." *Statistica Neerlandica* 38: 13–19. <https://doi.org/10.1111/j.1467-9574.1984.tb01092.x>.
- Guy Carpenter. 2024. "Guy Carpenter Global Property Catastrophe Rate on Line Index, 2000–2024." <https://www.guycarp.com/insights/2024/01/global-rate-on-line-index-2000-2024.html>.
- Hielkema, P. 2023. *The Role of Insurers in Tackling Climate Change: Challenges and Opportunities*. EIOPA. https://www.eiopa.europa.eu/publications/role-insurers-tackling-climate-change-challenges-and-opportunities_en/#description.
- Hürlimann, W. 2005. "Excess of Loss Reinsurance with Reinstatements Revisited." *ASTIN Bulletin* 35: 211–38. <https://doi.org/10.1017/S0515036100014136>.
- Khare, S. P. 2024. "Order Statistic Exceedance Probability Sensitivities to Alternative Model Assumptions." *Variance* 17 (1): 1–21. <https://doi.org/10.66573/001c.92094>.
- Khare, S. P., and K. Roy. 2021. "Quantifying the Role of Occurrence Losses in Catastrophe Excess of Loss Reinsurance Pricing." *Risks* 9 (3): 1–38. <https://doi.org/10.3390/RISKS9030052>.
- Klugman, S., H. Panjer, and G. Willmot. 2008. *Loss Models: From Data to Decisions*. 3rd ed. Wiley.
- Krishnamoorthy, K. 2015. *Handbook of Statistical Distributions with Applications*. Chapman and Hall.
- Lynas, M., B. Z. Houlton, and S. Perry. 2021. "Greater than 99% Consensus on Human Caused Climate Change in the Peer-Reviewed Scientific Literature." *Environmental Research Letters* 16: 114005. <https://doi.org/10.1088/1748-9326/ac2966>.
- Mata, A. J. 2000. "Pricing Excess of Loss Reinsurance with Reinstatements." *ASTIN Bulletin* 30: 349–68. <https://doi.org/10.2143/AST.30.2.504640>.
- Moody's. 2024. "Physical and Transition Risk." Moody's on Climate. <https://www.moody.com/web/en/us/capabilities/climate-risk.html>.
- PWC. 2024. *Reinsurers Rank Their Top Risks*. Reinsurance Banana Skins Survey 2023. <https://www.pwc.co.uk/insurance/assets/pdf/reinsurance-banana-skins-2023.pdf>.
- Seneviratne, S. I., X. Zhang, M. Adnan, et al. 2021. "Weather and Climate Extreme Events in a Changing Climate." Chap. 11 in *Climate Change 2021: The Physical Science Basis*. Cambridge University Press. <https://doi.org/10.1017/9781009157896.013>.
- Shevchenko, P. V. 2010. "Calculation of Aggregate Loss Distributions." *The Journal of Operational Risk* 5 (2): 3–40. <https://doi.org/10.21314/IOP.2010.077>.
- Shi, P., X. Feng, and A. Ivantsova. 2015. "Dependent Frequency-Severity Modeling of Insurance Claims." *Insurance: Mathematics & Economics* 64: 417–28. <https://doi.org/10.1016/j.insmatheco.2015.07.006>.
- Wasserman, L. 2004. *All of Statistics: A Concise Course in Statistical Inference*. Springer.
- WTW. 2024. *Pricing Climate Risks Effectively*. WTW. <https://www.wtwco.com/en-gb/insights/campaigns/pricing-climate-risk-effectively>.
- Young, D. H. 1970. "The Order Statistics of the Negative Binomial Distribution." *Biometrika* 57: 181–86. <https://doi.org/10.1093/BIOMET/57.1.181>.
- Zaitsev, S. 2021. "Impact of Climate Change on Reinsurance." Hanover Re. <https://piu.org.pl/wp-content/uploads/2021/09/Impact-of-climate-change-on-reinsurance.pdf>.

