

Modelling catastrophe claims with left-truncated severity distributions^(*)

Anna Chernobai¹, Krzysztof Burnecki², Svetlozar Rachev^{1,3},
Stefan Trück³ and Rafał Weron²

¹ University of California Santa Barbara, CA 93106, USA

² Hugo Steinhaus Center for Stochastic Methods, Institute of
Mathematics and Computer Science, Wrocław University of
Technology, 50-370 Wrocław, Poland

³ Institut für Statistik und Mathematische Wirtschaftstheorie,
Universität Karlsruhe, D-76128 Karlsruhe, Germany

Summary

In this paper, we present a procedure for consistent estimation of the severity and frequency distributions based on incomplete insurance data and demonstrate that ignoring the thresholds leads to a serious underestimation of the ruin probabilities. The event frequency is modelled with a non-homogeneous Poisson process with a sinusoidal intensity rate function. The choice of an adequate loss distribution is conducted via the in-sample goodness-of-fit procedures and forecasting, using classical and robust methodologies.

Keywords: Natural Catastrophe, Property Insurance, Loss Distribution, Truncated Data, Ruin Probability

^(*) We are thankful to S. Stoyanov of the FinAnalytica Inc. for partial computational assistance. E-mail addresses: A.Chernobai: annac@pstat.ucsb.edu, K.Burnecki: Krzysztof.Burnecki@pwr.wroc.pl, S.Rachev: rachev@pstat.ucsb.edu, S.Trück: stefan.trueck@statistik.uni-karlsruhe.de and R.Weron: Rafal.Weron@pwr.wroc.pl.

1 Introduction

Due to increasingly severe catastrophes in the last five years the property insurance industry has paid out over \$125 billion in losses. In 2004 property insured losses resulting from natural catastrophes and man-made disasters, excluding the tragic tsunami of December 26, amounted to \$42 billion, of which 95% was caused by natural disasters and 5% by man-made incidents (SwissRe 2004). These huge billion dollar figures call for very accurate models of catastrophe losses. Even small differences in model parameters can result in underestimation of risk leading to billion dollar losses of the reinsurer. Hence, sound statistical analysis of the catastrophe data is of uttermost importance.

In the property insurance industry the term “catastrophe” denotes a natural or man-made disaster that is unusually severe and that affects many insurers and policyholders. In this paper we analyze losses resulting from natural catastrophic events in the United States. Estimates of such losses are provided by ISO’s (Insurance Services Office Inc.) Property Claim Services (PCS). For each catastrophe, the PCS loss estimate represents anticipated industrywide insurance payments for different property lines of insurance covering, see Burnecki et al. (2000). An event is designated a catastrophe when claims are expected to reach a certain dollar threshold. Initially the threshold was set to \$5 million. However, due to changing economic conditions, in 1997 ISO increased its dollar threshold to \$25 million. In what follows we examine the impact of the presence of left-truncation of the loss data on the resulting risk processes.

The correct estimation of the claims frequency and severity distributions is the key to successfully running an insurance business. A naive and possibly misleading approach for modelling the claim magnitudes would be to fit the *unconditional* distributions. Since the lower quantiles of the actual catastrophe data are truncated from the available data set, ignoring the (non-randomly) missing data would result in biased estimates of the parameters leading to over-stated mean and understated variance estimates, and under-estimated upper quantiles, in general. Furthermore, treating the available frequency data as complete, results in under-estimated intensity of the events (e.g. in compound Poisson processes for aggregated insurance claims). One serious implication of such data misspecification could be wrong (under-estimated) ruin probabilities for the compound risk process. The estimation technique for loss data truncated from below is also useful when dealing with excess-of-loss reinsurance coverage where the data generally exceeds some underlying retention, see Klugman et al. (1998) and Patrik (1981).

The paper is organized as follows. In Sections 2 and 3 we give a brief overview of the insurance risk model and present a methodology of treating the loss data samples with non-randomly missing observations in which the number of missing data points is unknown. Necessary adjustments to the loss and

frequency distributions are discussed. In Section 4 we examine the theoretical aspects of the effects of such adjustment procedures to the severity and frequency distributions from Section 3 on the ruin probabilities. In Section 5 we present an extensive empirical study for the 1990-1999 U.S. natural catastrophe data. In this section we model the incidence of events with a non-homogeneous Poisson process and consider various distributions to fit the claim amounts. We then conduct the goodness-of-fit tests – in-sample and out-of-sample, select most adequate models, and examine the effects of model misspecification on the ruin probabilities. Finally, an additional forecasting methodology based on the robust statistics is proposed. Section 6 concludes and states final remarks.

2 Catastrophe insurance claims model

A typical model for insurance risk, the so-called collective risk model, has two main components: one characterizing the frequency (or incidence) of events and another describing the severity (or size or amount) of loss resulting from the occurrence of an event (Panjer & Willmot 1992). The stochastic nature of both the incidence and severity of claims are fundamental components of a realistic model. Hence, claims form the aggregate claim process $S_t = \sum_{k=1}^{N_t} X_k$, where the claim severities are described by the random sequence $\{X_k\}$ and the number of claims in the interval $(0, t]$ is modelled by a point process N_t , often called the claim arrival process. It is reasonable in many practical situations to consider the point process N_t to be a non-homogeneous Poisson process (NHPP) with a deterministic intensity function $\lambda(t)$. We make such an assumption in our paper.

The risk process $\{R_t\}_{t \geq 0}$ describing the capital of an insurance company is defined as:

$$R_t = u + c(t) - S_t. \quad (1)$$

The non-negative constant u stands for the initial capital of the insurance company. The company sells insurance policies and receives a premium according to $c(t)$. In the non-homogeneous case it is natural to set

$$c(t) = (1 + \theta)\mu \int_0^t \lambda(s)ds, \quad (2)$$

where $\mu = E(X_k)$ and $\theta > 0$ is the relative safety loading which “guarantees” survival of the insurance company.

