Modelling risk in forest insurance
- Extreme value and frequency analysis of insurance claims due to storm damaged forest

Anders Gustafsson
Kristoffer Fagerström
RH11

Faculty of Engineering, LTH
Lund University

May 7, 2013

Abstract

The last decade have brought some extremely large expenditures to cover damages caused by storm weather for the Swedish forest insurer Länsförsäkringar AB (LFAB). The extent of some of those damages have been on levels that were not fully covered by reinsurance and thus caused substantial economic losses for LFAB. It is therefore of great interest to be able to understand and predict those storm damages in terms of extent and frequency. The goal of this thesis was thus to derive statistical models for severity and frequency of larger storms by using data in the form of insurance claim payments due to storm damages on forest. A generalised pareto distribution (GPD) was fitted to large payments but could not explain the variance of the worst damages in a satisfying way. Different covariates were therefore introduced in the scale parameter to improve the model. Earlier studies have proposed that strong winds alone are not able fully explain the variance of those damages so for that reason a part of this thesis was dedicated to explore possible dependencies between suggested covariates and storm damages. These covariates included factors that could influence soil stability and tree stand composition. In addition to this the wind measure showing most correlation to severe damages had to be identified. Model testing led to the final model which used two different covariates in its scale parameter: counts of gust winds over 24.5 m/s and temperature. The frequency of events was expected to follow a Poisson distribution which is also the industry standard for this kind of events. It however turned out that the variance in the number of annual storms was significantly larger than the mean and the Poisson model thus had to be rejected. Instead, a negative binomial distribution that allows for more variance in the data proved to be a better fit and was therefore suggested as a model for the frequency of storm damages.

Acknowledgement

Contents

Chapter 1

Introduction

1.1 Background

Storm damage is one of the main causes of large expenditure for forest insurers. These damages occur irregularly and the extent of them tend to vary greatly. The worst and one of the most recent examples of this was the storm Gudrun/Erwin in January 2005 that caused forest damages twice as extensive as those caused by the second worst storm recorded in November 1969 (Bengtsson & Nilsson, 2007). This thesis is based on a dataset of insurance claim payments made by Länsförskringar AB (LFAB) due to storm damages during the years 1983-2011 for the southern parts of Sweden. See Section 1.3.1 for a map of the covered region. This means that the forest damages are not measured in terms of destroyed or damaged volumes of forest but in monetary measures of those volumes. The storms Gudrun and Per (January 2007) are coupled with the largest sums of claims during the period with Gudrun being coupled to a sum that is about tenfold of that of Per. The proportions between different storm damages can be seen in Figure 1.1. For further details on the LFAB claim dataset see Section 1.3.1.

Data such as this, with some values being much bigger than others and irregular event occurrences, can often be modelled with extreme value statistics (EVS). This is a commonly used method to model natural hazards which often seem to follow this behaviour and it has been shown to work fairly well both for a dataset of Swedish forest damages from 1965-2007 (Bengtsson & Nilsson, 2007) and for a storm claim damage dataset from 1982-2005 (Brodin & Rootzén, 2009). For an introduction to and theory of EVS, see Section 2.1. The event occurrence times, or frequency of storms, is often seen as a completely random process in time, i.e. a one dimensional homogeneous Poisson process. There is however some indications that this is a too simple model, and that time inhomogeneities or clustering have to be taken into account (Mailier et al, 2006). For more details and theory about the frequency of storms see Section 2.2 and for the and methods used in this thesis, both in the analysis of storm frequency and storm
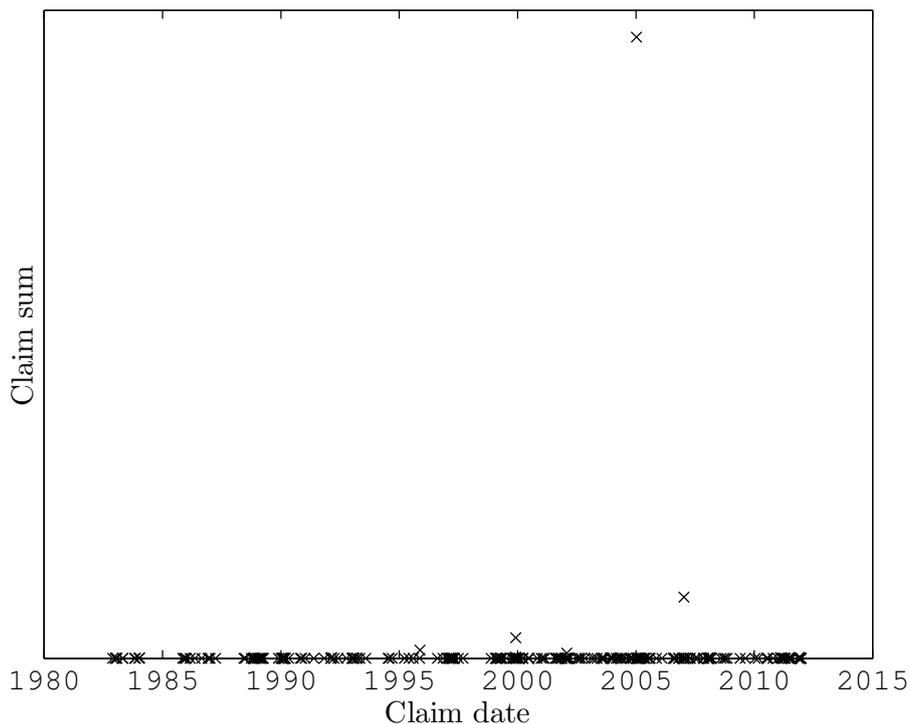
Figure 1.1: The LFAB Claim Data, linear y-axis.

severity, see Chapter 3.

The problem of managing financial risk due to storms is not a new problem for forest insurers but during the last two decades several storms resulting in very large and historically unprecedented storm damages have hit Europe. The global reinsurance company Swiss Re (2000) argues in a report to the Intergovernmental Panel of Climate Change (IPCC) that the loss potential of storms is huge and much underestimated and that forest insurance therefore has been, from an insurers perspective, too cheap. Also, several studies have shown that there is an increasing long-term trend in storm damage in Europe (Usbeck et al, 2010; Nilsson et al, 2001; European Forest Institute, 2010). This trend is also likely to continue in southern Sweden with continuing changes in global climate (Blennow et al, 2010). A literature review by Schleshaas et al (2003) found that wind storms account for more than 50% of the primary damage to forests during the 19th and 20th century. However, this doesn't necessarily mean that there is an increasing number, or increase in the severity of storms. An extensive study by the Wasa Group (1998) (updated with a newer dataset by Alexandersson et al (2000)) found no increasing trends in storminess in the last hundred years. This is also supported by the results of studies done with local measurements for Sweden (Wern & Bärring, 2009). There are several possible reasons for this. One could be a bias in reports and information in the sense that forest damages were not reported or even insured in the same extent during the early decades of the 20th century as they are today. Another likely reason is the increase in

productive forestry area, mostly in the form of larger areas of conifer stands, especially Norway spruce, which are more susceptible to storm damage (Nilsson et al, 2001; Schleshaas et al, 2003). For a more comprehensive list of factors, in addition to wind, that could affect forest susceptibility to storm damages see (Nilsson et al, 2001) and (Schmoeckel & Kottmeier, 2008). The only certain thing in this context is that more research is needed to understand which factors that are causing the largest storm damages. This research could then lead to knowledge about how to predict future losses.

We have found two previous studies using EVS to analyse Swedish storm damage insurance claims. The first was made by Rootzén & Tajvidi (2001) and is an EVS study of storm claim data in Scania, Sweden. It aimed to explain insurance claim payments due to storm damage by using wind speed as a covariate but was only able to explain a little more than 50% of the variance. The report ended with a call for more studies on the topic. The second study, by Brodin & Rootzén (2009), used an updated insurance claim dataset and EVS to predict future losses. They tried to assess the risk of future extreme storm damages using both univariate (which is the approach used in this thesis) and bivariate EVS distributions. They did not, however, use any covariates in their models to explain the variance. In addition to claims for forest damages, the datasets used in these studies also included insurance claims due to farm and property damages. This is an important distinction as the study by Brodin & Rootzén (2009) states that forest insurance claims are the ones that seem to have the most unpredictable extreme behaviour.

This thesis can mainly be seen as a continuation and extension of these two studies, and was done with one of the authors of the first study as supervisor.

1.2 Objective and Aims

LFAB is at present time the biggest actor in the business of forest insurance in Sweden. The irregularity and variance of claim payments has meant significant difficulty for them in setting customer premiums at an appropriate level and in knowing how much reinsurance they need to cover payments for extreme storm damages. The main goal of this thesis was therefore to create a mathematical model that was able to explain and predict both the frequency of storms and the quantitative size of damage that is caused to forests as a consequence of these storms. These models would be of great use for insurance companies in their work to better estimate future costs related to forest insurance and reinsurance matters. It would also benefit general storm research, the forestry industry and other parties interested in understanding the apparent randomness of storms and storm damages.

In an attempt to continue earlier studies made on this subject and explain the variance of data similar to that used before, an important field of study was to investigate different covariates with potential impact on the damage done by a

storm. Covariates of special interest, except wind speed, that has been suggested to be of importance for the extent of forest damages includes parameters for describing soil stability (Usbeck et al, 2010) and factors like stand composition and age (Nilsson et al, 2001; Schmoeckel & Kottmeier, 2008). In order to be able to fulfil our main goal, sources of relevant data for these covariates had to be found. Further, each covariate's significance to the claim sizes had to be shown to finally be able to describe as much as possible of the variance of the storm damages. A better understanding of the factors that have a significant impact on the storms that cause forest damage makes it possible to more accurately predict trends and future financial risks. This could in turn lead to better reinsurance management for insurers.

Our goal was that by having a dataset spanning over more years than previous studies, focusing only on forest insurance claims, and by taking more parameters than wind speed into account we would be able to get better models than was previously possible.

1.3 The Data

1.3.1 Claim data

The claim data was kindly provided to us by LFAB as a set of every payment made to customers due to forest damage from 1983 to 2012 with matching date and location for each province in Götaland, except Gotland. The exact monetary amounts in this data were considered as sensitive information and were by request from LFAB not presented in text and figures in this report. In Figure 1.2 the same data is plotted as in the previous Figure 1.1 but on a logarithmic y-axis to show the variation of the smaller claims. In Figure 1.3 the total sums of claims for each province is mapped.

1.3.2 Covariate data

Literature studies of storm damages and covariate data availability led to five possible covariate groups that were considered in this study:

1. Gust wind speed.

2. Soil stability in terms of soil being frozen or not.

3. Soil stability in terms of soil being wet or not.

4. Forest stand age.

5. Forest stand composition.
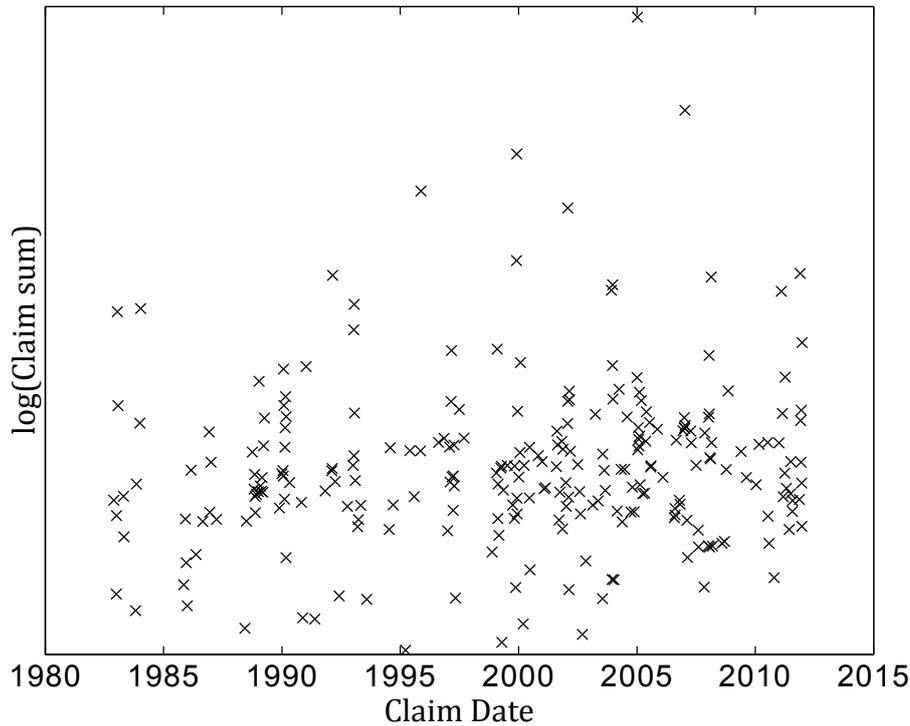
Figure 1.2: The LFAB Claim Data, logarithmic y-axis.

Meteorological data (1-3)

Meteorological data was obtained from SMHI (Swedish Meteorological and Hydrological Institute). Temperature and precipitation data used to estimate (2) and (3) was available on the SMHI website where 3-hour data of several different meteorological parameters for many measurement stations are freely available for non-commercial use[1]. For the procedure of estimating soil conditions from this data see Section 3.1.3

Records of gust winds were not freely available and had to be ordered for an administrative fee. The data per se was however free of charge for non-commercial use. Gust wind data had only been recorded at automatic weather stations that in most cases were set up during 1995. This limited the use of gust wind data to the time period of 1996-2012 but for this period hourly gust wind records for several stations in Götaland were available to us. For the procedure of creating covariate datasets from this data see Section 3.1.3.