APPENDICES

A. PRE-CATXL MEAN AND STANDARD DEVIATION

Our objective is to derive integral formulas for the expectation ($\mathbb{E}[Y]$) and standard deviation ($\mathbb{S}[Y]$) of the random variable Y from Definition 2.6 (annual aggregate loss prior to the application of the catXL). The strategy for the calculation is:

- Write down known analytical formulas for $\mathbb{E}[Y]$ and $\mathbb{S}[Y]$
- Derive equations for all the components of the analytical formulas

The above is possible for Y because we know the frequency and severity that generate the distribution of Y . A similar approach is possible when a catXL layer is applied but *only* when we have unlimited reinstatements with $A_{lim} \rightarrow \infty$. We take a different approach for the case with finite reinstatements and use discrete Fourier transforms (details in Appendix C).

For the multiperil model (Section 2.2) the frequency is $P_N(k)$ with severity $f_X(x)$ (which are independent) and together these generate the distribution of Y . We know that (see Klugman et al. 2008)

$$\mathbb{S}[Y] = \sqrt{\lambda(\mu^2 + \sigma^2)},$$

and

$$\mathbb{E}[Y] = \lambda\mu,$$

where for convenience we use the notation $\lambda = \mathbb{E}[N] = \mathbb{E}[N_{HU}] + \mathbb{E}[N_{WS}] + \mathbb{E}[N_{WF}] + \mathbb{E}[N_{SCS}] = 178$, $\mu = \mathbb{E}[X]$ and $\sigma^2 = \mathbb{V}[X]$.

We first compute $\mu = \mathbb{E}[X] = \int_0^\infty x f_X(x) dx$. As $f_X(x)$ is a mixture density, we know that its expectation is a weighted sum of expectations for the individual peril severities, so we immediately obtain

$$\mu = \sum_{i=1}^4 w_i \alpha_i \beta_i,$$

where $i = 1$ corresponds to HU, $i = 2$ corresponds to WS, $i = 3$ corresponds to WF and $i = 4$ to SCS. The weights are $w_1 = w_{HU} = \frac{\mathbb{E}[N_{HU}]}{\mathbb{E}[N]}$, $w_2 = w_{WS} = \frac{\mathbb{E}[N_{WS}]}{\mathbb{E}[N]}$, $w_3 = w_{WF} = \frac{\mathbb{E}[N_{WF}]}{\mathbb{E}[N]}$, and $w_4 = w_{SCS} = \frac{\mathbb{E}[N_{SCS}]}{\mathbb{E}[N]}$, and the peril-specific gamma distribution parameters (α, β) are from Table 1 (recalling that $\alpha_{HU}\beta_{HU}$ is the expectation of the gamma distribution for the hurricane severity).

The variance is $\sigma^2 = \mathbb{V}[X] = \mathbb{E}[X^2] - (\mathbb{E}[X])^2$. With $\mathbb{E}[X] = \mu$ in hand, to complete the calculation we must derive $\mathbb{E}[X^2]$ which is

$$\mathbb{E}[X^2] = \sum_{i=1}^4 \int_0^\infty w_i x^2 f_{i,X}(x) dx,$$

using the same indexing for i . Using integral tables we obtain

$$\mathbb{E}[X^2] = \sum_{i=1}^4 w_i \frac{\beta_i^2 \Gamma(2 + \alpha_i)}{\Gamma(\alpha_i)},$$

where $\Gamma()$ is the gamma function. With the above we have all the ingredients to compute λ, μ, σ^2 required to obtain $\mathbb{E}[Y]$ and $\mathbb{V}[Y]$.