In examining the nature of the risk associated with a portfolio of business, it is often of interest to assess how the portfolio may be expected to perform over an extended period of time. One approach involves the use of ruin theory, which is concerned with the excess of the income $c(t)$ (with respect

to a portfolio of business) over the outgo, or claims paid, $S(t)$. This quantity, referred to as insurer's surplus, varies in time. Specifically, ruin is said to occur if the insurer's surplus reaches a specified lower bound, e.g. minus the initial capital. One measure of risk is the probability of such an event, clearly reflecting the volatility inherent in the business. In addition, it can serve as a useful tool in long range planning for the use of insurer's funds. The ruin probability in finite time T is given by

$$\psi(u, T) = \mathbb{P} \left(\inf_{0 < t < T} \{R_t\} < 0 \right).$$

In non-life insurance, it is natural to regard T equal to four or five years as reasonable (Grandell 1991). Naturally, the choice of the intensity function and the distribution of claim severities heavily affects the simulated values and, hence, the ruin probability.

In what follows, we denote the distribution of claims by F_γ and its probability density function by f_γ . We assume that the family of distributions of F_γ is sufficiently well behaved so that the parameter (vector) γ can be estimated consistently by maximum likelihood (MLE). To avoid the possibility of negative losses we restrict the distribution to be concentrated on the positive half line. Independence between frequency (N_t) and severity (F_γ) distributions is generally assumed and we assume it also in this paper. Given a sample $\mathbf{x} = (x_1, x_2, \dots, x_n)$ containing all losses which have occurred during a time interval $[T_1, T_2]$, the task of estimating γ can be performed with MLE, while $\lambda(t)$ can be calibrated by fitting a deterministic function to the loss frequency process via, for example, a least squares procedure (see Section 5.2).

In reality not all insurance losses over a certain time interval are recorded accurately. In the framework of catastrophic losses, the losses of magnitudes not exceeding \$5 million (until 1996) or \$25 million (since 1997) are not recorded in the PCS databases. The problem of catastrophe insurance claims data thus lies in the presence of *non-randomly missing data* on the left side of the loss distribution. The question addressed in the subsequent analysis is whether ignoring the missing data has a significant impact on the estimation of the intensity function $\lambda(t)$ and the severity parameter γ , and hence the ruin probability estimates. From the statistical viewpoint, with non-randomly missing data, all estimates would be biased if the missing data is not accounted for. However in practical applications a possible rationale for ignoring the missing data would be as follows: since the major part of catastrophic insurance losses is in excess of the \$5 and even the \$25 million threshold, then losses below it can not have a significant impact on the ruin probabilities, that is largely determined by the *upper* quantiles of the loss distribution.

3 Estimation of loss distributions

In this section we present a procedure for a consistent estimation of the claims size distribution. Given a time interval $[T_1, T_2]$ – the sample window – the collected data which is available for estimating $\lambda(t)$ and γ , is considered *incomplete*: there exists one non-negative pre-specified threshold $H \geq 0$, that defines a partition on $\mathbb{R}_{\geq 0}$: $[0, H]$ and (H, ∞) . If a random outcome of the loss distribution belongs to $[0, H]$ then it does not enter the data sample: neither the frequency nor the severity of losses not exceeding H are recorded (missing data). Realizations in (H, ∞) are fully reported, i.e. both the frequency and the loss amounts are specified. Hence, we are dealing with *truncated data*.

Let the observed sample in $[T_1, T_2]$ be of the form $\mathbf{x}^o = (x_1, x_2, \dots, x_n) \in (H, \infty)$, where n denotes the number of observations in (H, ∞) and x_1, x_2, \dots, x_n the values of the observations. The corresponding sample space is denoted as \mathcal{X} . Given that the total number of observations in the complete sample is unknown, the *joint* density on \mathcal{X} (with respect to the product of counting and Lebesgue measures) can be given by the following expression:

$$g_{\lambda(\Delta T), \gamma}(\mathbf{x}^o) = \frac{1}{n!} \left(\int_{T_1}^{T_2} \lambda^o(t) dt \right)^n e^{-\int_{T_1}^{T_2} \lambda^o(t) dt} \prod_{k=1}^n \frac{f_\gamma(x_k)}{1 - F_\gamma(H)} \quad (3)$$

where $F_\gamma(H)$ denotes the probability for a random realization to fall into the interval $[0, H]$. Other model specifications, such as a renewal process or a Cox process are also possible, but not considered in this study. By representation (3), the Poisson process $N_{\Delta T}^o$ that counts only the losses of magnitudes greater than H is interpreted as a *thinning* of the original (complete) process $N_{\Delta T}^c$ governed by $\lambda^c(t)$, with a new intensity (rate) function $\lambda^o(t) = \{1 - F_\gamma(H)\} \lambda^c(t)$. The time frame is $\Delta T = T_2 - T_1$. The superscripts “o” and “c” refer to “observed” (the incomplete data set), and “complete” or “conditional” (the complete data set), respectively. The maximization of the corresponding log-likelihood function of the compound process is done only with respect to γ :

$$\hat{\gamma}_{\text{MLE}}^c = \arg \max_{\gamma} \log \left(\prod_{k=1}^n \frac{f_\gamma(x_k)}{1 - F_\gamma(H)} \right). \quad (4)$$

In this study, the estimation of the intensity function does not require MLE. It is obtained directly by fitting a deterministic function to the aggregated numbers of events per unit interval over the time frame of interest. The true intensity of the complete data set is obtained by $\hat{\lambda}^c(t) = \hat{\lambda}^o(t) / \{1 - \hat{F}_{\gamma^c}(H)\}$. The unknown value $F_{\gamma^c}(H)$ needs to be estimated from (4). Such an adjustment allows to “add back” the fraction of the missing data. Certainly, $F_{\gamma^c}(H)$ would vary for various distributions. The crucial assumption of such an amendment is that the underlying distribution of the loss magnitudes is indeed the true distribution.

In the cases where no closed-form expression for the MLE estimate of γ in expression (4) is available, as applies to most of the distributions considered in the empirical part of this paper, we have to solve for it numerically via direct numerical optimization. However, in the cases where a closed-form expression for both the unconditional MLE estimate of γ as well as the expectations $E_\gamma(\log f_\gamma(\mathbf{x}^o))$ and $E_\gamma(\log f_\gamma(\mathbf{x}^c))$ for given value of γ are available (e.g. for distributions such as Gaussian, Lognormal and Exponential), it might be optimal to apply the *Expectation-Maximization* algorithm (*EM*-algorithm), see Chernobai et al. (2005a) and Chernobai et al. (2005c). Surprisingly, for the distributions under study (Lognormal and Exponential) no advantage over direct integration in either speed or accuracy was observed. Hence, only the latter method was used for calibration throughout the paper.