Forestry data (4-5)

The data for forest stand age and forest stand composition were obtained from the Swedish University of Agricultural Sciences (SLU) yearly report "Riksskogstax-

---

[1]Available at: `http://www.smhi.se/klimatdata/meteorologi/dataserier-2.1102`
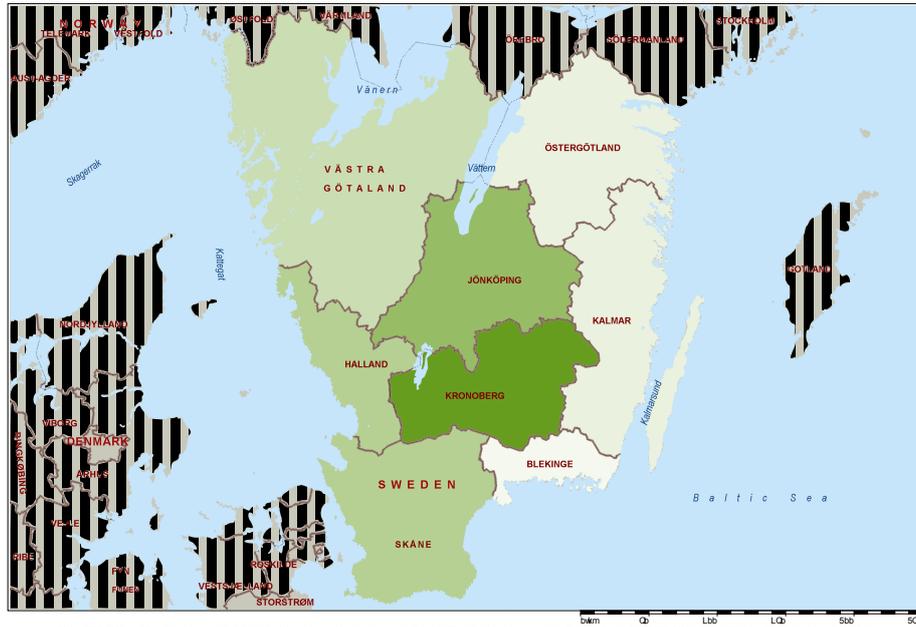
Figure 1.3: Map of southern Sweden and surrounding areas. The non-striped areas are the eight provinces in Götaland from which data in the LFAB dataset was gathered. The sum of total claims from the LFAB dataset over the whole period of time from 1983-2011 for each of these provinces are mapped on a logarithmic scale from light to dark, where the darkest areas corresponds to the largest sums of claims.

eringen". A web-based tool called Taxwebb[2] for searching this publicly available data was used to retrieve yearly data on a provincial spatial scale. For the procedure of creating covariate datasets from this data see Section 3.1.3.

1.4   Disposition

The following chapters are ordered so that Chapter 2 explains the fundamentals of extreme value theory on a level that is needed to understand the method and results of this thesis. The method is described in Chapter 3 and provides the reader with a description of data handling, important definitions and the work of finding suitable models for extreme events and their frequency. Chapters 4 and 5 contains the results of the study and a discussion of those. Conclusions and topics for further studies are also included in the discussion. Finally, all references that have been used or mentioned can be found in the bibliography last in this report.

---

[2]Available at: `http://www-taxwebb.slu.se/Taxwebb/TabellForm/index.html`

Chapter 2

Theory

The following sections are dedicated to the fundamental theory that this thesis is based on and aims to give the reader a better understanding of the method, results and discussion that are presented below. The reader is expected to have basic knowledge of probability theory and statistics. While the majority of this chapter is devoted to extreme value statistics, a brief section describing the Poisson and the negative binomial distribution used in the frequency analysis is included together with a description of generalised linear models.

2.1 Classical Extreme Value Theory

Theory on extreme value statistics presented below is based on Coles (2001), which is also recommended for a more thorough review of the theory of statistical modelling of extreme values. For the truly dedicated reader Leadbetter et al (1983) is recommended for a profound guide to the mathematics behind extreme value models.

2.1.1 Introduction

The classical extreme value theory is based on the asymptotic behaviour of

$$M_n = \max\{X_1, \ldots, X_n\} \tag{2.1}$$

where $X_1, \ldots, X_n$ is a sequence of independent and identically distributed (iid) variables with the distribution function F_X. If F_X is known, the distribution

function of M_n could be derived as follows:

$$F_{M_n} = P\{M_n \leq z\}$$
$$= P\{X_1 \leq z, \ldots, X_n \leq z\}$$
$$= \prod_{i=1}^{n} P\{X_i \leq z\}$$
$$= F_X(x)^n$$

Unfortunately, F_X is normally not known which means that one would have to rely on an approximated model to be able to determine F_{M_n} in this way. However, a seemingly small error in the estimate of F_X could then have a significant impact on the estimate of F_{M_n} because of the power function. This method for approximation is thus not preferable.

Extreme value theory suggests another way to deal with this problem that doesn't involve an estimation of F_X. This builds on the fact that F_X is unknown and that F_{M_n} belongs to a family of distributions, the extreme value distributions, which could be estimated from a sample of extreme value data. The argument behind this can be likened with the assumption of a normal distribution for a sample mean as stated by the central limit theorem. If x_+ is the smallest value for which $F_X(x) = 1$, $F_{M_n}(x)$ will converge to 0 as $n \to \infty$ for all $x < x_+$. The consequence of this is that the distribution of M_n will degenerate to a point mass at $x = x_+$. However, this problem is removed by a linear renormalisation of M_n according to:

$$M_n^* = \frac{M_n - b_n}{a_n}$$

where $a_n > 0$ and b_n are sequences of constants that, if correctly chosen, stabilises the location and scale of M_n^* as n increases. For that reason it is the limiting distribution of M_n^* instead of M_n that is of interest in this case. It can be shown that regardless of the initial distribution F_X the limiting distribution of M_n^* will always belong to one of the following distributions:

1. $G(x) = \exp\left[-\exp\left[-\left(\frac{x-b}{a}\right)\right]\right] \quad -\infty < x < \infty$

2. $G(x) = \begin{cases} 0 & x \leq b \\ \exp\left[-\left(\frac{x-b}{a}\right)^{-\alpha}\right] & x > b \end{cases}$

3. $G(x) = \begin{cases} \exp\left[-\left[-\left(\frac{x-b}{a}\right)\right]^{\alpha}\right] & x < b \\ 1 & x \geq b \end{cases}$

These distribution families together form the extreme value distributions and are known as the Gumbel, Fréchet and Weibull distributions respectively. Each of these families has a location parameter and a scale parameter corresponding to b and a, where $a > 0$. Families two and three also have a shape parameter $\alpha > 0$.

2.1.2 The Generalised Extreme Value Distribution

The possible limiting distributions for M_n^* are summarised into a single distribution family called the generalised extreme value distribution (GEV) defined as:

$$G(x) = \exp\left[-\left(1 + \gamma\left(\frac{x-\mu}{\sigma}\right)\right)^{-1/\gamma}\right] \qquad (2.2)$$

where μ, σ and γ are location, scale and shape parameters and $-\infty < \mu < \infty$, $0 < \sigma$ and $-\infty < \gamma < \infty$. Depending on the shape parameter the GEV can represent any of the extreme value distributions. For $\gamma > 0$ or $\gamma < 0$ the GEV turns into a Fréchet or a Weibull distribution, respectively. In the case of $\gamma = 0$ the limit $\gamma \to 0$ applied on the GEV turns it into a Gumbel distribution. Theory states that if there exists sequences of constants $a_n > 0$ and b_n so that

$$P\left\{\frac{M_n - b_n}{a_n} \leq x\right\} \longrightarrow G(x) \quad \text{as } n \to \infty \qquad (2.3)$$

for a non-degenerate distribution function $G(x)$, then must $G(x)$ belong to the GEV family stated in (2.2). The limit in (2.3) implies that the GEV family of distributions would be a good choice to model the distribution of maxima for long sample sequences. Moreover, it also implies that if the distribution of M_n^* can be approximated by a member of the GEV family, the distribution of M_n can be modelled with another member of the same family. Since the constants a_n and b_n are included in the model parameters, this leaves us with the task of parameter estimation which could be done using a maximum likelihood estimator for example. To sum up this basic form of extreme value modelling, the general working procedure can be explained as follows: A given set of independent observations $X_1, X_2 \ldots$ are blocked into n long sequences of equal length, often corresponding to a certain time span as a year. Each block then contributes with its maximum value, M_n, to a dataset which constitutes the basis for the GEV parameter estimation. It is important to note that the choice of block size will have a big impact on this estimation. Large blocks will render few block maxima, which will bring a greater variance to the estimated parameters. On the other hand, small blocks means that the approximation by the limit model in (2.3) will deteriorate resulting in a biased estimate.

2.1.3 The Generalised Pareto Distribution

The blocking of data performed when a GEV model is used may entail that important information in the data is overlooked. For example, if the two most extreme data points in the same sequence are about the same only the most extreme will be included in the block maxima dataset. For this reason an alternative way of modelling the extreme values is offered by the Generalised Pareto Distribution (GPD). Let $X_1, \ldots, X_n$ be a sequence of iid variables with the distribution function F_X and define an extreme value as a value of X_i exceeding a

threshold u. Conditional probability then yields that

$$P\{X > y + u | X > u\} = \frac{P\{X > y + u\}}{P\{X > u\}} = \frac{1 - F_X(y + u)}{1 - F_X(u)} \quad y > 0 \qquad (2.4)$$

Since the original distribution, F_X, is not known in this case either it has to be approximated. Theory, that is not covered here, leads to the following result: Let M_n be defined as in (2.1) and suppose F_X satisfies (2.3) so that $P\{M_n \leq x\} \approx G(x)$ as $n \to \infty$, where $G(x)$ is the GEV with parameters μ, $\sigma > 0$ and γ. It can then be shown that, provided a threshold large enough,

$$P\{X - u > y | X > u\} \longrightarrow H(y) \quad \text{as } u \text{ increases} \qquad (2.5)$$

where $H(y)$ is the Generalised Pareto family of distributions defined as

$$H(y) = 1 - \left(1 + \gamma \frac{y}{\tilde{\sigma}}\right)^{-1/\gamma} \quad y \geq 0 \qquad (2.6)$$

and

$$\tilde{\sigma} = \sigma + \gamma(u - \mu) \qquad (2.7)$$

If $\gamma > 0$ the maximum of the GPD is unbounded and the distribution is said to be heavy tailed. On the other hand, if $\gamma < 0$ the distribution has an upper bound $u - \tilde{\sigma}/\gamma$. In the case of $\gamma = 0$, $H(y)$ reduces to

$$H(y) = 1 - \exp\left(-\frac{y}{\tilde{\sigma}}\right) \quad y \geq 0 \qquad (2.8)$$

which is recognised as the exponential distribution with parameter $1/\tilde{\sigma}$.

The choice of threshold when modelling the exceedances corresponds to the choice of block size when modelling block maxima. A higher threshold means fewer exceedances and thus a greater variance, while a lower threshold leads to more bias in the estimate. This trade-off between variance and bias calls for a method that makes it possible to identify an optimal threshold. In practice one is looking for the lowest possible threshold that still provides a decent approximation in (2.5) as this yields a relatively low variance and acceptable bias. Principally there are two methods for deciding a threshold level that takes this trade-off into account. The first method makes use of the fact that the expected value of the GPD is a linear function of the threshold level. The expected value of a GP distributed variable X is defined as

$$E(X) = \frac{\sigma}{1 - \gamma} \qquad (2.9)$$

where $\gamma < 1$. For a GPD where $\gamma \geq 1$ the expected value is infinite. If a GPD with scale parameter σ_{u_0} is a valid model for excesses over a threshold u_0, another GPD could be used to model excesses over a threshold $u > u_0$. That GPD would have the same shape parameter but according to (2.7) its scale parameter would be

$$\sigma_u = \sigma_{u_0} + \gamma u \qquad (2.10)$$

13

and the expected value of exceedances over the threshold u would then be

$$E(X - u | X > u) = \frac{\sigma_u}{1 - \gamma} = \frac{\sigma_{u_0} + \gamma u}{1 - \gamma} \qquad (2.11)$$

Hence, by plotting the mean residual life (MRL) plot containing the set of points

$$\left(u, \frac{1}{n} \sum_{i=1}^{n} (x_{(i)} - u) \right) \quad \text{for } u < x_{max}$$

where $x_{(i)}$ corresponds to the i:th ordered value of x, the values of u corresponding to a linear area in the resulting graph constitutes valid choices of threshold levels. Since the lowest possible threshold is desired, a suitable choice would be the value of u that marks the beginning of the linear area. In Figure 2.1 that point would probably lie around 16 m/s, suggesting a threshold in the vicinity of this value to be used if a GPD were to be fitted to the data.
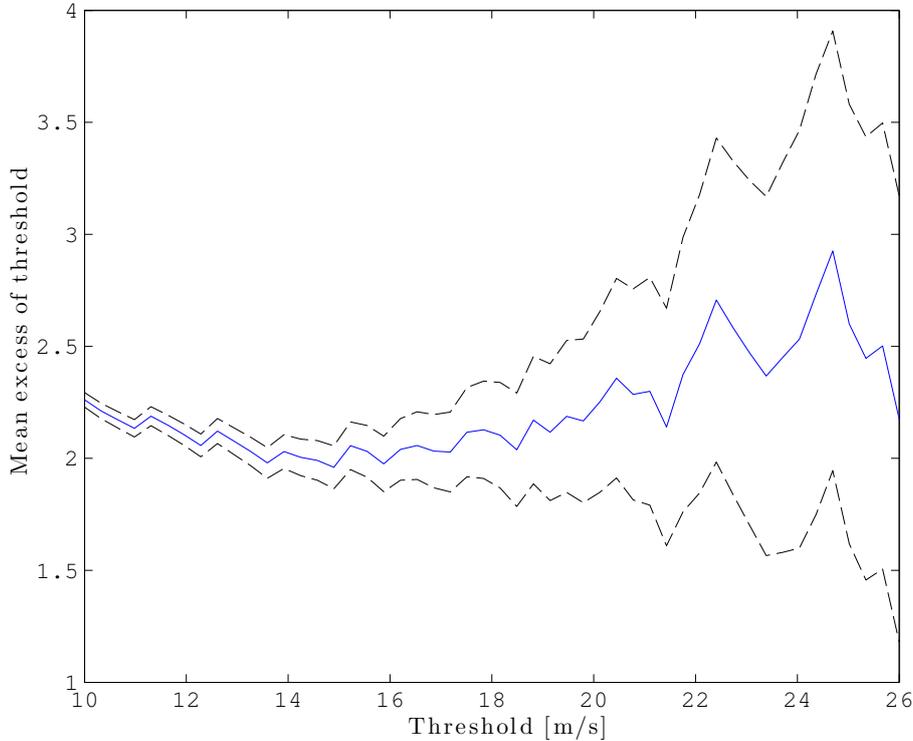
Figure 2.1: MRL plot with confidence bounds of gust winds recorded in Växjö between 1996 and 2012. The plot has been truncated at 10 m/s and 26 m/s for illustrative purposes.