B. POST-CATXL MEAN AND STANDARD DEVIATION VIA INTEGRATION (UNLIMITED REINSTATEMENTS)

Our objective is to compute the expectation ($\mathbb{E}[Y_{AE}]$) and standard deviation ($\mathbb{S}[Y_{AE}]$) for the post-catXL random variable Y_{AE} in the case where the aggregate limit $A_{lim} \rightarrow \infty$ (unlimited reinstatements). Recall that the severity random variable required to compute the distribution of Y_{AE} is $X_{AE} \equiv \min[\max[0, X - A], E - A]$. Let this random variable have generalized distribution $f_{X_{AE}}(x)$, which combined with the frequency $P_N(k)$ generates the distribution of Y_{AE} (with unlimited reinstatements). This is *not* the case when A_{lim} is finite (discussed in Appendix C). As in Appendix A, we know that Klugman et al. (2008)

$$\mathbb{S}[Y_{AE}] = \sqrt{\lambda(\mu_{AE}^2 + \sigma_{AE}^2)},$$

and

$$\mathbb{E}[Y_{AE}] = \lambda\mu_{AE}.$$

Above, λ is the base model rate (as in Appendix A), the severity mean is $\mu_{AE} = \mathbb{E}[X_{AE}]$, and the variance is $\sigma_{AE}^2 = \mathbb{V}[X_{AE}]$. Using known results (Klugman et al. 2008), the integral for μ_{AE} is given by

$$\mathbb{E}[X_{AE}] = \int_A^E (x - A) f_X(x) dx + (E - A) \int_E^\infty f_X(x) dx,$$

where $f_X(x)$ is the multiperil mixture density. Analogous to Appendix A, to compute the variance $\sigma_{AE}^2 = \mathbb{E}[X_{AE}^2] - (\mathbb{E}[X_{AE}])^2$ we need to compute $\mathbb{E}[X_{AE}^2]$, which, using standard results (Klugman et al. 2008), is given by

$$\mathbb{E}[X_{AE}^2] = \int_A^E (x - A)^2 f_X(x) dx + (E - A)^2 \int_E^\infty f_X(x) dx,$$

with $f_X(x)$ again denoting the multiperil mixture. The above shows all the necessary ingredients to obtain $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ as desired.

C. CATXL LAYER AGGREGATE LOSS DISTRIBUTIONS VIA DISCRETE FOURIER TRANSFORMS (ALL CASES)

Our objective is to compute the post-catXL pricing metrics $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ for any number of reinstatements, n with aggregate limit A_{lim} . Appendix B provided integral formulas for computing $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ in the case where $n \rightarrow \infty$. However, the method in Appendix B for computing $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ does *not* apply when n and A_{lim} are finite. Instead we use Fourier transform methods. Several key points must be noted before diving into the details.

- Our use of a Fourier transform based method was a pragmatic choice to solve the problem of computing $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$ for quantifying sensitivities. Our aim is to simply show that it is accurate. We assume that other methods could be applied. The method we chose is simply an (accurate) means to our end of understanding pricing metric sensitivity to model parameter perturbations.
- The overall strategy for the calculation is to first compute the distribution of Y_{AE} where $n \rightarrow \infty$. Practi-

cally this means we calculate a probability vector on a loss grid running from 0 to a specified maximum loss. A finite A_{lim} can then be applied to this “full” probability vector by assigning all probability beyond A_{lim} to the A_{lim} point. We then have a discrete probability vector corresponding to the full distribution of Y_{AE} , but crucially for any finite n and corresponding A_{lim} .

- We gauged the accuracy of the Fourier transform method by comparing it with results from numerical integration in the case where $n \rightarrow \infty$. We assumed that the Fourier transform results are accurate for all cases if accurate for the $n \rightarrow \infty$ case. Why is this valid? As shown in the calculations to follow, the first step in the Fourier transform method is to compute the full distribution for Y_{AE} where $n \rightarrow \infty$. But including a finite A_{lim} is straightforward, and we see no reason why this “truncation” as a final step should hamper the accuracy of our metrics $\mathbb{E}[Y_{AE}]$ and $\mathbb{S}[Y_{AE}]$.