4 Impact of density misspecification on the ruin probability

We examine two possible approaches insurance companies may undertake for the parameter estimation and subsequently the ruin probability determination. The *first*, and correct, approach involves finding the estimates $\hat{\lambda}^c(t)$ and $\hat{\gamma}_{\text{MLE}}^c$ for the unknown function $\lambda^c(t)$ and parameter γ^c with the direct numerical optimization (or the *EM*-algorithm) and determine the ruin probability. For simulations, losses can be drawn from the distribution with the complete-data estimated parameters, and use the complete-data frequency parameter. The *second* approach would be to use the observed frequency estimate $\hat{\lambda}^o(t)$ and fit the unconditional distribution to the truncated data. Evidence from the literature indicates that this naive approach has been widely used in practice.

For reinsurance, when only large losses matter, a *third* approach could be relevant. It involves estimating the frequency function $\lambda(t)$ by the observed frequency $\hat{\lambda}^o(t)$, and estimating the complete-data conditional distribution. For the calculation of the ruin probabilities with the Monte Carlo method (which is most commonly used), losses above the threshold should be simulated from the conditional distribution, and the observed frequency should be used.

In our subsequent analysis and estimations of the ruin probabilities, we compare the first two approaches. The second oversimplified and misspecified approach will lead to biased estimates for the rate function $\lambda(t)$ and the parameter γ of the loss distribution. The bias can be expressed analytically. For a Lognormal $\mathcal{LN}(\mu, \sigma)$ loss distribution, for example, the bias is expressed as follows:

$$\begin{aligned}
E\widehat{\lambda}^\sigma(t) &= \lambda^c(t) \cdot (1 - F_{\gamma^c}(H)) = \lambda^c(t) - \lambda^c(t) \cdot \Phi(a), \\
E\widehat{\mu}^\sigma &= E\left(\frac{1}{n} \sum_{k=1}^n \log X_k \mid X_k > H\right) = \mu^c + \sigma^c \cdot \frac{\varphi(a)}{1 - \Phi(a)}, \\
E(\widehat{\sigma}^\sigma)^2 &= E\left(\frac{1}{n} \sum_{k=1}^n \log^2 X_k - (\widehat{\mu}^\sigma)^2 \mid X_k > H\right) \\
&= (\sigma^c)^2 + (\sigma^c)^2 \cdot \left[a \cdot \frac{\varphi(a)}{1 - \Phi(a)} - \left\{ \frac{\varphi(a)}{1 - \Phi(a)} \right\}^2 \right],
\end{aligned}$$

where $a = (\log H - \mu^c)/\sigma^c$, φ and Φ denote the density and d.f. of the standard Normal law and μ^c and σ^c are the true (complete data) Lognormal parameters. In the first expression above, the oversimplified approach leads to a misspecification bias that will be less than 0 always. Since the bias of the location parameter $\widehat{\mu}^\sigma$ is always positive, then the observed $\widehat{\mu}^\sigma$ is always overstated. For practical purposes, since $\log H < \widehat{\mu}^c$ (the threshold level is relatively low), then the bias of the scale parameter is negative, and so the true $(\sigma^c)^2$ is underestimated under the unconditional fit. The effect on the ruin probability would depend on the values of H , $\widehat{\mu}^c$ and $(\widehat{\sigma}^c)^2$.

5 Empirical analysis of catastrophe data

We take for our study the PCS data covering losses resulting from natural catastrophe events in USA that occurred between 1990 and 1999. The data were adjusted using the Consumer Price Index provided by the U.S. Department of Labor, see Burnecki et al. (2005). For the calibration and in-sample validation we consider all claim amounts exceeding \$25 million between 1990 and 1996. For the forecasting part of the paper we consider the losses over a three year period 1997-1999. The goal of the subsequent empirical study is three-fold: we aim at 1^o examining the effect of ignoring the threshold (missing data) on distributional parameters, 2^o obtain the best model via the goodness-of-fit tests, and 3^o examine the effect of the data misspecification (from part 1^o) on ruin probability under the threshold \$25 million.

5.1 Loss distributions

The following severity distributions (and following parameterizations) are considered in the study:

- Exponential: $\mathcal{Exp}(\beta) = \beta e^{-\beta x}$; $\beta > 0$.

	$\gamma, F(H)$	Unconditional	Conditional
<i>Exp</i>	β	$2.7912 \cdot 10^{-9}$	$3.0006 \cdot 10^{-9}$
	$F(H)$	6.74%	7.23%
<i>LN</i>	μ, σ	18.5660, 1.1230	17.3570, 1.7643
	$F(H)$	8.63%	42.75%
<i>Gam</i>	α, β	$0.5531, 1.5437 \cdot 10^{-9}$	$2.155 \cdot 10^{-8}, 0.8215 \cdot 10^{-9}$
	$F(H)$	18.34%	$\approx 100\%$
<i>Weib</i>	β, τ	$2.8091 \cdot 10^{-6}, 0.6663$	0.0187, 0.2656
	$F(H)$	21.23%	82.12%
<i>GPD</i>	ξ, β	$0.5300, 1.2533 \cdot 10^8$	0.8090, $0.5340 \cdot 10^8$
	$F(H)$	17.27%	32.77%
<i>Burr</i>	α, β	0.1816, $3.0419 \cdot 10^{35}$	0.1748, $1.4720 \cdot 10^{35}$
	τ	4.6867	4.6732
	$F(H)$	2.58%	3.87%
log <i>S</i>	α, β	1.4265, 1	1.9165, 1
	σ, μ	0.5689, 18.8584	0.9706, 17.9733
	$F(H)$	0.005%	23.27%

Table 1: Estimated parameters and $F(H)$ of the fitted distributions to the PCS data.