In addition to the MRL plot one should always study the behaviour of parameter estimates for different thresholds. Theory states that given that a GPD can be used as a model for points over a threshold u_0, points exceeding a higher threshold u can also be modelled with a GPD. These models will share the same shape parameter but their scale parameters will differ from each other. Suppose that points exceeding a threshold u_0 follows a GPD with parameters σ_{u_0} and γ

14

and that $\gamma \neq 0$. Points exceeding a higher threshold u will then, according to (2.7), follow a GPD with scale parameter

$$\sigma_u = \sigma_{u_0} + \gamma(u - u_0) \tag{2.12}$$

This expression can be reparameterised so that

$$\sigma^* = \sigma_{u_0} - \gamma u_0 \tag{2.13}$$

where σ^* is called modified scale. This expression is only depending on u_0 meaning that the modified scale will be constant for values of u exceeding u_0. Hence, by plotting the modified scale against a range of different thresholds, the modified scale should be constant for thresholds above u_0 meaning that values exceeding this threshold can be modelled with a GPD. As the shape parameter should be unaffected by a threshold shift, the corresponding plot for the shape parameter should also be constant above u_0. Plots for the modified scale and the shape parameter based on the same data as Figure 2.1 can be seen in Figure 2.2. As can be seen in these plots, they both give support for a threshold choice around 16 m/s as the parameters appear to be rather constant after that point. These plots also makes the wish for a low as possible threshold obvious as the variance of the parameter estimates increases rapidly with higher thresholds.

2.1.4 Dependency

The theory behind the extreme value distributions assumes that there exists no dependencies in the data they are based upon. However, as far as extreme values concerns this is in reality usually not the case. For the GEV this problem is handled by selecting a block size that is large enough to make the block maxima independent of each other, provided that long range dependencies of extreme levels are weak. As for modelling with GPD another method is required since all values exceeding a threshold are used. Often these values appear in small groups with strong intra-dependencies that need to be taken care of. This is most often done using a method called declustering which can be simply described by the following steps:

1. Define clusters by an empirical rule

2. Pick out the maximum value inside each cluster

3. Assume that cluster maxima are independent and fit the GPD to those values

The cluster definition in the initial step above is often made so that consecutive values above the threshold are considered to belong to the same cluster. A value r is then decided so that a cluster is ended when there have been at least r values observed below the threshold. Again, the choice of r will mean a trade-off between bias and variance as a too small r will bring more dependence between clusters while greater values of r will render fewer clusters and a possible loss of usable data.
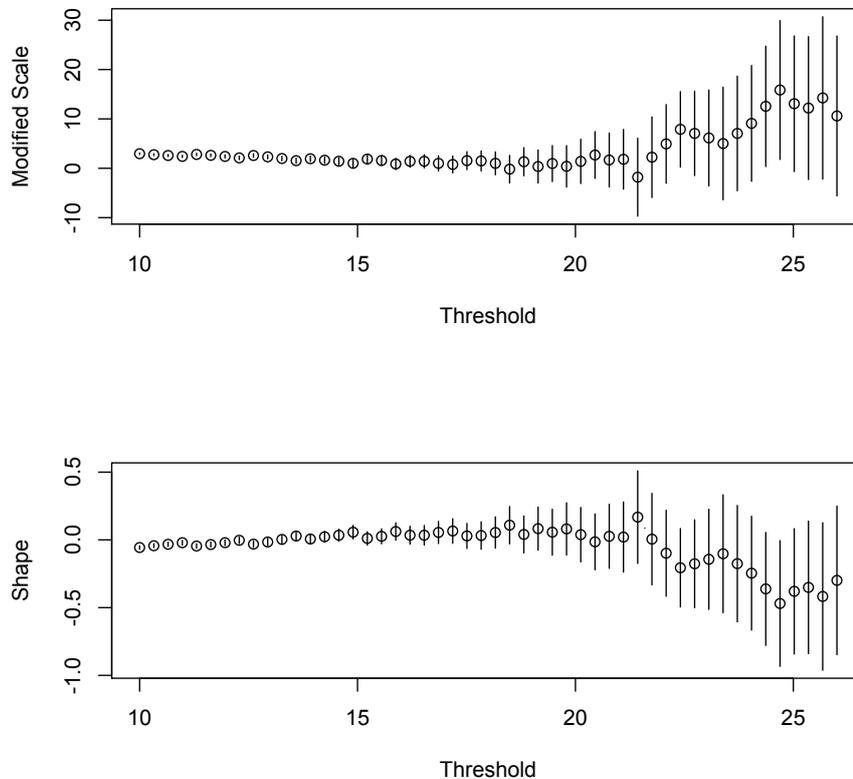
Figure 2.2: The modified scale and shape parameter with confidence bounds for different thresholds.

2.1.5 Non-Stationarity

Stationarity is an important attribute of the underlying stochastic process that generates the extreme values. In such a process the stochastic properties are expected to be homogeneous through time. That is, if $X_1, X_2, \ldots$ constitutes the process, $X_1, X_2, \ldots$ must be identically distributed and the joint distribution between any of the variables must not react to shifts in time or space. It is important to note that there may exist dependencies in a stationary process and that such dependencies doesn't contradict the stationarity property of the process. Extreme value theory is however often applied to physical processes where short and long term trends are plausible and sometimes apparent as in the case of meteorological observations. In the analysis of extreme values it is of interest to identify unknown trends and, of course, to consider known trends in the statistical model. This is done by assuming that the parameters in the model are related to time in a certain way. For example, if wind speeds are studied one would expect higher speeds during autumn and winter months than during the summer. Assuming a GPD model is used one could take this into account

by applying a lower threshold during the summer than during the autumn and winter. Another possibility is that yearly maxima of wind speeds are studied and the general wind climate is expected to get windier with time because of climate change or for some other reason. If a GEV model is used, a linear time trend could then be introduced in the location parameter μ as

$$\mu(t) = \beta_0 + \beta_1 t \qquad (2.14)$$

where β_1 corresponds to the yearly increase ratio for the maximum wind speed. More complex models such as quadratic or trigonometric models could also be used to mirror other dependencies that may exist. Similar models may be applied for the GEV and GPD scale parameter σ although these models are often exponential to ensure that σ is always a positive value. Since the shape parameter γ is hard to estimate with high precision a time dependent model for the parameter would normally not be realistic. Apart from time, a process may be related to other variables usually called covariates. For example, if race times for sprinters are expected to correlate with variables connected to their preparation or gear, one would want to introduce factors corresponding to such variables in the model as well. This is done in a similar way as with time dependencies but instead of a time function an indicator function is used.

2.1.6 Model Testing

In the process of deriving a suitable model for extreme values several plausible variants must be compared to each other to be able to decide which model that best describes the data. This is done by a hypothesis test that compares one model with another to decide if there's a significant difference between their abilities to explain data variation. The main rule when looking for the best model is parsimony; the simplest model explaining as much variation as possible is preferred. If M_0 and M_1 are two different models where $M_0 \subset M_1$, a simple test for comparison of nested models is done using the deviance test statistic

$$D = 2(l_1(M_1) - l_0(M_0)) \qquad (2.15)$$

where $l_0(M_0)$ and $l_1(M_1)$ are the maximised log-likelihood functions for model M_0 and M_1 respectively. A large value of D suggests that model M_1 explains the data better than model M_0 while the opposite holds for small values of D. To decide wether the difference is significant or not the value of D is compared to the distribution of the deviance function. This means that model M_0 is rejected at the significance level α if D is bigger than the $(1 - \alpha)$ quantile of the χ_k^2 distribution, where $0 < \alpha < 1$ and k is the difference in dimensions between M_1 and M_0.

2.2 Frequency Models

2.2.1 The Poisson Distribution

While the reader is likely to be familiar with the theory of the Poisson distribution a brief summary might be in order. Following theory may be found in any book treating basic statistical theory such as Blom et al (2009). The Poisson distribution is a discrete probability distribution describing the probability of a number of events or arrivals occurring randomly in a specified amount of time or space. To consider an event as random it must be able to occur at any time or any place and the occurrences has to be independent of each other. Further, two events may not occur at exactly the same time or place. If X is a Poisson distributed random variable, its probability density is defined as

$$P\{X = x\} = \frac{e^{-\lambda}\lambda^x}{x!} \quad x = 0, 1, 2, \ldots \text{ and } \lambda > 0 \tag{2.16}$$

where λ is describing the intensity of events. If n values from a Poisson distributed variable X has been measured, the maximum likelihood estimator of λ is given by

$$\hat{\lambda} = \frac{1}{n} \sum_{i=1}^{n} x_i \tag{2.17}$$

corresponding to the mean of the measurements, often referred to as the intensity of events. As a characteristic property for the Poisson distribution, the single parameter λ is equal to both mean and variance of the distribution.

2.2.2 The Negative Binomial Distribution

The fact that the mean and variance of the Poisson distribution are the same makes it a bad model for populations where the observed variance is bigger than its observed mean. In these cases the negative binomial distribution is often used as an alternative to the Poisson distribution as it offers a higher degree of flexibility in the shape of the distribution, courtesy of a second parameter (Klugman et al, 2004). The negative binomial probability function is defined as

$$P\{X = x\} = \binom{x + r - 1}{x} p^r (1-p)^x \tag{2.18}$$

where $x = 0, 1, 2, \ldots$ and the parameters $r > 0$ and $0 < p < 1$. This describes the probability of x number of failures before the r:th success where the probability of success in each experiment is equal to p. Corresponding mean and variance of the distribution are given by

$$E(X) = \frac{pr}{1-p} \tag{2.19}$$

$$V(X) = \frac{pr}{(1-p)^2} \tag{2.20}$$

The negative binomial can be seen as a Poisson distribution with a random gamma distributed parameter, and is often used by insurance companies to model claim frequency when there is a heterogenous risk among the insured (Klugman et al, 2004). That is, each of the insureds is expected to make a claim according to a Poisson process with an intensity λ corresponding to the insured's risk. In a pool with many insureds some will have a high claim intensity and others a lower depending on different circumstances, making the negative binomial a better option for modelling of insurance claim frequency. In other words, if the number of claims from an insured is distributed as $N|_{\Lambda=\lambda} \sim \text{Po}(\lambda)$ and $\Lambda \sim \text{Ga}(\alpha, \beta)$, then is

$$\text{Po}(\text{Ga}(\alpha, \beta)) \overset{\text{d}}{=} \text{NegBin}(\alpha, \frac{\beta}{\beta + 1}) \tag{2.21}$$

where the gamma distribution, $\text{Ga}(\alpha, \beta)$, is defined as

$$\text{Ga}(x; \alpha, \beta) = \frac{x^{\alpha-1}}{\Gamma(\alpha)} \beta^\alpha \exp(-\beta x) \quad x \geq 0 \text{ and } \alpha, \beta > 0 \tag{2.22}$$

and

$$\Gamma(\alpha) = (\alpha - 1)! \tag{2.23}$$

This is known as the representation of a negative binomial as a Poisson mixture with a gamma distributed variable (Klugman et al, 2004).

2.2.3   Non-Stationarity and Generalised Linear Models

As for data modelled with an extreme value distribution, potential trends or other sources for non-stationarity may exist in data describing the frequency of an event. An example: if the decay of a radioactive isotope is studied and its decays are counted for ten minutes every hour, the number of decays during a specific ten minute interval could probably be modelled with a Poisson distribution. But since the intensity of the decay declines as time progresses the parameter of the Poisson distribution would also decline for each ten minute period as the experiment went on. A suitable model for the number of decays during an arbitrary time interval would therefore require that a function of time is introduced in the Poisson parameter. Standard linear models assumes that the errors, that is the difference between the estimated mean and the data, are normally distributed and are therefore inappropriate to use for data that is expected to come from a Poisson distribution for example. Intuitively, this is easy to understand as a linear model of the mean of a Poisson parameter would then be negative for some values and thus contradict (2.16). Instead, a generalised linear model (GLM) can be used in these cases. This model was introduced by Nelder & Wedderburn (1972) and is made up by three components:

1. A response variable, Y, from the exponential family of distributions with parameter θ.

2. Independent variates $x_1, \ldots, x_n$, and a linear model $\mu_i = \beta_0 + \sum_{i=1}^n \beta_i x_i$.

3. A link function that connects the parameter θ to μ so that $\theta = f(\mu)$.

For the Poisson example above these components would be $Y_i \sim \text{Po}(\theta_i = \lambda_i)$. Since the expected value of a Poisson distribution is equal to the parameter λ and time is the only considered variate in this case, the linear model would be

$$\lambda_i = \beta_0 + \beta_1 t_i \tag{2.24}$$

Since λ must be positive the natural linking function would be

$$\log(\lambda_i) = \beta_0 + \beta_1 t_i \tag{2.25}$$

The deviance for a fitted model gives a measure of discrepancies between the fitted model and the data. For a large sample with n observations and a model with p estimated parameters, the deviance is approximately chi-square distributed with $n - p$ degrees of freedom (Nelder & Wedderburn, 1972). The deviance may thus be used to test the goodness of fit of the model and a hypothesis test for two nested models can be constructed in the same way as described in Section 2.1.6.