We now provide the details behind our computation of the aggregate annual loss distribution for a given catXL layer with attachment A , exhaustion E and number of reinstatements n with annual limit $A_{lim} = (n + 1)(E - A)$ (Section 2.3). This requires four steps:

- Step 1 – Create a discrete representation of the multiperil severity $f_X(x)$ (Definition 2.3) but transformed through a catXL layer for a specific E and A .
- Step 2 – Recall the definition of the discrete Fourier transform (DFT) and its inverse (IDFT), the key tool in the calculation of the aggregate loss probabilities.
- Step 3 – Derive a formula for the discrete representation of the aggregate annual loss distribution for a given E and A under the assumption that $n \rightarrow \infty$.
- Step 4 – Account for finite reinstatements n and annual limit A_{lim} by truncating the probabilities calculated in Step 3. We note that the following is also applicable to the pre-catXL case by setting $A = 0$ and ignoring any exhaustion E or aggregate limit A_{lim} .

Step 1: We first create an M -dimensional grid consisting of points $x_{[0]}, \dots, x_{[M-1]}$ (or compactly $x_{[j]}$ with $j = 0, \dots, M - 1$), where $x_{[0]} = 0$, and the maximum loss is assigned $x_{[M-1]} = L_{max} - A$, with discretization size $x_{[j+1]} - x_{[j]} = \Delta x = \frac{L_{max} - A}{M} > 0$ for all j . Integration on the multiperil severity $f_X(x)$ is used to assign probabilities on the grid $x_{[j]}$ accounting for the layer. The multiperil severity $f_X(x)$ is a mixture of gamma distributions with support on $[0, \infty)$. Of course, our discrete representation has finite support, hence our choice of a maximum loss $L_{max} - A$. We now formulate the probabilities $\tilde{p}_{[0]}, \tilde{p}_{[1]}, \dots, \tilde{p}_{[M-1]}$ on the grid (where $\tilde{p}_{[0]}$ is associated with $x_{[0]} = 0$). To account for the catXL layer attachment A , we set $\tilde{p}_{[0]}$ to all the mass from 0 to A plus the additional mass from $x_{[0]}$ to $x_{[1]}$ and therefore,

$$\tilde{p}_{[0]} = \int_0^{A+\Delta x} f_X(x) dx.$$

In this work we use numerical integration to compute the above and all similar types of integrals. Next we must find

the grid point index $m \in [1, \dots, M - 1]$, which we can associate with the exhaustion E (L_{max} is always chosen so that $E < (L_{max} - A)$). We do so by finding the largest m such that $x_{[m]} \leq (E - A)$. With m in hand, we assign the following probability to grid point $x_{[m]}$

$$\tilde{p}_{[m]} = \int_E^{L_{max}} f_X(x) dx,$$

which is all the probability beyond the exhaustion point (up to L_{max} , which is presumed to be large enough so that $\tilde{p}_{[m]}$ well approximates the case where we integrate to ∞ , further discussion on accuracy is provided below). Finally, we take care of points with index $j = 1, \dots, m - 1$ (within the layer) and $j = m + 1, \dots, M - 1$ (beyond the exhaustion). Within the layer for $j = 1, \dots, m - 1$, we integrate the severity to compute the probabilities

$$\tilde{p}_{[j]} = \int_{x_{[j]}}^{x_{[j+1]}} f_X(x) dx,$$

and beyond the exhaustion there is zero probability, and hence, $\tilde{p}_{[j]} = 0$ when $j = m + 1, \dots, M - 1$. Due to numerical rounding, numerical integration errors, and the fact L_{max} is finite, $\sum_{j=0}^{M-1} \tilde{p}_{[j]} \equiv \Phi$ will not generally add up to exactly 1 (to machine precision). We can minimize these errors by enforcing normalization so that we obtain the final probabilities $p_{[j]} = \frac{\tilde{p}_{[j]}}{\Phi}$. With the grid $x_{[0]}, x_{[1]}, \dots, x_{[M-1]}$, and its associated post-catXL layer severity probabilities $(p_{[0]}, p_{[1]}, \dots, p_{[M-1]}) \equiv \mathbf{p}_{AE}$, we are set up to calculate the annual aggregate loss probabilities (subscript AE is used to denote the layer under consideration).