- Lognormal: $\mathcal{LN}(\mu, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2x}} e^{-(\log x - \mu)^2/2\sigma^2}$; $\mu, \sigma > 0$.
- Gamma: $\mathcal{Gam}(\alpha, \beta) = \frac{\beta^\alpha x^{\alpha-1}}{\Gamma(\alpha)} e^{-\beta x}$; $\alpha, \beta > 0$.
- Weibull: $\mathcal{Weib}(\beta, \tau) = \tau \beta x^{\tau-1} e^{-\beta x^\tau}$; $\beta, \tau > 0$.
- Burr: $\mathcal{Burr}(\alpha, \beta, \tau) = \tau \alpha \beta^\alpha x^{\tau-1} (\beta + x^\tau)^{-(\alpha+1)}$; $\alpha, \beta, \tau > 0$.
- Generalized Pareto: $\mathcal{GPD}(\xi, \beta) = \beta^{-1} (1 + \xi x \beta^{-1})^{-(1+\frac{1}{\xi})}$; $\beta > 0$.
- log- α Stable: no closed-form density for $\log \mathcal{S}_\alpha(\beta, \sigma, \mu)$; $\alpha \in (0, 2)$, $\beta \in [-1, 1]$, $\sigma > 0$, $\mu > 0$.

In Table 1 we demonstrate the change in the parameter values when the conditional (truncated) distribution is fitted instead of the unconditional. The location parameters are lower and the scale parameters are higher under the correct data specification. In addition, the shape parameter which is present in the relevant distributions is lower (except for log- α Stable; note that the shape parameter of the GPD distribution is ξ^{-1}) under the conditional fit, indicating a heavier tailed true distribution for the claim size data.

The estimated fraction of the missing data $F(H)$ is larger under the conditional fit, as expected. This could be considered as evidence for the fact that




Figure 1: *Left panel:* Quarterly number of losses for PCS data. *Right panel:* Periodogram of PCS quarterly number of losses. A distinct peak is visible at frequency $\omega = 0.25$ implying a period of $1/\omega = 4$ quarters, i.e. one year.

conditional estimation accounts for true “information loss” while the unconditional fit underestimates the fraction of missing data. We point out that the estimates of $F(H)$ are explicitly dependent on the choice of the distribution (and certainly the threshold H). Further, we find that for the light-tailed Exponential distribution the estimated fraction of data below threshold H is almost negligible while for more heavy-tailed distributions like Weibull or log- α Stable the estimated fraction is significantly higher for the conditional case. The results are consistent with the findings in Chernobai et al. (2005c) with operational loss data.

For the purpose of our subsequent analysis, we decided to exclude the Gamma distribution for the following reasons. The Gamma distribution produced the true “information loss” nearly equal to 100%, which means that if Gamma is the true distribution for the data, then nearly all data is considered missing, which is unfeasible. The true estimate of the intensity rate would blow up to infinity.

5.2 Intensity function

We model the frequency of the losses with a non-homogeneous Poisson process (NHPP), in which the intensity of the counting process varies with time. The time series of the quarterly number of losses does not exhibit any trends but an annual seasonality can be very well observed using the periodogram, see Figure 1. This suggests that calibrating a NHPP with a sinusoidal rate function would give a good model. Following Burnecki & Weron (2005) we estimate the parameters by fitting the cumulative intensity function, i.e. the mean value function $E(N_t)$, to the accumulated quarterly number of

PCS losses. The least squares estimation is used to calibrate $\lambda(t) = a + b \cdot 2\pi \cdot \sin\{2\pi(t - c)\}$ yielding parameters $a = 30.8750$, $b = 1.6840$ and $c = 0.3396$. This form of $\lambda(t)$ gives a reasonably good fit measured by the mean square error $MSE = 18.9100$, and the mean absolute error $MAE = 3.8385$. It is notable, that if, instead, a homogeneous Poisson process (HPP) with a constant intensity was considered for the quarterly number of losses, then the respective error estimates would yield $MSE = 115.5730$ and $MAE = 10.1308$. The latter values are based on the Poisson parameter estimated to be $\lambda = 33.0509$ for the data set, obtained by fitting the Exponential distribution to the respective inter-arrival times, in years. Alternatively, the mean annual number of losses, can be obtained by multiplying the quarterly number of points by four and averaging, yielding 31.7143. These result in $MSE = 38.2479$ and $MAE = 5.3878$. In either case, significantly higher values for MSE and MAE under HPP, lead to the conclusion that NHPP with the intensity rate of a functional form described above, results in a reasonably superior calibration of the loss arrival processes.

To adjust for the missing data, we adjust the parameters a , b and c , according to the procedure described in Section 3. Using the estimates of the missing data, $F(H)$, from Table 1, straightforward calculations result in the conclusion that the true frequency of the loss events is highly underestimated.

5.3 Backtesting

In this section of our empirical study we aim at determining which of the considered distributions is most appropriate to use for the catastrophe loss data. The ultimate choice of a model can be determined via backtesting. We conduct two types of test: in-sample and out-of-sample goodness-of-fit tests.

5.3.1 In-Sample Goodness-of-Fit Tests

We test a composite null hypothesis that the empirical distribution function (EDF) $F_n(x)$ belongs to an entire family of hypothesized truncated distributions. We consider four kinds of the EDF-based tests: Kolmogorov-Smirnov (D), Kuiper (V), Anderson-Darling (A^2) and Cramér-von Mises (W^2).¹ For numerical implementation of the above formulas consult e.g. Burnecki et al. (2005). For the truncated case, to account for the missing data the corresponding statistics were computed and the simulations were carried out according to the procedure described in Chernobai et al. (2005b). Since the limiting distributions of the test statistics are not parameter-free, the p -values and the critical values were obtained from Monte Carlo simulations, using 1,000 samples.

¹For D and V statistics we use the scaling factor \sqrt{n} , and for A^2 and W^2 we use n .