Chapter 3

Method

3.1 Data handling

A large part of the work with this thesis was the gathering and handling of data. Further, a rule that could be applied on our original claim dataset to decide which claims that were likely to be a result of the same storm had to be defined. By doing this the sizes of claims belonging to the same storm could be summarised resulting in a quantitative measure of storm severity that could be analysed. All storm losses then had to be adjusted for time inhomogeneities (e.g. inflation, see Section 3.1.1) before any analysis of the claim data could be done. In addition to this, data for different covariates had to be collected from various sources and matched in time, and to some extent geographic origin, with the storm losses. The challenge in this primarily lay in the different resolution in space and time of the different covariate datasets, requiring different approaches for each of them. The managing and creating of the needed datasets for analysis was mostly done in *Matlab*.

3.1.1 Inflation and Portfolio Changes

The assumption of time-homogeneity needed for much of the EVS analysis requires adjustment for factors that will change the size and numbers of claims over time. Inflation is the most obvious of these factors to the claim sizes and the claim data was adjusted by the Swedish Consumer price index (KPI) available from Statistics Sweden (SCB)[1], see Figure 3.1.

A second possible factor that would change both the total sums of claim sizes and also the number of claims is a change in portfolio over time, i.e. that the amount of forest insured by LFAB has changed over the years. Unfortunately LFAB does not distinguish between forest insurance and farm insurance in the portfolio so we can not know for sure how the portfolio has changed over the

---

[1]Available at: `http://www.scb.se/Pages/Product____33769.aspx`
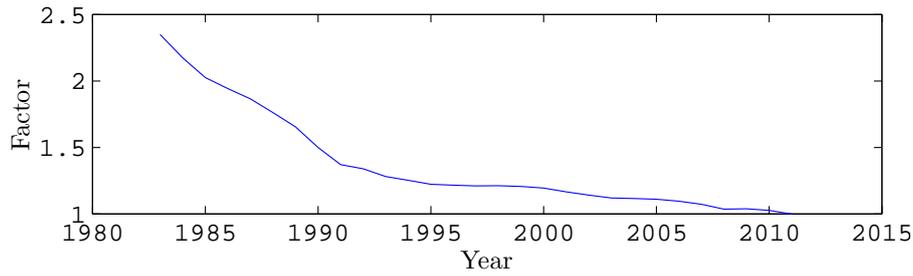
Figure 3.1: Yearly multiplication factor for claim sizes due to changes in consumer price index

years. While LFAB's total portfolio has increased this is because household and company insurance has increased, farm insurance have stayed pretty much stationary (Brodin & Rootzén, 2009). After discussions with LFAB we therefore decided to assume that the forest insurance portfolio had been close enough to constant so that no portfolio data adjustment were needed.

3.1.2 Definition storm events

A fundamental part of this thesis was to decide which of the individual claims in the LFAB dataset that could be expected to have some kind of relation in terms of common factors regarding their existence and size. The most important of these factors was believed to be strong winds as the dataset was based on forest damages in the form of fallen trees. The empirical rule that was defined for identifying related claims and merge them into storm events was therefore primarily based on assumptions regarding meteorological conditions for the south of Sweden. Our first assumption was that it was unrealistic that two separate storm weathers with strong winds would occur in Götaland on the same day because of its relatively small geographical extension compared to an average European wind storm (Bonazzi et al, 2012). In a first step we therefore assumed that claims registered on the same day, irrespective of their geographical origin, were caused by the same storm weather and summed claims that were registered on the same day. Secondly we wanted to set an upper time limit for which a storm weather could be expected to last. By studying the dataset from LFAB, this time limit was set to three days, as there on very few occasions were more than three subsequent days where claims had been registered. This definition was also in line with another study made on the subject by Brodin & Rootzén (2009) as well as with the policy of LFAB's reinsurers allowing claims filed under up to three subsequent days to be merged. Lastly, we adopted the same rule for isolating storm events as Brodin & Rootzén (2009), saying that two storm events must be separated by at least two days to be considered isolated from each other. It is also worth mentioning that these rules served as a declustering algorithm for the data corresponding to that described in the theory Section 2.1.4. With these rules in mind we let a three day window move over all days in our claims dataset, summing all days with losses during the window which resulted in a list

of total storm damages for different storm events.

All storm events that were separated according to the third rule above were then included in our storm events dataset together with their respective date of occurrence. However, a couple of special cases still remained after this. On the occasion of four consecutive days with claims two candidates was proposed by this method. In these cases the candidate with the largest damages was chosen and the other was discarded. Further, on the occasion of five or more consecutive days with claims, isolated events were formed inside these periods so that for example one lasted day 1 and 2 and another on day 5 during a five day period. These were constructed so that the total storm damages were maximised. These storms were then added to the storm events dataset. Finally, as the worst storms historically were winter storms in the Götaland region we chose to exclude all storm damages registered from April to September to ensure better stationarity in our data series, see theory Section 2.1.5. After applying these rules on the data, a set of 171 separate storm events with corresponding damage and dates distributed over the period 1983 to 2011 remained.

3.1.3 Covariates

Gust Wind

While the EVS analysis was made on the spatial scale of Götaland as a whole, the wind measurements were taken from weather stations on specific locations. To obtain data sets of daily gust wind speeds representative of the spread in geographic origin of the claims we used five different stations chosen from the available data that was spread out over the area of Götaland, see Figure 3.2. There are two things to note in their placement: First, we only chose inland stations and two of the stations were situated close to each other, both of them in the province of Kronoberg. This was due to the fact that a majority of the claims in terms of total costs came from this province. Second, while it is quite possible, even likely, that there were higher gust wind speeds recorded along the coasts they would not be as likely to coincide with large forest damages as coastal areas are usually not forested, something SMHI (2011) also have noted.

In addition to the choice of stations, measures of wind speeds in Götaland were needed. Della-Marta et al (2007) suggests two measures for this: a weighted spatial mean of wind speeds from the area and a spatial 95% quantile of wind speeds in the area. The mean is expected to be sensible to both wind speed and the spatial extent of a storm while the 95% quantile is expected to give a better estimate of the storm severity. Lastly, the highest maximum gust wind speed in Götaland during each storm was also analysed so that possible threshold effects were not missed. Such effects has been suggested by European Forest Institute (2010), meaning that very high wind speeds are needed if large scale damages are to occur. To take all of these aspects into account we created the following wind covariate datasets:

$Wind_{Mean}$: Maximum mean gust wind speed over Götaland during the storm.

$Wind_{95}$: 95% quantile of all daily gust wind speeds over Götaland during the storm.

$Wind_{Max}$: Maximum gust wind speed observed anywhere in Götaland during the storm.

As a further development of the datasets above, datasets containing the squared values of each dataset were created since basic physics states that the kinetic energy of a moving mass is proportional to its squared velocity. These datasets were called $Wind^2_{Mean}$, $Wind^2_{95}$ and $Wind^2_{Max}$.

Another approach to use the wind data to create a covariate dataset is to consider the meteorological threshold-definition of storms and count the number of occurrences of wind speeds over that threshold. This index would give both a measure of the severity of the storm and its spatial size. This was done for true meteorological storm observations (24.5 m/s) as well as a slightly lower but still high gust wind speed (20 m/s) and stored in the dataset as $WindCount_{24.5}$ and $WindCount_{20}$ respectively.
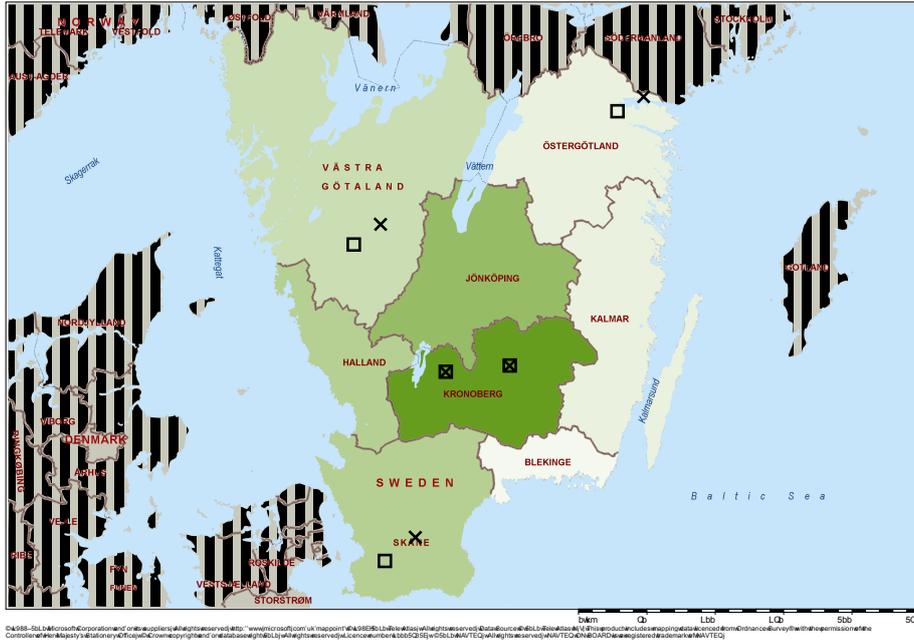
Figure 3.2: Map of southern Sweden and surrounding areas. Crosses represent the location of stations used as source for gust wind data while squares represent the location of stations used as source for temperature and precipitation data. Reasons for differences in station locations are explained in Section 3.1.3 under Soil Stability.

Soil stability

This was done with a semi-qualitative approach using data from the same station or a station nearby that recorded the gust winds, see Figure 3.2. The reason for differences in the locations of the stations (except the two in Kronoberg) was simply that from the data available to us stations recording gust winds in some cases were different to those recording temperature and precipitation. Indices of soil stability and thus storm resistance by using data for temperature and rain wasn't easily defined and required guesswork. This is because of the many other different processes and interactions that affect soil stability in addition to the amount of precipitation and temperature which was the only data available for this thesis. For example, while much precipitation in the form of rain is likely to loosen the soil and thereby lessen the grip of tree roots, an equal amount of precipitation in the form of snow would perhaps have no influence on soil stability at all. That is of course unless the snow melts and instead decreases the stability. For the scope of this thesis we chose a rather rough approach on this subject that follows below, for more on the subject see Section 5.2.2 in the discussion.

By looking at the data seven days prior to and during each storm event, the accumulated precipitation and temperature mean for each station was calculated. This was in turn used to calculate a single mean for all stations reflecting the temperature and precipitation conditions for Götaland during each storm. This information was stored in the datasets $Temp$ and $Precip$. Two interval datasets was also created as it was also possible that there were threshold effects in that data. For example, as long as the soil is frozen it doesn't matter if it is $-5°$ or $-10°$. Thus two interval datasets for temperature and precipitation were created with three intervals each:

Temperature:

$T_{hi} : x > 2°$

$T_{mid} : -2° < x < 2°$

$T_{low} : x < -2°$

In this dataset the levels can be seen as three levels of probability of frozen soil: On the T_{hi} level the soil was likely to not be frozen, on the T_{mid} level the soil could be frozen and on the T_{low} level the soil was likely to be frozen. When defining these levels we took into account that the soil is generally colder than the air due to heat emissions during the night [2]. The other dataset was defined upon the following intervals:

Precipitation:

$P_{hi} : x > 40mm$

$P_{mid} : 15mm < x < 40mm$

---

[2] http://nsidc.org/cryosphere/frozenground/how_fg_forms.html

$$P_{low} : x < 15mm$$

where the level P_{low} corresponds to precipitation amounts under the approximate weekly mean for Sweden[3]. Weekly amounts of precipitation below this level were assumed to not affect soil stability in any notable fashion. P_{mid} on the other hand, collected those storms where the precipitation levels stretched from normal to more than double of that amount, raising our suspicion of more unstable soil under those circumstances. P_{hi} with precipitation amounts even higher than that therefore reflected storms where it was most likely that the precipitation actually had influenced the soil stability.

Forestry

The forestry data was used to create yearly age group and stand type fractions in Götaland. As mentioned in Section 1.2), there is literature that singles out spruce as the stand type being most susceptible to storm damage and its fraction of the total forest volume for all storms was therefore stored in a separate dataset. Other stand types were not treated exclusively and thus fractions for conifer (except spruce) and deciduous stand types were summed and added in two separate datasets.

$$FS_{spruce}$$

$$FS_{conif}$$

$$PS_{decid}$$

As noted in Section 1.2 another suggested factor for forest storm susceptibility is the growing stock age. Thus the forest age data were grouped in four equal 40 year spans to be able to check for correlation against each age group. The datasets created were:

$$FA_{0-40}$$

$$FA_{41-80}$$

$$FA_{81-120}$$

$$FA_{>120}$$

3.2 Analysis

3.2.1 Workflow

The analysis of the claims dataset can be divided into two sub-analyses dealing with storm severity and storm frequency respectively, see Figure 3.3. The storm severity analysis was done using the POT method by fitting a GPD to storm

---

[3]`http://www.smhi.se/klimatdata/meteorologi/nederbord/1.2887`

damages above a certain threshold. For details on setting this threshold, see Section 3.2.2 below. Of the original 171 separate storm events 41 were over this threshold and was put into a separate dataset. The extreme behaviour in this data was apparent as the sum of damages during these 41 storms turned out to embody 99.98% of the total damages for all 171 storms. The identification of covariates that could possibly explain some damages was thus focused on the explanation of storm damages above the threshold. However, the same covariates were also used to see if they could explain smaller damages below the threshold or if they were tightly coupled to only severe damages. Since records of gust winds were limited to the period between 1996 and 2011 they could only be matched to storm damages during the same period. That also meant that GPD models based on the complete dataset of storm damages could only use covariates from the soil stability and forestry categories while models with all covariates were limited to the period with records of gust winds.