Step 2: For completeness we review the definitions of DFT and IDFT, which are critical to our calculation of the post-catXL aggregate loss probabilities. We review the definitions provided in section 6 of Shevchenko (2010), cast into our particular notation. The DFT is defined as

$$\psi_j = \sum_{l=0}^{M-1} p_{[l]} e^{\frac{2\pi i}{M} lj},$$

where we recall that $j = 0, 1, \dots, M - 1$, and the complex number $i = \sqrt{-1}$. The IDFT, which recovers the original probabilities, is defined by,

$$p_{[j]} = \frac{1}{M} \sum_{l=0}^{M-1} \psi_l e^{-\frac{2\pi i}{M} lj}.$$

Shevchenko (2010) pointed out that when M is a power of 2, the DFT and IDFT can be implemented efficiently with order $M \log_2(M)$ calculations. The efficient implementation is the well-known fast Fourier transform (FFT). A complete treatment of the FFT from an applied perspective is provided in Brigham (1988) (interested readers can explore the references therein for rigorous treatment of Fourier transforms). With the FFT in hand, we move onto to calculate the aggregate loss probabilities for a catXL layer.

Step 3: Consider the catXL layer with attachment A , exhaustion E and (temporarily) ignoring the annual aggregate limit $A_{lim} = (n + 1)(E - A)$ (Section 2.3) so that $A_{lim} \rightarrow \infty$. Our aim is to compute the aggregate probabilities that we stack into a vector $(\tilde{p}_{S,AE,[0]}, \tilde{p}_{S,AE,[1]}, \dots, \tilde{p}_{S,AE,[M-1]}) \equiv \tilde{\mathbf{p}}_{S,AE}$ (on the numerical grid $x_{[0]}, x_{[1]}, \dots, x_{[M-1]}$ defined above, using the subscript S to make clear the annual sum, and the $\sim$ denotes

probabilities before applying a final normalization step). The vector of probabilities $\tilde{\mathbf{p}}_{S,AE}$ is defined by

$$\tilde{\mathbf{p}}_{S,AE} = \sum_{k=0}^{\infty} P_N(k) \mathbf{p}_{S,AE|N=k}$$

where $\mathbf{p}_{S,AE|N=k}$ (defined on the same grid $x_{[j]}$ with $j = 0, \dots, M-1$) are the aggregate layer probabilities given exactly $N = k$ draws from the severity (pictured as years of simulation in Figure 1 with $N = k$ events). The strategy is to take the FFT of $\tilde{\mathbf{p}}_{S,AE}$, then show how we can use the linearity and convolution properties of the FFT to derive an expression where a well-known identity can be used to eliminate the infinite sum under the assumption that $P_N(k)$ is Poisson. The final step is to take the inverse FFT to yield the desired probability vector. Proceeding with the calculation we take the FFT of $\tilde{\mathbf{p}}_{S,AE}$ to obtain,

$$FFT(\tilde{\mathbf{p}}_{S,AE}) = FFT\left(\sum_{k=0}^{\infty} P_N(k) \mathbf{p}_{S,AE|N=k}\right).$$

The linearity property enables us to write,

$$FFT(\tilde{\mathbf{p}}_{S,AE}) = \sum_{k=0}^{\infty} P_N(k) FFT(\mathbf{p}_{S,AE|N=k}).$$

The convolution theorem applied to $FFT(\mathbf{p}_{S,AE|N=k})$ above yields,

$$FFT(\mathbf{p}_{S,AE|N=k}) = (FFT(\mathbf{p}_{AE}))^k.$$

This enables us to use the k th power of the FFT of the layer transformed severity $\mathbf{p}_{AE}$, which we have already computed in Step 1. (This conveniently avoids dealing with the conditional probabilities $\mathbf{p}_{S,AE|N=k}$.) Plugging the result from the convolution theorem into the infinite sum for $FFT(\tilde{\mathbf{p}}_{S,AE})$, we obtain

$$FFT(\tilde{\mathbf{p}}_{S,AE}) = \sum_{k=0}^{\infty} P_N(k) (FFT(\mathbf{p}_{AE}))^k.$$