Unconditional				Conditional			
D	V	A^2	W^2	D	V	A^2	W^2
<i>Exp</i>							
5.1234	6.1868	48.9659	10.1743	5.5543	5.9282	72.2643	13.1717
[<0.005]	[<0.005]	[<0.005]	[<0.005]	[<0.005]	[<0.005]	[<0.005]	[<0.005]
<i>LN</i>							
1.5564	2.9710	4.4646	0.7139	0.6854	1.1833	0.7044	0.0912
[<0.005]	[<0.005]	[<0.005]	[<0.005]	[0.256]	[0.315]	[0.080]	[0.133]
<i>Weib</i>							
3.2755	5.5430	14.2197	2.3859	0.8180	1.5438	1.3975	0.1965
[<0.005]	[<0.005]	[<0.005]	[<0.005]	[0.107]	[0.053]	[0.006]	[0.008]
<i>GPD</i>							
2.7084	3.9240	7.6731	1.1013	0.4841	0.8671	0.3528	0.0390
[<0.005]	[<0.005]	[<0.005]	[<0.005]	[0.795]	[0.847]	[0.487]	[0.666]
<i>Burr</i>							
0.8876	1.3149	0.8983	0.1469	0.4604	0.8668	0.2772	0.0342
[0.054]	[0.106]	[0.014]	[0.022]	[0.822]	[0.793]	[0.560]	[0.659]
<i>log S</i>							
4.9863	4.9953	94.6918	11.3252	0.8961	1.2111	0.8062	0.1535
[0.596]	[0.616]	[0.483]	[0.575]	[0.456]	[0.470]	[0.484]	[0.444]

Table 2: Results of the in-sample goodness-of-fit tests. p -values were obtained via 1,000 Monte Carlo simulations, and are given in the square brackets.

It is evident from Table 2 that under the data misspecification none of the considered distributions (except for Burr to some extent) appear to provide a good fit, as indicated by the near-zero p -values. Note that the log- α Stable distribution's fit is poor under the unconditional fit, because, despite the high p -values, the observed test statistic values are high. With the truncated calibration, Burr and Pareto distributions show a very good fit around both the median and the tails as indicated by the low statistic values and high p -values. The fit of the light-tailed Exponential distribution is again poor, on the basis of which we exclude it from further analysis. The Weibull distribution provides moderately low values of the test statistics, however in terms of p -values the results are rather poor. As for log- α Stable and Lognormal, relatively low observed statistic values and often high p -values suggest an acceptable fit. In general, the data seems to follow a very heavy-tailed law.

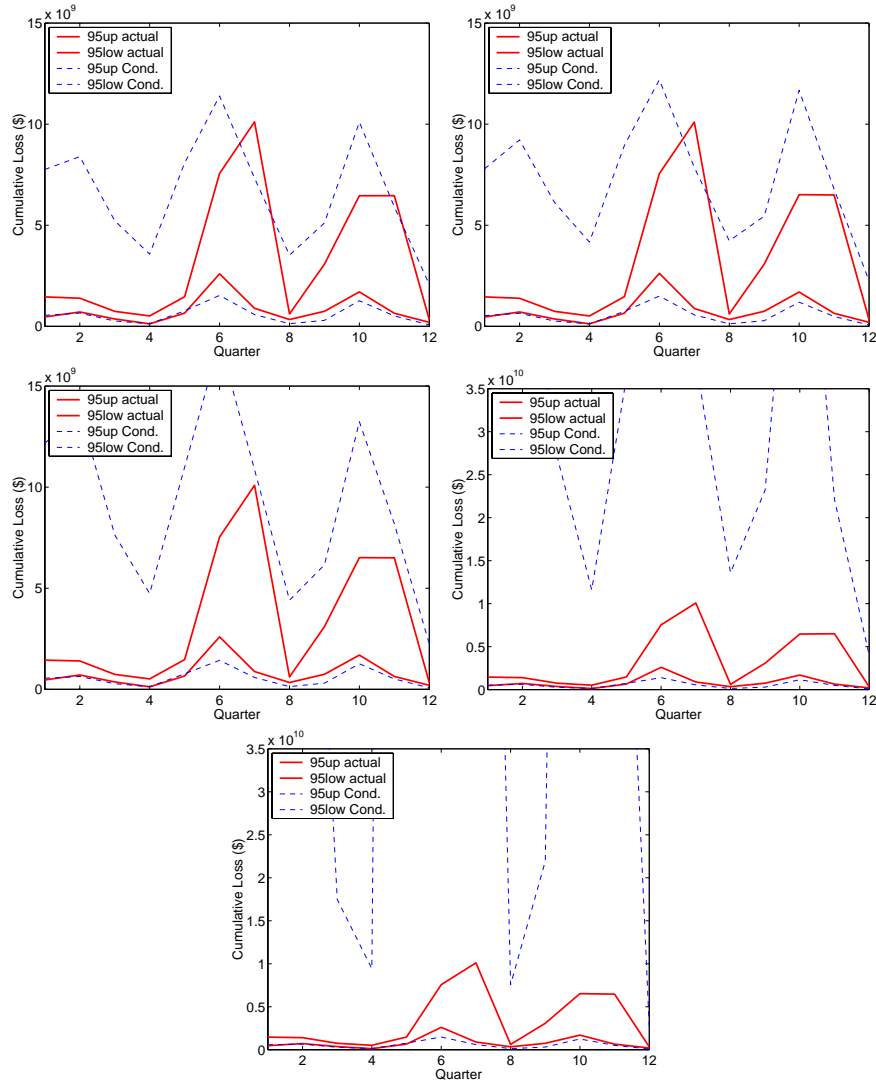


Figure 2: 95% confidence intervals for conditional (dashed blue) forecasted cumulative losses and corresponding 95% bootstrap confidence intervals for realized losses (solid red): Lognormal (*top left*), Weibull (*top right*), Generalized Pareto (*mid left*), Burr (*mid right*), and log- α Stable (*bottom*).

5.3.2 Out-of-Sample Goodness-of-Fit Tests

Examining how well or how badly various considered models predict the true future losses, is, we believe, the key to determining which of the loss distributions is the best to be used for practical purposes. For our loss distribution estimations in Section 5.1 we used the data set from 1990 to 1996. We now analyze our models' predicting power regarding the data between 1997 and 1999. For the out-of-sample backtesting, we assume that our model has a one-step ahead predicting power, with one step equal to one quarter. The window length of the sample used for calibration is taken to be six years. We start with the data from the first quarter of 1990 until the fourth quarter of 1996, in order to conduct the forecasting about the first quarter of 1997. First, we estimate the unknown parameters of truncated distributions. Next, to obtain the distribution of the quarterly aggregated losses we repeat the following a large number (10,000) of times: use the estimated parameters to simulate N losses exceeding the \$25 million threshold, where N is the actual number of losses in the quarter that we perform forecasting on, and aggregate them. At each forecasting step (twelve steps total) we shift the window by one quarter forward and repeat the above procedure. Note, that in this way we test the model for the severity distribution but not for the entire risk process or the intensity itself.