The storm frequency analysis, see Section 3.2.3, aimed to model the expected annual number of storms and, in particular, storms above threshold. A possible time trend in the respective Poisson parameter was also investigated and tested for. Apart from time, no other covariates were used in this analysis and it could therefore be done using data for all storms between 1983 and 2011.

While the data handling was mostly done in *MATLAB* (The MathWorks, 2012), the major part of the analysis was performed using the open source software *R* (R Core Team, 2012) and the *Extremes toolkit* (Gilleland & Katz, 2004).

3.2.2 Extreme Value Analysis

The extreme value analysis was done on the storm losses dataset described in Section 3.1.2 above. The *Extremes toolkit* contains a tool set for extreme value analysis including functions for fitting and model diagnostics based on theory described in Section 2.1.

Covariate identification

The covariates used in the models in the extreme value analysis above were based on previous studies of storm damages and possible covariates. Still, plausible relations between a possible covariate and claim sizes had to be examined. This was done by plotting scatter plots and box plots for values of covariate indices against the corresponding claim values to see if any correlation could be distinguished. If, for example, linearity was suspected between a covariate index and the logarithm of claim sizes, a linear model for the logarithm of the scale parameter was proposed for further testing. In addition to the covariates we also wanted to test for time trends in the GPD scale parameter. This was done by simply introducing a time factor that grew linearly over the timespan of the POT dataset.
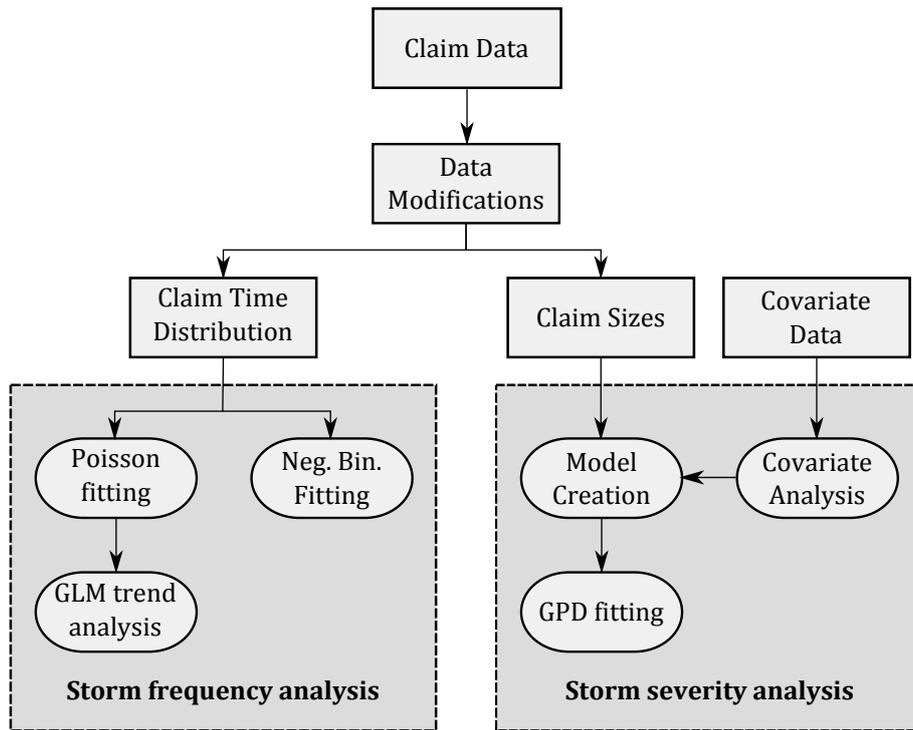
Figure 3.3: Workflow visualised. Rectangular boxes show data and data handling. Boxes with rounded corners show analysis' steps.

GPD fitting and analysis

The POT approach is more demanding in the way that it requires an appropriate threshold to be set. As a first step of the POT analysis such a threshold was identified and set using methods described in Section 2.1.3. See Figure 3.4 and Figure 3.5. The actual level of this threshold would not be meaningful to the reader as the claim sizes' actual values are hidden. The simplest possible model was then fitted to the claims exceeding the threshold. This model did not take any trends or covariates into account to any of its parameters. A range of different models, taking different covariates into consideration in the scale parameter, were then introduced. In a first step, models with log-linear functions of single covariates were studied and tested against the null model as described in Section 2.1.6. The covariates of the models that proved to be significantly better than the null model were then combined into a big composite model given that the covariate was not strongly correlated with any of the other covariates in the composite model. For example, if two different wind measures were used in separate models and both models turned out to be significantly better than the null model, only the wind measure of the "best" model was used in the composite model since the measures were strongly correlated with each other. The composite model was then reduced with one parameter at a time until the optimal model that explained as much as possible in relation to its complexity was found. See Theory Section 2.1.6.

28
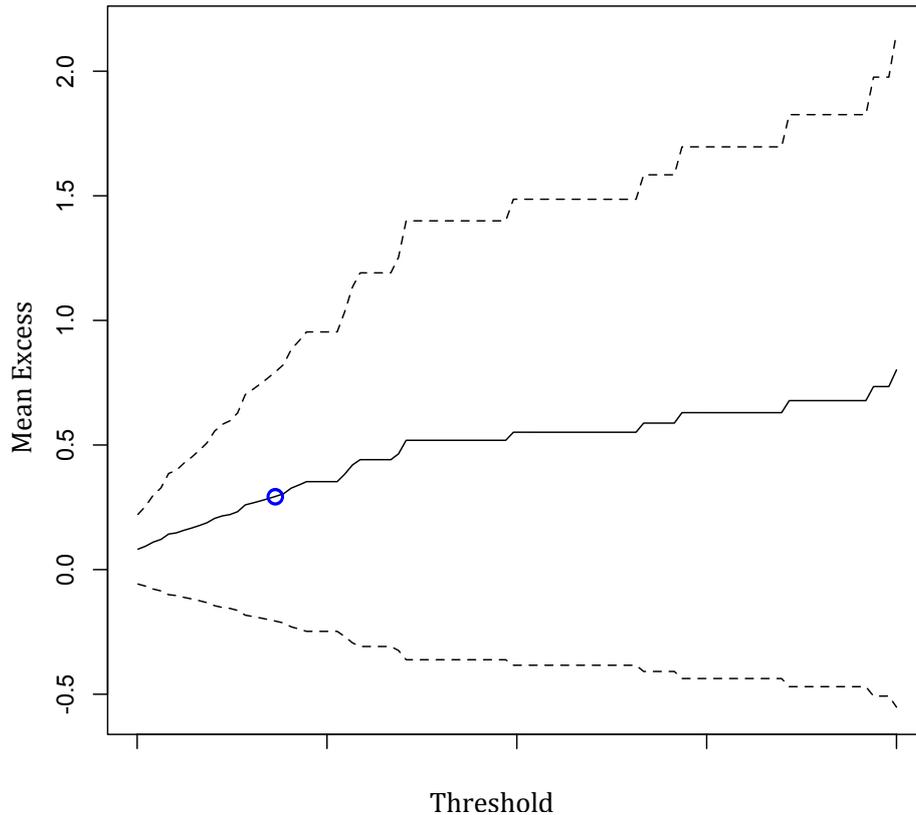
Figure 3.4: The MRL plot for claims registered between january 1996 and december 2011. The circle represents our choice of threshold.

3.2.3 Frequency analysis

The Poisson and negative binomial distributions have both been suggested as appropriate models for the frequency of extreme events as noted in Section 1.1. We therefore wanted to fit each of these models to the annual number of storm damages and storm damages over threshold to see which model that provided the best fit to the data. The Poisson distribution is the simpler of these distributions, as it only has one parameter, and is probably the most well known when it comes to frequency modelling. It also has a big advantage to the negative binomial as it is easy to introduce trends or other covariates into its parameter. This was done for a time trend using a GLM as described in Section 2.2.3. The time homogeneous Poisson model and negative binomial model were both evaluated through a χ^2-test for goodness of fit. The Poisson model with a time dependent parameter was compared to the time homogeneous ditto using theory described in Section 2.1.6.

The fitting and testing of Poisson and negative binomial was done using the *R* package *fitdistrplus* (Delignette-Muller et al, 2012) while the test for a time trend in the Poisson parameter was done using the *Extremes toolkit*. The results with plots can be found in results Section 3.2.3.
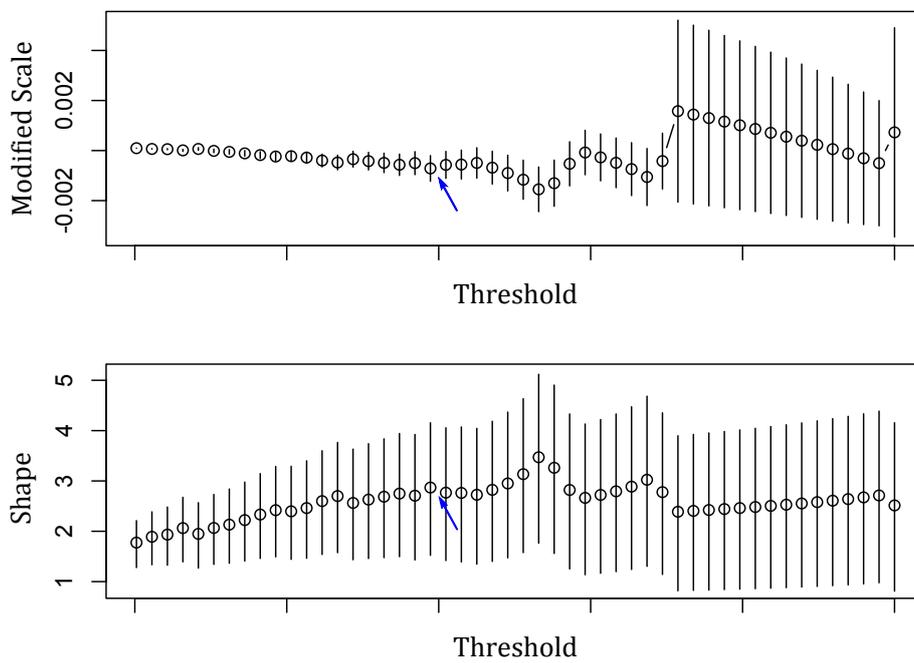
Figure 3.5: The parameter plots for GPD models fitted to different thresholds. The plots have the same scale on the threshold axis and the arrows are pointing at the same threshold level that were marked in Figure 3.4.

Chapter 4

Results

Presented below are various plots and tables describing the different covariates and models that were fitted to the claims dataset. The first section is devoted to the results of the covariates study that were described in Section 3.1.3. To get the full picture of the different covariates and to be able to see differences in dependence between a covariate and storms over threshold as well as the covariate and all storms, scatter- and boxplots for both these datasets have been included. The second section contains the results from the different GPD models that were fitted to the data. Resulting plots and table for goodness of fit is presented for the basic fit and the best fit respectively. In the last section are the results from the frequency analysis and accompanying plots for goodness of fit and an analysis of variance table.

4.1 Covariate analysis

4.1.1 Storms over threshold

Wind

Figure 4.1 shows gust wind data distribution for each of the storms and as expected there seems to be some correlation between high gust wind speeds and storm damages. Figure 4.2 depicts this further by plotting correlation for the different wind indices defined in Section 3.1.3 on page 23.

Further, scatter plots for $WindCnt_{24.5}$ and $WindCnt_{20}$, Figures 4.11a and 4.11b respectively, show that there seems to be a positive correlation between the claim sizes and gust wind speed and also between claim sizes and the number of wind speed observations over each threshold. Figure 4.4 shows the contribution of number of gust winds over 24.5 m/s split from each observation station, thus giving a hint about the spatial resolution of the worst storms. Read more in Section 5.2.1.
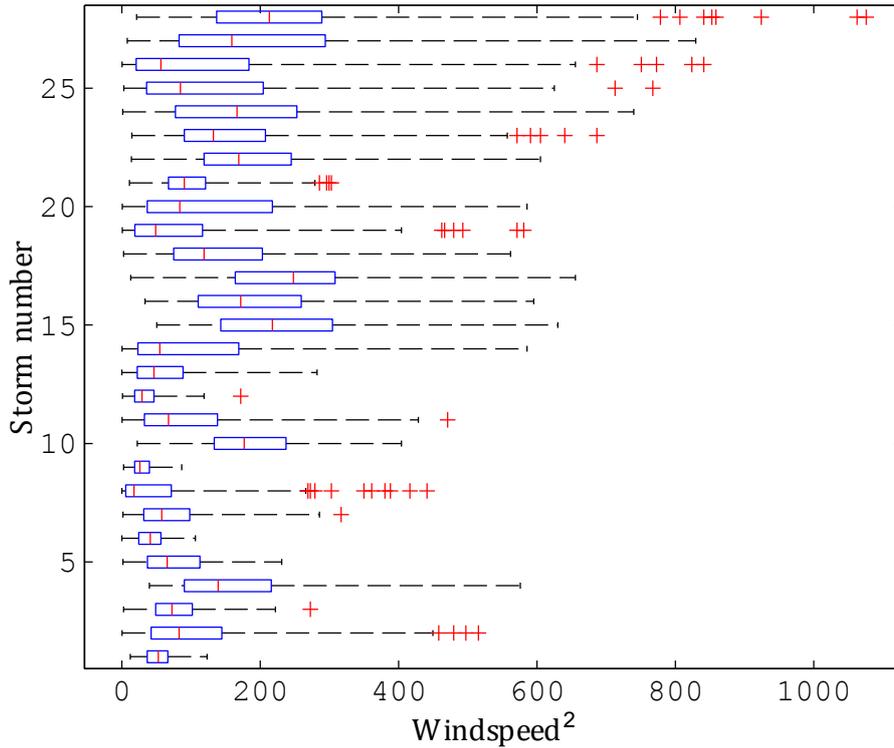
Figure 4.1: Box plot for gust winds during days when storm damages over threshold has been registered. The storm damages are ascending with the storm number. The central mark marks the median of the distribution. The whiskers marks the highest and lowest value in the distribution, if not the used algorithm sees a value as an outlier, then it's displayed as a cross. This means that the $Wind^2_{Max}$ dataset corresponds to the upper whisker limit or, as in most cases, by the cross furthest to right. The boxes show the 25% and 75% quantiles.