Under the assumption that $P_N(k)$ is Poisson this becomes

$$\begin{aligned} FFT(\tilde{\mathbf{p}}_{S,AE}) &= \sum_{k=0}^{\infty} \frac{e^{-\lambda} \lambda^k}{k!} (FFT(\mathbf{p}_{AE}))^k \\ &= \sum_{k=0}^{\infty} \frac{e^{-\lambda} (\lambda FFT(\mathbf{p}_{AE}))^k}{k!}, \end{aligned}$$

which as shown in Klugman et al. (2008) has a closed form,

$$FFT(\tilde{\mathbf{p}}_{S,AE}) = e^{-\lambda(1-FFT(\mathbf{p}_{AE}))}.$$

We can, in turn “recover” the desired aggregate probabilities $\tilde{\mathbf{p}}_{S,AE}$ by taking the inverse FFT, which yields the final prescription for computing the post-catXL aggregate loss probabilities (for $A_{lim} \rightarrow \infty$),

$$\tilde{\mathbf{p}}_{S,AE} = Re \left[IFFT \left(e^{-\lambda(1-FFT(\mathbf{p}_{AE}))} \right) \right],$$

where Re indicates that we take the real valued component of each element of the vector given that they will generally have complex components resulting from discrete computation and finite precision. As well, due to numerical errors, the sum of the probabilities in the vector $\tilde{\mathbf{p}}_{S,AE}$ will not generally add up to 1 (to machine precision), we therefore apply an additional normalization step to yield the final aggregate probability vector $\mathbf{p}_{S,AE}$ (which sums to 1 to machine precision). See Brigham (1988) and references therein for a rigorous discussion and proofs of the linear property and convolution theorem used above.

Step 4: The probabilities $\mathbf{p}_{S,AE}$ live on the grid $x_{[0]}, x_{[1]}, \dots, x_{[M-1]}$. To apply the aggregate limit $A_{lim} = (n+1)(E-A)$, we search for the largest index $q \in [1, \dots, M-1]$ such that $x_{[q]} \leq A_{lim}$. The final vector of probabilities is defined as $\mathbf{p}_{S,AE,n}$ (where subscript n indicates the numbers of reinstatements). We first set $\mathbf{p}_{S,AE,n} = \mathbf{p}_{S,AE}$. We then set the q^{th} element of $\mathbf{p}_{S,AE,n}$ to $\sum_{j=q}^{M-1} p_{S,AE,n,[j]}$ (summing the probabilities beyond A_{lim}), and all elements from $j = q+1, \dots, M-1$ are assigned zero probability, namely $p_{S,AE,n,[j]} = 0$ (there is zero probability of loss beyond A_{lim}). This prescription yields the final post-catXL aggregate probabilities for the layer $\mathbf{p}_{S,AE,n}$ characterized by A, E, n , from which we derive reinsurance pricing metrics.

The accuracy of the above calculations depend on the maximum loss L_{max} , the number of grid points M , and the features of the numerical integration schemes. The numerical integration schemes we used are highly accurate (see Section 3.3). Informally, both L_{max} and M should be large enough so that the integral $\int_{L_{max}}^{\infty} f_X(x)$ and grid size $\triangle x$ are “small.” In this study we chose fixed values for L_{max} and number of grid points M that allowed accurate computation of pricing metrics for all catXL layers (characterized by A, E, n) studied in this work. See Section 2.4.3, which discusses our method of measuring accuracy, and Section 3.3 for numerical results that quantify the accuracy of our numerical setup. Finally, our numerical scheme has been set up with issues specifically related to the FFT (Brigham 1988) in mind, including “zero padding,” which is achieved by the assignment of zero probabilities for grid points indexed by $j = m+1, \dots, M-1$, and setting M to a power of 2.