We break the analysis of the forecasting results into two main parts. First, we look at how good are the different model assumptions in predicting the distribution of accumulated losses around the center, or the main body, of the actual loss distribution. Then we examine how well they predict the true aggregated losses around the tails. At each step, we estimate the 2.5th and 97.5th percentiles² of the distribution of the sum of realized losses for the corresponding quarter. This is conducted using a non-parametric bootstrapping technique. We use the previous 24 quarters of data (six years) to estimate the parameters of the conditional distributions and determine, based on the parameter estimates. Then we determine the 95% confidence intervals (2.5th and 97.5th percentiles) for the quarterly cumulative losses, by simulating 10,000 samples (of losses that exceed the threshold of \$25 million) of the size equal to the number of realized losses in the forecasted quarter. We compare the intervals with the bootstrapped intervals based on the actual data.

Figure 2 portrays the forecasting ability of the considered five models around the tails of the actual loss distribution. It is notable, that for 1997 (quarters 1-4) all models tend to overestimate the cumulative losses, to various extents. For the remaining part of the test period, Lognormal and Weibull assumptions generally result in quite accurate forecasted 95% confidence interval estimates. As for Pareto, the forecasted 95% confidence intervals'

²Results for the 25th, 50th and 75th percentiles are available in the original Working Paper, see www.statistik.uni-karlsruhe.de/technical_reports/catastrophe.pdf.

	$\overline{\text{MSE}}$	$\overline{\text{MAE}}$
\mathcal{LN}	$6.8803 \cdot 10^{18}$	$1.2969 \cdot 10^9$
$Weib$	$5.7869 \cdot 10^{18}$	$1.3745 \cdot 10^9$
\mathcal{GPD}	$3.2915 \cdot 10^{20}$	$1.9123 \cdot 10^9$
$Burr$	$1.2562 \cdot 10^{24}$	$1.4986 \cdot 10^{10}$
$\log \mathcal{S}$	$1.8654 \cdot 10^{89}$	$1.2468 \cdot 10^{42}$

Table 3: Estimates for the average MSE and MAE of forecasted aggregated losses above \$25 million, obtained via 10,000 Monte Carlo simulations.

upper bounds are much higher than the true bootstrapped bounds. The effect is even more significant under the Burr and log- α Stable laws. Hence, the ruin probabilities may be unnecessarily high under these assumptions. Overall, the figures indicate that Lognormal and Weibull assumptions possess a reasonably good predicting power.

Finally, we use the 10,000 Monte Carlo samples to compute the MSE and MAE for the estimates of the sum of losses, with respect to the true realized losses of each quarter. We present the estimates of the average MSE and MAE over the 3-year forecasting time interval in Table 3. In terms of the average MSE and MAE, the Lognormal, Weibull and Pareto distributional assumptions result in relatively accurate predictions regarding the aggregated losses one quarter ahead, for the considered three year period. In particular, the Weibull model is most optimal in terms of the MSE and the Lognormal model appears most optimal in terms of the MAE. The errors under the Burr and log- α Stable assumptions are considerably higher.

Summarizing, based on the in-sample goodness of fit tests, the observed test statistic values are the lowest for the GPD and Burr laws, indicating their superior in-sample fit. However, the forecasting analysis rather supports the use of the less heavy-tailed Lognormal and Weibull model, as is indicated by the smallest error estimates. Since Lognormal also provided an acceptable in-sample fit it could be considered as the overall most appropriate distribution.

It is notable that, although Pareto, Burr and log- α Stable distributions showed acceptable in-sample fit, they highly overpredict the actual losses in the forecasting 3-year period. This is due to the following: in the data sample used for calibration, there were few losses whose magnitude was of significantly higher order than the rest of the data. Such losses can be categorized as “low frequency/high severity” losses. The two largest losses and their magnitudes (after the inflation adjustment) were Hurricane Andrew that hit Florida and Louisiana on August 24-26, 1992 (appr. \$18.5 billion losses) and Northridge

	Unconditional		Conditional		Increase (# times)	
	5-year	10-year	5-year	10-year	5-year	10-year
\mathcal{LN}	0.00545 [0.00122]	0.00669 [0.00092]	0.10443 [0.00197]	0.13137 [0.00205]	19	20
$Weib$	0.00754 [0.00799]	0.00874 [0.00056]	0.10785 [0.00317]	0.13077 [0.00142]	14	15
GPD	0.07938 [0.00238]	0.10376 [0.00197]	0.15997 [0.00418]	0.20434 [0.00428]	2	2

Table 4: 5-year and 10-year ruin probability estimates, standard errors (in square brackets) and the degree of increase for $u = \$10$ billion and $\theta = 0.3$.

earthquake that struck in California on January 17, 1994 (appr. \$14.4 billion losses). The latter loss entered all twelve shifted sample windows used for forecasting the quarterly losses, while the former entered the first seven. As a consequence very heavy-tailed distributions yielded good in-sample fits. However, since losses of a comparable magnitude were not present in the 1997-1999 sample, this resulted in an overestimation of the confidence intervals' upper bounds. We will return to this issue in Section 5.5.

5.4 Ruin Probability

We consider a hypothetical scenario where the insurance company insures losses resulting from catastrophic events in the United States. The company's initial capital is assumed to be $u = \$10$ billion and the relative safety loading used is $\theta = 30\%$. We choose different models of the risk process whose application is most justified by the statistical results described above. We decided to exclude both the Burr and the log- α Stable distributions as they produce an infinite first moment, making them inapplicable for the purpose of estimating the ruin probability (recall the premium formula (2)). This leaves us with three models: Lognormal, Weibull and GPD (Pareto).

In this paper, the ruin probability is approximated by means of Monte Carlo simulations. We generated $10 \times 10,000$ simulations that were also used to calculate the standard errors. Table 4 demonstrates the effects on the ruin probability due to the data (mis)specification. Based on the results with the considered three distributions, the ruin probability tends to be significantly underestimated when the naive approach is used (i.e. when unconditional distribution is fitted to the loss magnitudes, and the frequency rate function estimates are left unadjusted) instead of the correct approach, in which both the frequency and severity account for the missing data. Note that for the Lognormal and Weibull distribution, the ruin probability increased about 20, respectively 15, times when the correct estimation approach is employed while for the more heavy-tailed GPD distribution the probability only doubled.