Ground stability

Following, in Figures 4.5 and 4.6, are the resulting box and scatter plots respectively for the soil stability covariate datasets. The box plots were based on the datasets $Temp_I$ and $Precip_I$ while the scatter plot show the full scale of temperature and precipitation data, i.e. the $Temp$ and $Precip$ datasets, see Section 3.1.3 on page 25 for a definition of these datasets. Both the temperature and precipitation data seemed to have some correlation with storm damages, especially from the looks of the box plots. The scatter plots does however raise a question mark for both temperature and precipitation and their respective correlation with storm damages.
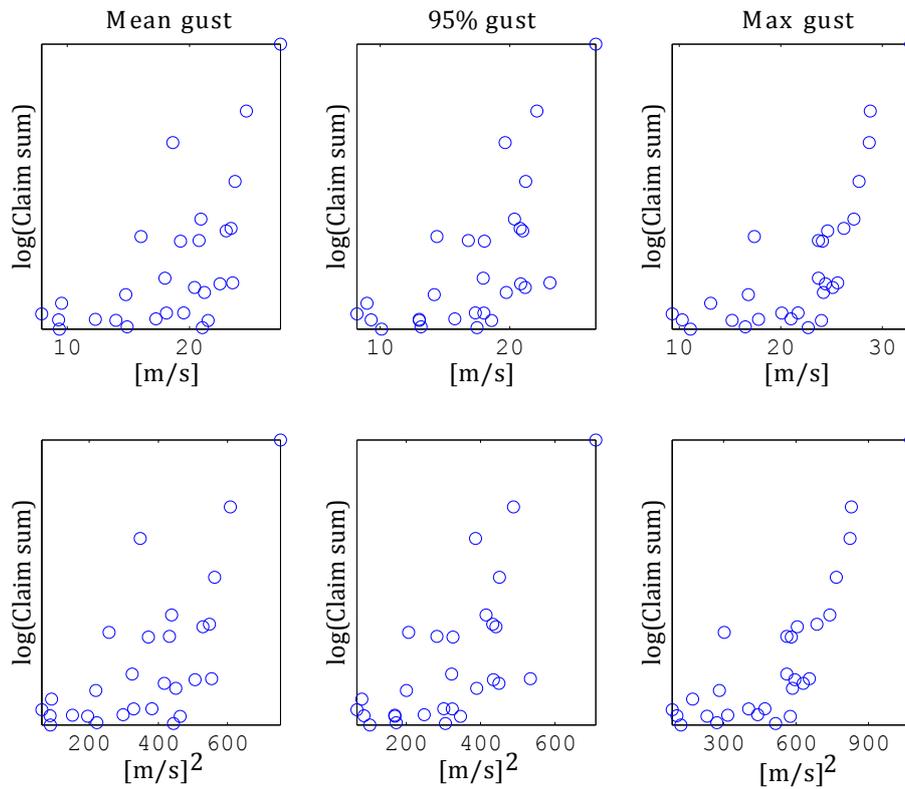
Figure 4.2: Scatter plots for claim sizes for storms over threshold after 1996 against the different wind indices and squared wind indices.
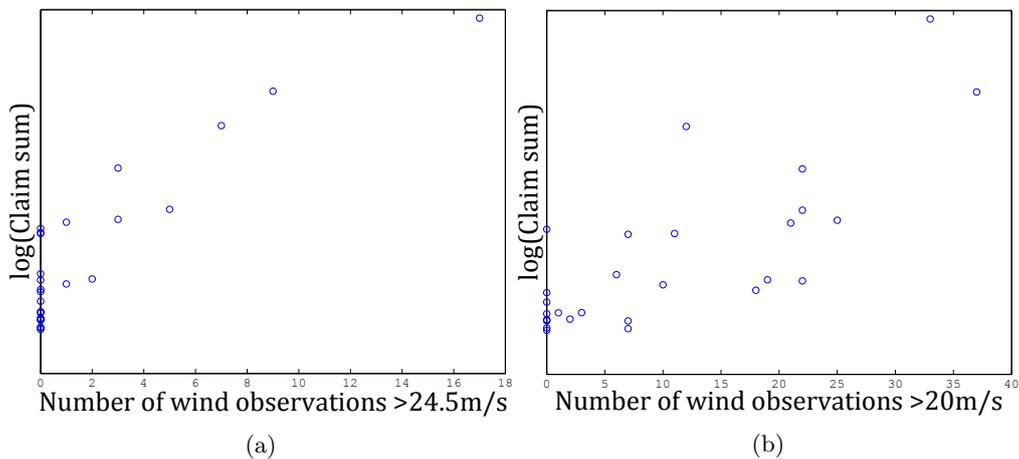
Figure 4.3: Scatter plots between logarithmic claim sizes for storms over threshold (a) $WindCnt_{24.5}$ and (b) $WindCnt_{20}$. For an explanation of these datasets see Section 3.1.3 on page 23.

Forestry

In Figures 4.7 and 4.8 are box plots for the different groups of stand types and stand ages. These show little or no correlation with storm damages.
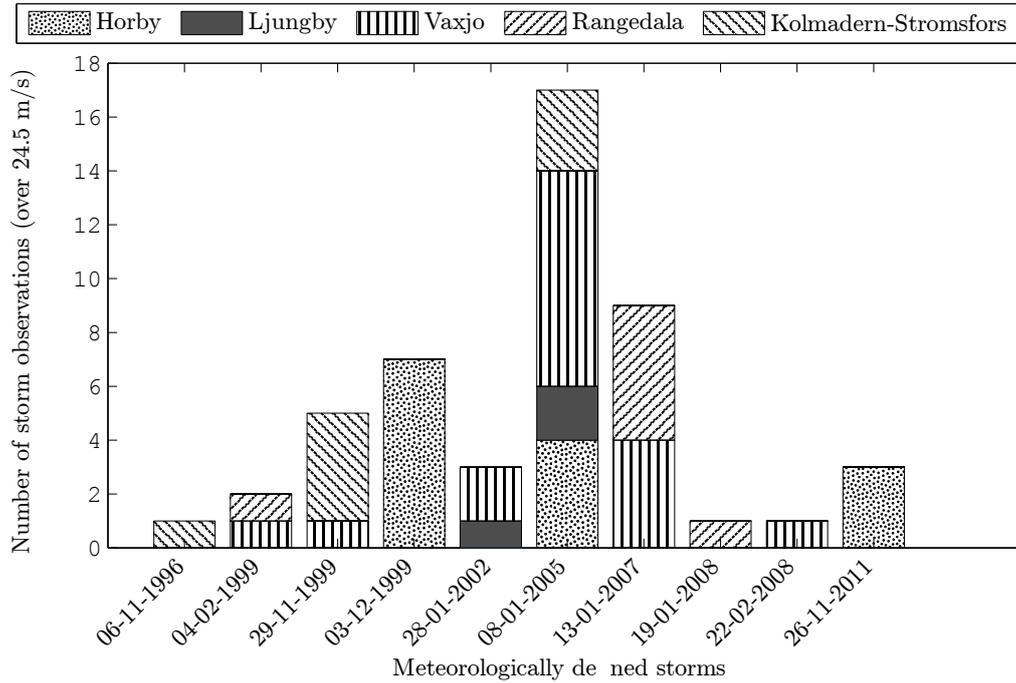
Figure 4.4: Bar plot showing the $WindCnt_{24.5}$ dataset split over each of the stations used in the analysis.
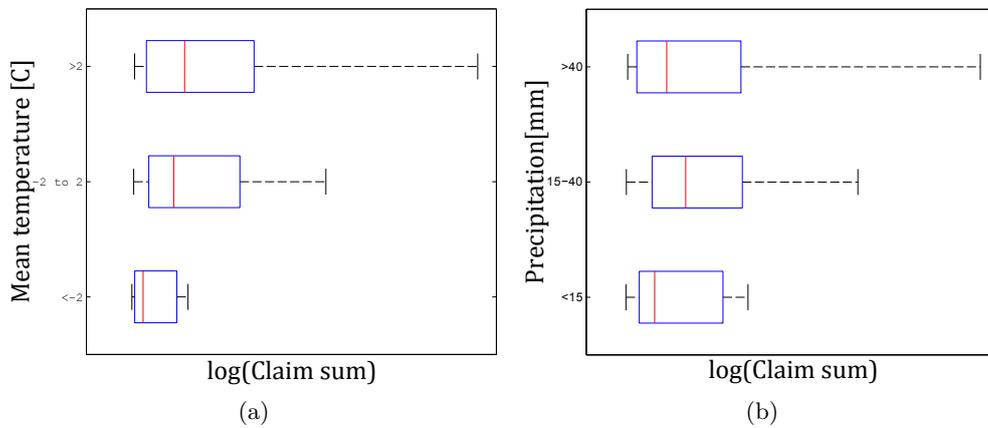
Figure 4.5: Box plots showing correlation of the claims sizes for storms over threshold against (a) $Temp_I$ and (b) $Precip_I$ seven days before and during the claim date. The central mark marks the mean of the claim sizes. The whiskers marks the highest and lowest value in the distribution not considered outliers by the algorithm, these outliers are shown as red crosses. The boxes show the 25 % and 75% quantiles.
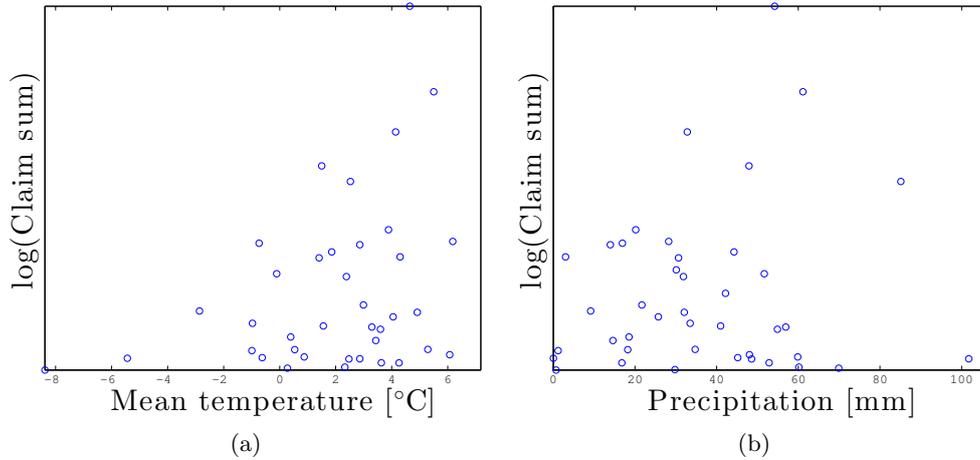
Figure 4.6: Plots showing correlation of the claims sizes for storms over threshold for the whole period against (a) *Temp* and (b) *Precip* seven days before the claim date.

4.1.2 All storms

Wind

Figure 4.9 show gust wind data distribution for each storm. While it is storms coupled with the biggest claim sizes that show the most obvious correlation, the correlation is also visible for all storms. The scatter plot was done for this dataset and is shown in 4.10 and shows that the correlation with wind speeds is evident for all storms as well.

Correlation for the wind observation counts for all storms is seen in Figure 4.11 and this shows as well evident correlation with the full dataset.

Ground stability

The ground stability box plots and scatter plots for the full storm dataset can be seen in Figure 4.12 and 4.13 respectively. The box plots still gave some support for correlation, although somewhat less for precipitation. In the scatter plot it was even harder to see if there was an actual correlation. See the continued analysis in Section 4.2.

Forestry

Considering the poor correlation between storm damages over the threshold and data for forest stand type and age, the corresponding plots for all claims, 4.14 and 4.15, are not surprising. As in the case of storms over threshold there is
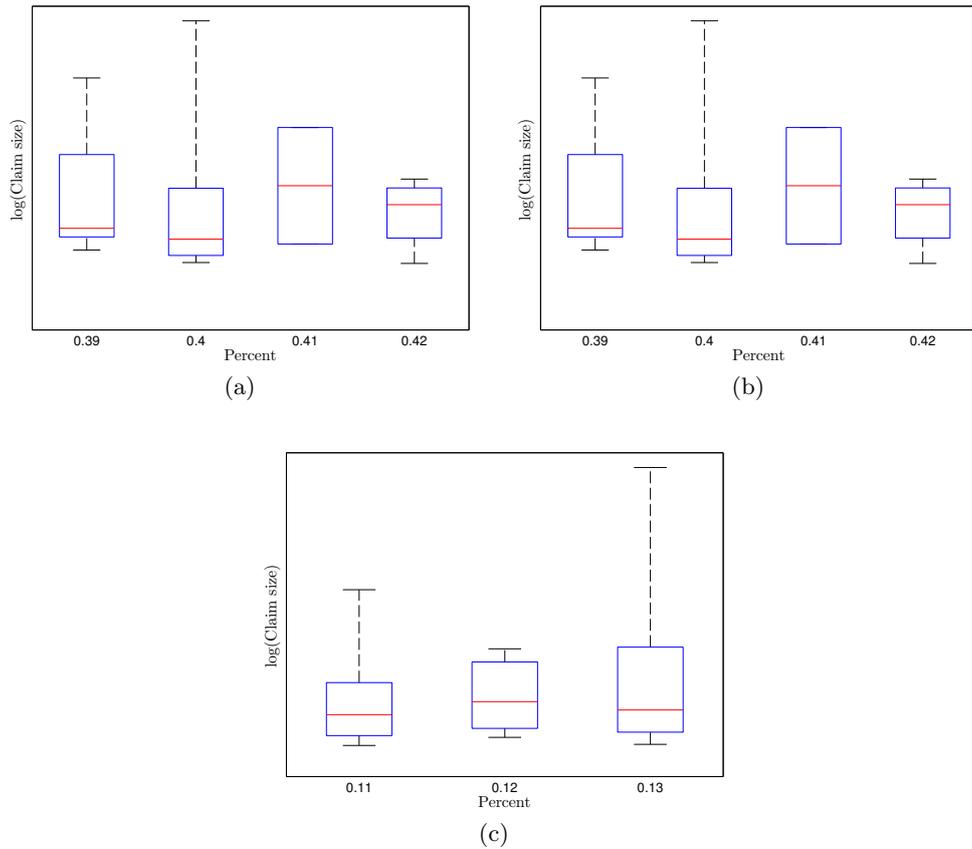
Figure 4.7: Box Plots of stand type datasets, (a) FS_{spruce}, (b) $FS_{conifer}$ and (c) $FS_{deciduous}$, percentage versus distributions of claim sizes for storms over threshold

no apparent correlation between forest stand types and forest stand ages for all storms. These datasets were thus not included in further analysis.

4.2 GPD analysis

Plots showing goodness of fit for the most basic GPD model for storm damages over threshold after 1996 can be seen in Figure 4.16. While the probability plot suggests a decent fit, the quantile plot makes it evident that the discrepancies between model and empirical values are large, especially for the higher quantiles that represent the more extreme values. Hence, the use of this model would probably mean a considerable underestimation of the probability of future storm damages on the same levels as those during the storms Gudrun and Per for example. The estimated parameters, with standard errors in parentheses, and
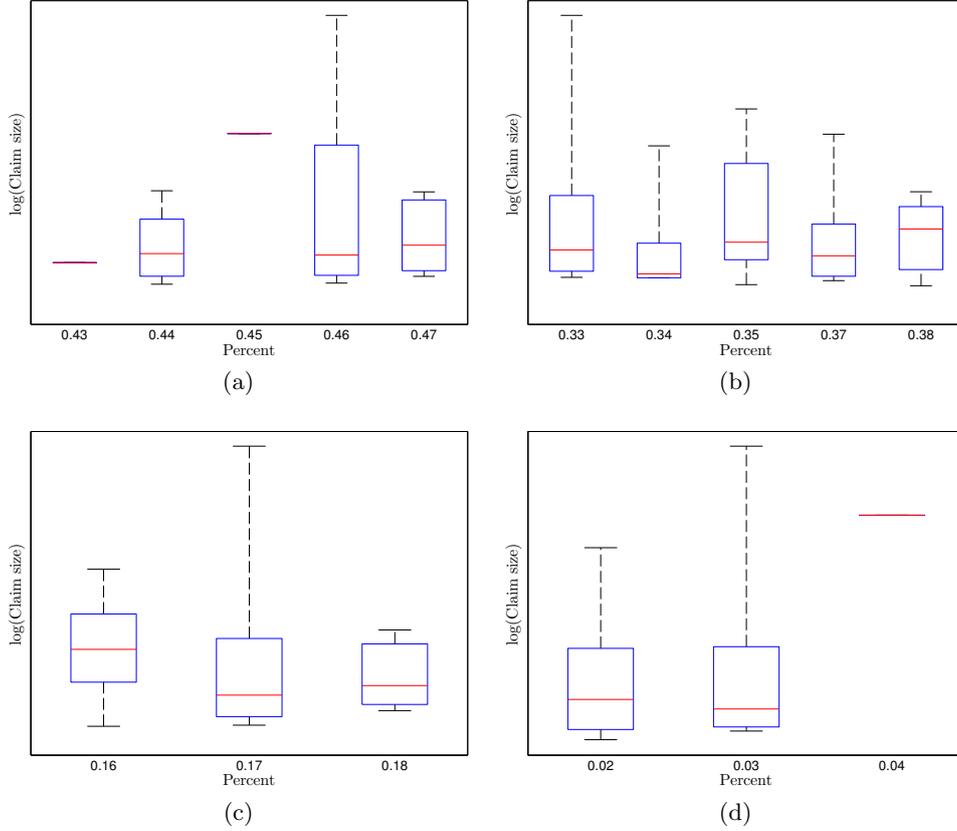
Figure 4.8: Box Plots of forest age group datasets, (a) FA_{0-40}, (b) FA_{40-80} and (c) FA_{80-120} and (d) $FA_{>120}$, percentage versus distributions of claim sizes over threshold

log-likelihood for the GPD null model were

$$\hat{\sigma} = 4.6 \cdot 10^{-4} \ (2.0 \cdot 10^{-6})$$
$$\hat{\gamma} = 2.84 \ (0.68)$$
$$\text{Llh} = 107.5$$

The estimated shape parameter was positive and greater than one, thus implying a heavy tailed and unbounded distribution with an infinite mean and variance.

Models that each take one of the different covariates described previously into account are summarised in Table 4.1 and could be considered as a complement to the scatter plots above to see which covariates that are able to explain the claims to some extent. A high log-likelihood value indicates that the covariate is explaining more of the variation in claim sizes and the value $\boldsymbol{D}$ is the deviance test statistic for the actual model compared to the null model. This value should be compared to the χ^2_k distribution with k degrees of freedom. The degrees of freedom corresponds to the difference in estimated parameters between the two models that are compared and were in all but two cases equal to one, yielding
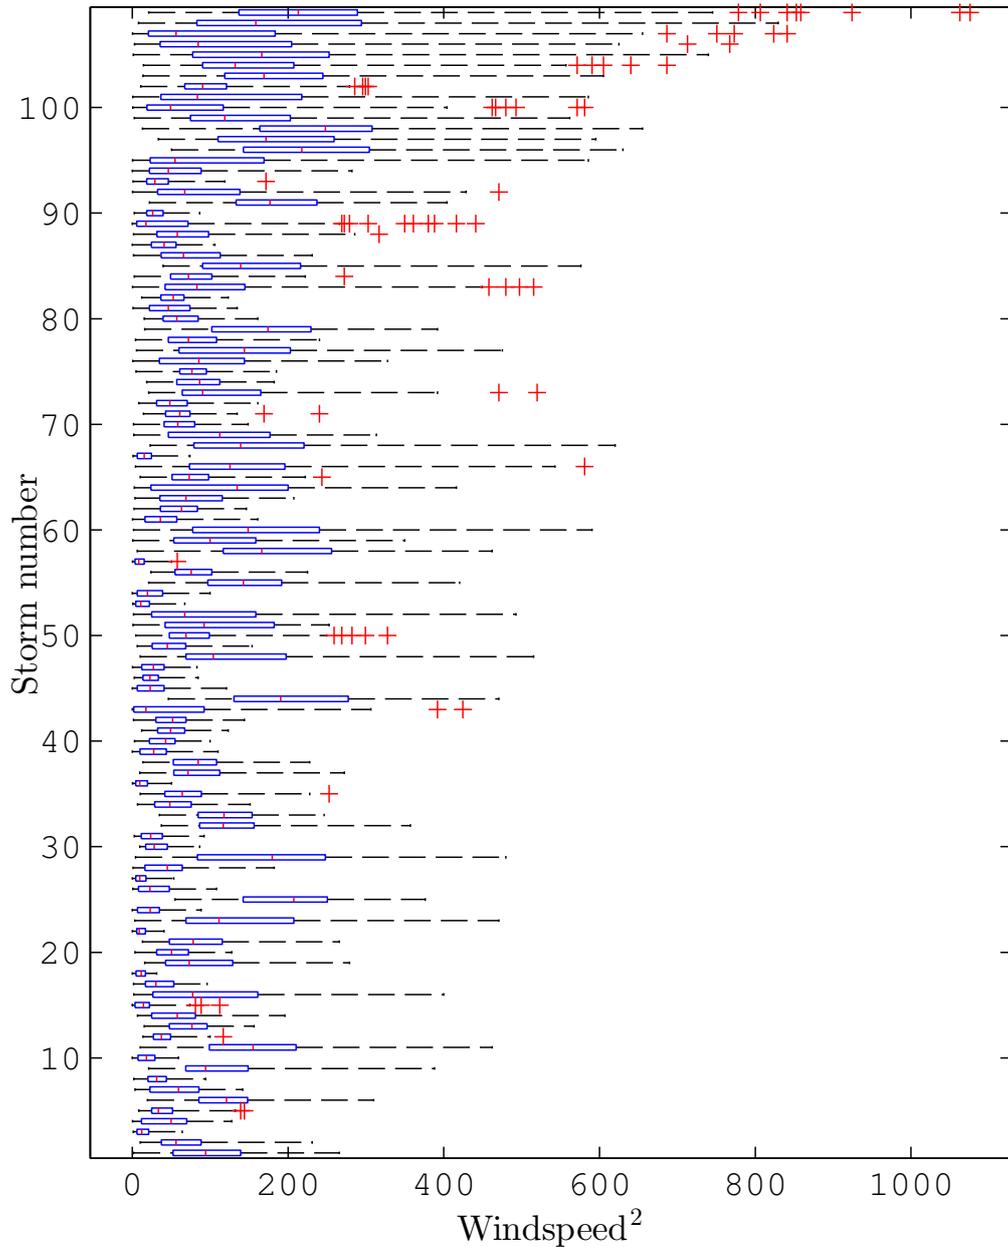
Figure 4.9: Box plot for gust winds during days when storm damages has been registered. The storm damages are ascending with the storm number. The central mark marks the median of the distribution. The whiskers marks the highest and lowest value in the distribution, if not the used algorithm sees a value as an outlier, then it's displayed as a cross. This means that the $Wind_{Max}^2$ dataset corresponds to the upper whisker limit or by the cross furthest to right. The boxes show the 25 % and 75% quantiles.
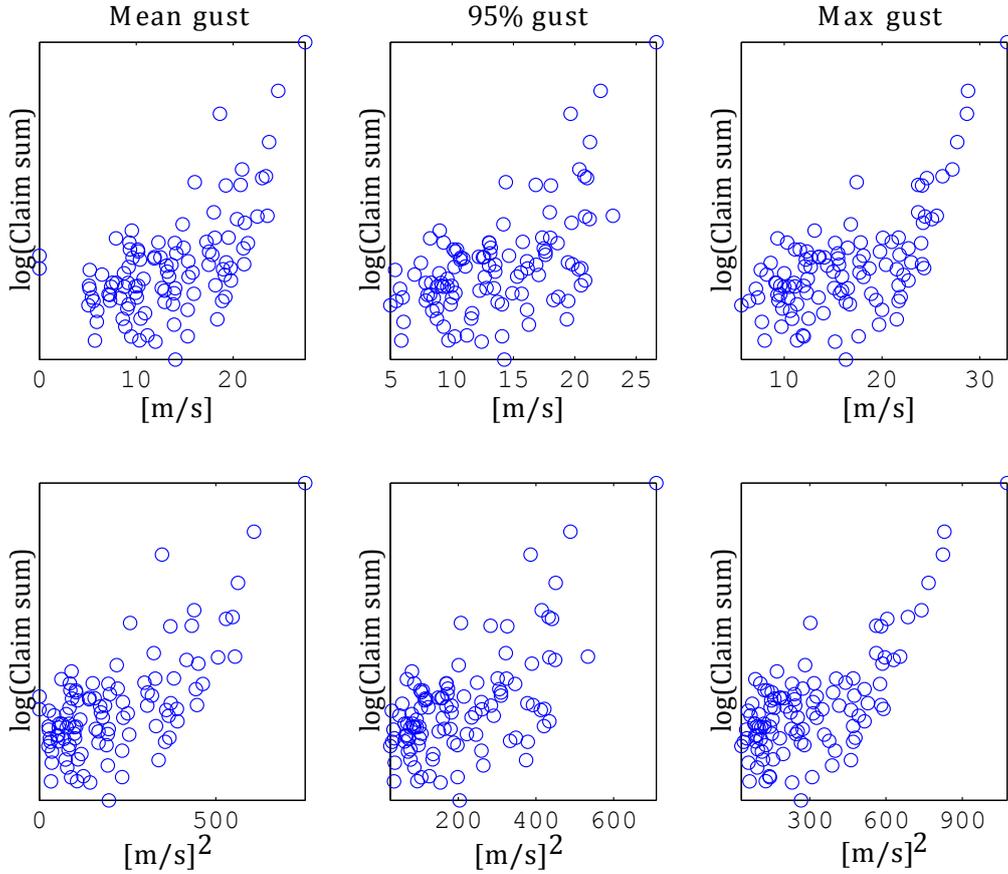
Figure 4.10: Scatter plots for claims sizes for storms over threshold after 1996 against the different wind indices and squared wind indices.

$\chi_1^2 \approx 3.84$. When the models that used $Temp_I$ and $Precip_I$ were compared to the null model the difference in estimated parameters was two and their deviance statistics were therefore compared to $\chi_2^2 \approx 5.99$.