5.5 Robust Estimation Approach

In recent years outlier-resistant or so-called robust estimates of parameters are becoming more wide-spread in risk management. Such models – called robust (statistics) models – were introduced by P.J.Huber in 1981 and applied to robust regression analysis, more recent references on robust statistics methods include Huber (2004), Rousseeuw & Leroy (2003), Martin & Simin (2003), Knez & Ready (1997) and Hampel et al. (1986). Robust models treat extreme data points as outliers (or some standard procedure is used to detect outliers in the data) which distort the main flow of the loss process. Under the robust approach, the focus is on modelling the major bulk of the data that is driving the entire process. Robust models help protect against the outlier bias in parameter estimates and provide with a better fit of the loss distributions to the data than under the classical model. Moreover, outliers in the original data can seriously drive future forecasts in an unwanted (such as worst-case scenario) direction, which is avoided by the robust approach models.

Following the idea of robust statistics, for the forecasting purposes we offer a second methodology that involves determining outliers and trimming the top 1-5% of the data. This data adjustment results in a more robust outlook regarding a general future scenario. Excluding the outliers in the original loss data noticeably improves the forecasting power of considered loss distributions, and can be used for forecasting of the generic (most likely) scenario of future losses within reasonable boundaries. The resulting ruin probabilities will be more optimistic than otherwise predicted by the classical model. We emphasize, however, that we are not recommending the use of only one of the two approaches – classical or robust – instead of the other. Rather, in the presence of outliers, we encourage the use of both models for the analysis, and use the robust model as the complement to the classical. We wish to stress that the reason for trimming the top several percent of the data is independent from the consideration that, since the highest claims are transferred to the reinsurers, they need to be taken out from the insurers' database. This transfer of risk from the insurers to the reinsurers is not treated in this paper.

Here we exclude the top 1% of the data, which corresponds to the two highest losses described earlier as outliers. We reproduce Tables 2, 3, 4 now for the trimmed data. They are presented respectively in Tables 5, 6, 7.³ We would like to briefly comment on the results. Table 5 provides in-sample goodness of fit results. Under the conditional approach, even after excluding the two highest points, the Pareto model results in the lowest distance statistics and highest p -values. Also for the heavy-tailed log- α Stable p -values are quite

³Parameter estimates for the robust estimation approach and Figures for out-of-sample bootstrapped confidence intervals are available in the original Working Paper that can be downloaded from www.statistik.uni-karlsruhe.de/technical_reports/catastrophe.pdf.

Unconditional				Conditional			
D	V	A^2	W^2	D	V	A^2	W^2
<i>Exp</i>							
2.8508	4.5958	13.5090	2.6223	3.3324	4.0820	23.2179	4.3263
[<0.005]	[<0.005]	[<0.005]	[<0.005]	[<0.005]	[<0.005]	[<0.005]	[<0.005]
<i>LN</i>							
1.3512	2.5616	3.7007	0.5581	0.5693	0.9528	0.4288	0.0502
[<0.005]	[<0.005]	[<0.005]	[<0.005]	[0.545]	[0.685]	[0.334]	[0.460]
<i>Weib</i>							
2.4880	4.5218	9.6317	1.5862	0.6615	1.1256	0.7045	0.0906
[<0.005]	[<0.005]	[<0.005]	[<0.005]	[0.334]	[0.433]	[0.102]	[0.145]
<i>GPD</i>							
2.6100	3.8319	7.0758	0.9959	0.4696	0.9133	0.3363	0.0331
[<0.005]	[<0.005]	[<0.005]	[<0.005]	[0.831]	[0.766]	[0.529]	[0.751]
<i>Burr</i>							
0.8914	1.3929	1.0414	0.1600	0.4799	0.9472	0.3972	0.0466
[0.043]	[0.058]	[0.005]	[0.010]	[0.757]	[0.607]	[0.274]	[0.399]
<i>log S</i>							
4.2920	4.3149	65.8189	8.5308	1.3800	1.4265	4.4837	0.7335
[0.587]	[0.603]	[0.459]	[0.549]	[0.730]	[0.753]	[0.699]	[0.711]

Table 5: Results of the in-sample goodness-of-fit tests for the *trimmed* data. p -values in square brackets were obtained via 1,000 Monte Carlo simulations.

high, followed by Burr and Lognormal, which contrary to the former also give good results for the test statistics. The fit of the Weibull distribution is not so good while the results for the Exponential remain very poor for the trimmed data. Another main observation is that the forecasts for all distributions show clearly lower MSEs and MAEs for all distributions (Table 6). While the Weibull distribution provides the best forecasts in terms of the MSE and MAE, the log- α Stable and Lognormal give only slightly worse results. Like in the case of the original data the forecasts of the Burr distribution are very conservative and give the highest MSE and MAE error. As regards the finite time ruin probabilities, the figures are much lower than under the classical model with untrimmed data, even in the “conditional” case (we note that the gap between the figures under unconditional and conditional models is remarkably big), because the robust methods put more probability to the medium and small size losses and thus increase the probabilities of ruin resulting from them (Table 7).

	$\overline{\text{MSE}}$	$\overline{\text{MAE}}$
\mathcal{LN}	$2.6103 \cdot 10^{18}$	$1.0635 \cdot 10^9$
$Weib$	$2.0225 \cdot 10^{18}$	$1.0135 \cdot 10^9$
\mathcal{GPD}	$2.3701 \cdot 10^{19}$	$1.2851 \cdot 10^9$
$Burr$	$5.9534 \cdot 10^{23}$	$1.0456 \cdot 10^{10}$
$\log \mathcal{S}$	$2.6091 \cdot 10^{18}$	$1.0701 \cdot 10^9$

Table 6: Average forecast errors for aggregated losses above \$25 million, obtained via 10,000 Monte Carlo simulations, for the *trimmed* data.

Overall, evidence confirms the classical model result that the Lognormal assumption can be considered as plausible providing both an acceptable in-sample and good out-of-sample results. For the trimmed data, however also the log- α Stable provides good results for both criteria. The Pareto model provides the best in-sample fit but higher forecast errors than Lognormal and log- α Stable while Weibull is convincing in terms of forecast MSE and MAE but doesn't provide good in-sample results.

6 Conclusions

This paper focused on analyzing the effects of data misspecification, under which the loss data available from a loss database is truncated from below at a predetermined threshold level. Such thresholds are often ignored in practice. We examined the consequences on the estimates of the ruin probabilities and the choice of claim size distribution for the catastrophic claims model. The theoretical study proposed a practical solution to the problem and suggested that using truncated (conditional) distributions instead of regular (unconditional) distributions provides for more accurate distributional parameters. In-sample goodness-of-fit tests confirmed that fitting a wrong distribution results in an unacceptable fit and the rejection of the null hypothesis that the data is drawn from a hypothesized family of distributions. Fitting truncated distributions and using a correct testing procedure, considerably improved the goodness-of-fit. The truncated Lognormal distribution showed a good fit – both in-sample and out-of-sample (forecasting). Pareto, Burr and log- α Stable showed a good in-sample fit but a poor forecasting power. The Weibull distribution provided good forecasting results but the in-sample fit of the distribution was rather poor. Overall, Lognormal was concluded to be superior on the basis of the in-sample and out-of-sample goodness of fit testing procedures.

	Unconditional		Conditional		Increase (# times)	
	5-year	10-year	5-year	10-year	5-year	10-year
\mathcal{LN}	0.00077 [0.00035]	0.00103 [0.00092]	0.03040 [0.00223]	0.03719 [0.00205]	40	36
$Weib$	0.00004 [0.00007]	0.00006 [0.00010]	0.00702 [0.00082]	0.00835 [0.00087]	176	139
GPD	0.01487 [0.00095]	0.01911 [0.00132]	0.10434 [0.00281]	0.14079 [0.00337]	7	7

Table 7: 5-year and 10-year ruin probability estimates, standard errors (in square brackets) and the degree of increase for $u = \$10$ billion and $\theta = 0.3$, for the *trimmed* data.

A methodology for performing necessary adjustments to the frequency parameter of the loss events was also proposed. It was demonstrated that given the estimated fraction of missing data, the true total number of claims is significantly underestimated, if the missing data is left unaccounted for. The paper then argued and provided empirical evidence that the finite-time ruin probabilities are always seriously underestimated under the model misspecification, i.e. when unconditional distributions are wrongly used. For the several considered distributions, the true ruin probability is up to 20 times higher for 5 and 10 years, than what one would obtain using a wrong model. This has a variety of serious implications for the insurance purposes, and can be of high interest for both insurers and reinsurers. As a counter measure, for a more solid protection against the risk of ruin, they might need to consider increasing the safety loading parameter θ .

This paper finally advocates the use of the “robust” approach, as a complement to the classical approach, under which top 1% of the loss data was treated as outliers and was excluded from the database, and all model parameters were reestimated. Such models are more pertinent to the mainstream loss events and possess an advantage of a more credible forecasting capacity. The use of the robust methodology improved remarkably the forecasts. It also improved the fit of the log- α Stable distribution (making it roughly the second-best candidate), and confirmed the good results of the Lognormal law.

References

- Burnecki, K., Härdle, W. & Weron, R. (2004), An introduction to simulation of risk processes, *in* J. Teugels & B. Sundt, eds, ‘Encyclopedia of Actuarial Science’, Wiley, Chichester, pp. 1564–1570.
- Burnecki, K., Kukla, G. & Weron, R. (2000), ‘Property insurance loss distributions’, *Physica A* **287**, 269–278.

- Burnecki, K., Misiolek, A. & Weron, R. (2005), Loss distributions, *in* P. Cizek, W. Härdle & R. Weron, eds, ‘Statistical Tools for Finance and Insurance’, Springer, Berlin, pp. 289–317.
- Burnecki, K. & Weron, R. (2005), Modeling of the risk process, *in* P. Cizek, W. Härdle & R. Weron, eds, ‘Statistical Tools for Finance and Insurance’, Springer, Berlin, pp. 319–339.
- Chernobai, A., Menn, C., Trück, S. & Rachev, S. (2005a), ‘A note on the estimation of the frequency and severity distribution of operational losses’, *Mathematical Scientist* **30**(2).
- Chernobai, A., Rachev, S. & Fabozzi, F. (2005b), Composite goodness-of-fit tests for left-truncated loss samples, Technical report, University of California Santa Barbara.
- Chernobai, A., Trück, S., Menn, C. & Rachev, S. (2005c), Estimation of operational Value-at-Risk in the presence of minimum collection thresholds, Technical report, University of California Santa Barbara.
- Embrechts, P., Klüppelberg, C. & Mikosch, T. (1997), *Modeling Extremal Events for Insurance and Finance*, Springer-Verlag, Berlin.
- Grandell, J. (1991), *Aspects of Risk Theory*, Springer-Verlag, New York.
- Hampel, F. R., Ronchetti, E. M., Rousseeuw, R. J. & Stahel, W. A. (1986), *Robust Statistics: The Approach Based on Influence Functions*, John Wiley & Sons.
- Huber, P. J. (2004), *Robust Statistics*, John Wiley & Sons, Hoboken.
- Klugman, S. A., Panjer, H. H. & Willmot, G. E. (1998), *Loss Models: From Data to Decisions*, Wiley, New York.
- Knez, P. J. & Ready, M. J. (1997), ‘On the robustness of size and book-to-market in cross-sectional regressions’, *Journal of Finance* **52**, 1355–1382.
- Martin, R. D. & Simin, T. T. (2003), ‘Outlier resistant estimates of beta’, *Financial Analysts Journal* **59**, 56–69.
- Panjer, H. & Willmot, G. (1992), *Insurance Risk Models*, Society of Actuaries, Schaumburg.
- Patrik, G. (1981), Estimating casualty insurance loss amount distributions, *in* ‘Proceedings of the Casualty Actuarial Society’, Vol. 67, pp. 57–109.
- Rousseeuw, P. J. & Leroy, A. M. (2003), *Robust Regression and Outlier Detection*, John Wiley & Sons, Hoboken.
- SwissRe (2004), ‘Sigma preliminary report’.