The low deviance value for the time model implied that there was no trend in the claim sizes. Both precipitation models also had low deviance values meaning that they were disqualified for inclusion in the composite model. The $Temp_I$ model was better but a deviance value of 5,43 was still too low in comparison to the χ_2^2 which is approximately 5,99 and was therefore not included. All wind models, including the gust wind count models, proved to explain the storm damages significantly better than the null model. However, since strong correlation was expected between those covariates only one of them could be picked for the composite model. Since the model with the storm wind counts scored the highest log-likelihood value it was the only wind covariate that was included. Lastly, measures of temperature on the continuous scale was also included. Finally, we thus had a model were storm wind counts and temperature were included in the
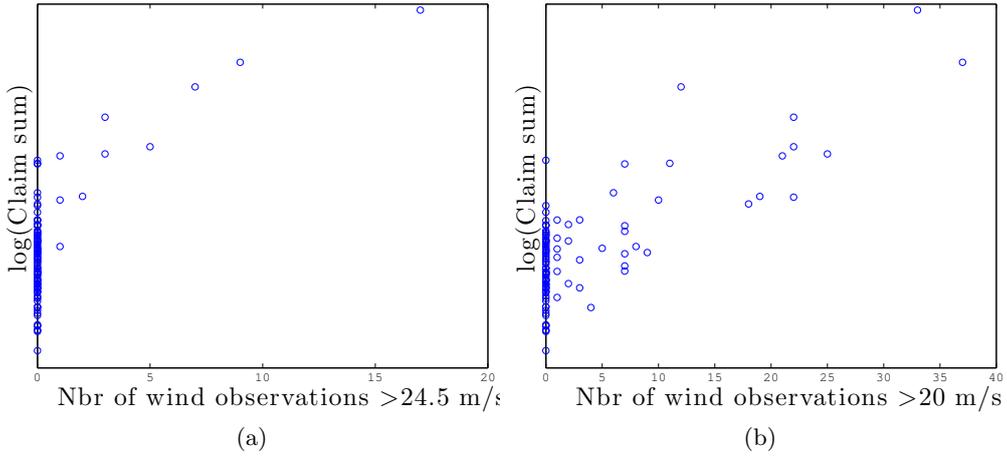
Figure 4.11: Scatter plots between logarithmic claim sizes for all storms and (a) $WindCnt_{24.5}$ and (b) $WindCnt_{20}$. For an explanation of these datasets see Section 3.1.3 on page 23
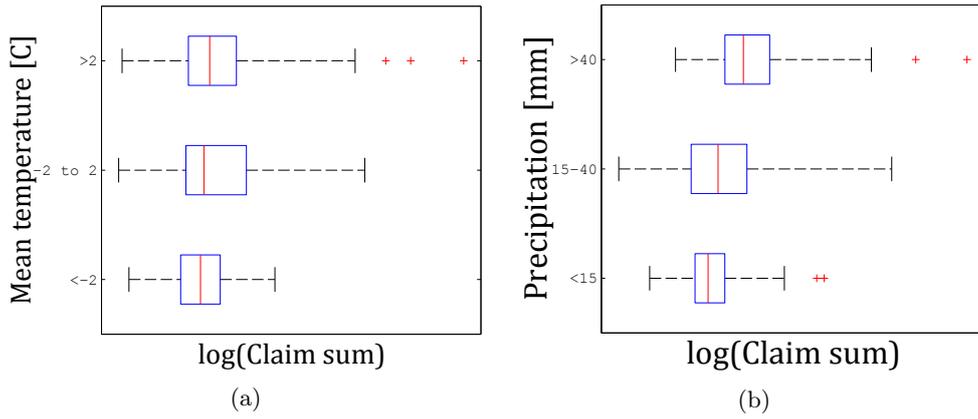
Figure 4.12: Box plots of the ground stability interval datasets (a) $Temp_I$ and (b) $Precip_I$ against the corresponding distribution of claim sizes for all storms. The central mark marks the mean of the claim sizes. The whiskers marks the highest and lowest value in the distribution not considered outliers by the algorithm, these outliers are shown as red crosses. The boxes show the 25 % and 75% quantiles.

scale parameter as

$$\log(\sigma) = \begin{bmatrix} 1 & WindCnt_{24.5} & Temp \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \end{bmatrix}$$
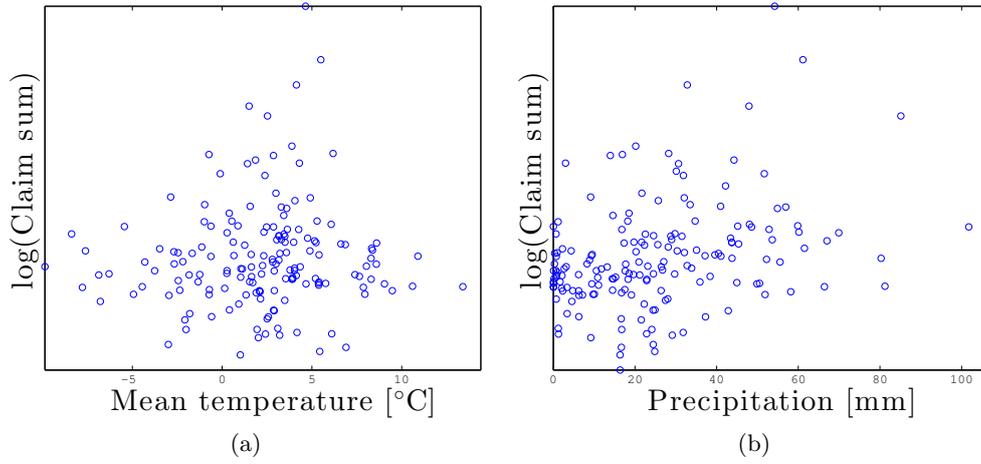
Figure 4.13: Plots showing correlation of the claim sizes against mean temperature (a) and precipitation (b) seven days before the claim date.
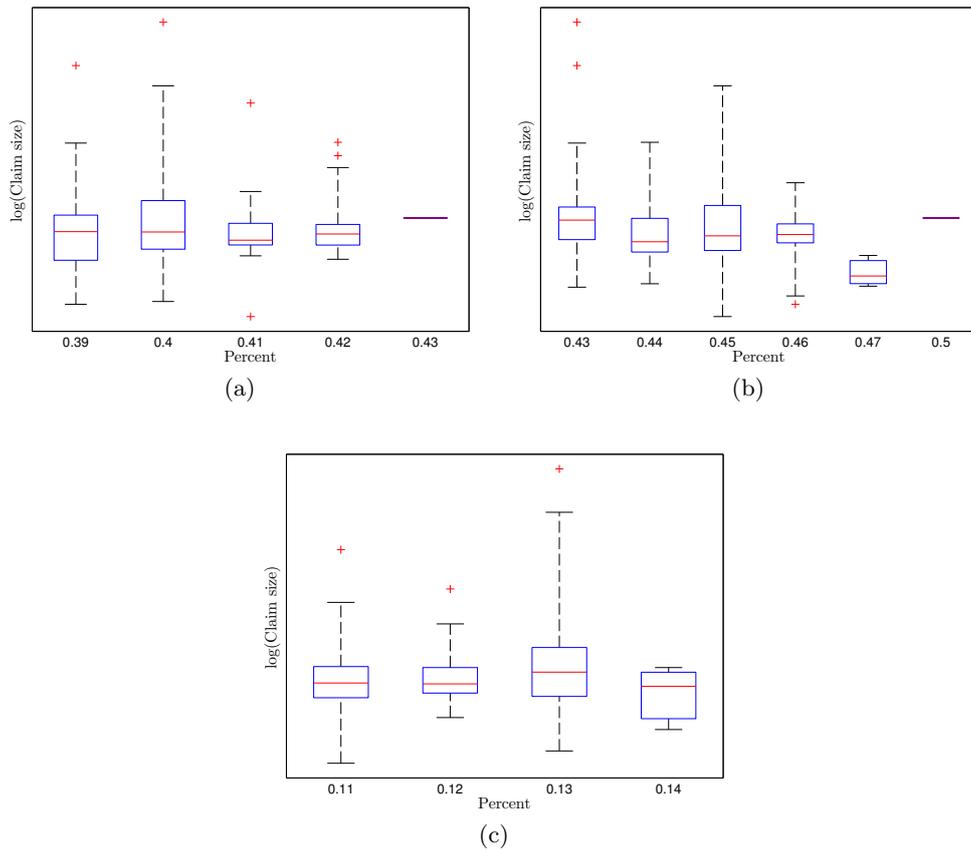
Figure 4.14: Box Plots of stand type datasets, (a) FS_{spruce}, (b) $FS_{conifer}$ and (c) $FS_{deciduous}$, percentage versus distributions of claim sizes for all storms
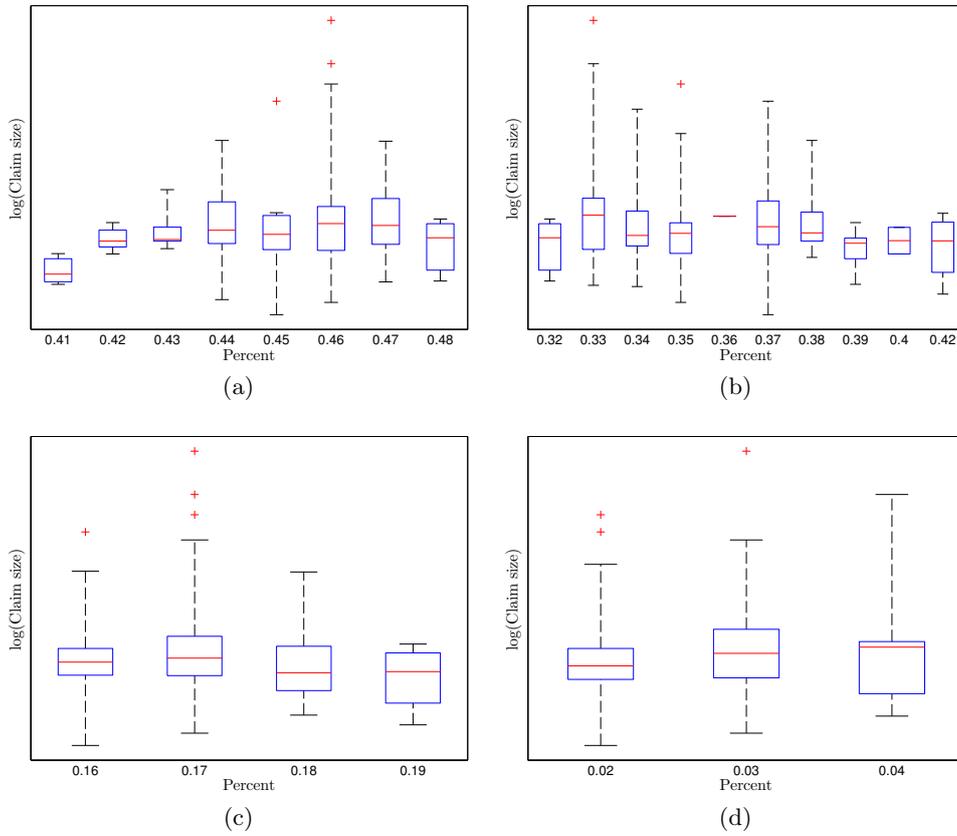
Figure 4.15: Box Plots of forest age group datasets, (a) FA_{0-40}, (b) FA_{40-80} and (c) FA_{80-120} and (d) $FA_{>120}$, percentage versus distributions of claim sizes for all storms.
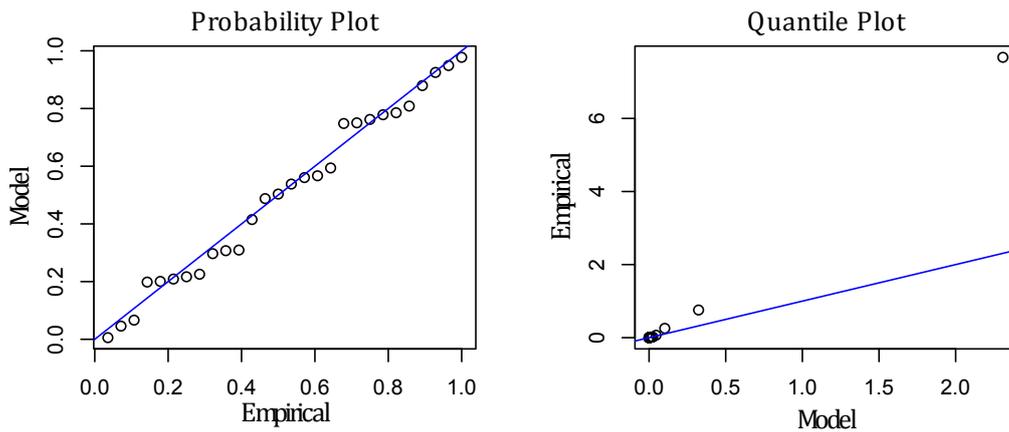
Figure 4.16: Goodness of fit diagnostics for a basic GPD model lacking any trend or covariates in any of its parameters.

Table 4.1: Single covariate GPD models and their respective deviance statistic compared to the null model. The value of df is the degrees of freedom which corresponds to the difference in number of estimated parameters between the actual model and the null model. The p-value corresponds to the probability that the difference in log-likelihood between the actual model and the null model is due to chance.

Covariate	Llh	D	df	p-value
$Time$	107.56	0.14	1	0.71
$Wind_{Mean}$	113.22	11.45	1	0.0007
$Wind^2_{Mean}$	114.23	13.47	1	0.0002
$Wind_{95}$	114.05	13.12	1	0.0003
$Wind^2_{95}$	115.15	15.31	1	< 0.0001
$Wind_{Max}$	116.73	18.48	1	< 0.0001
$Wind^2_{Max}$	119.71	24.45	1	< 0.0001
$WindCnt_{20}$	118.93	22.88	1	< 0.0001
$WindCnt_{24.5}$	120.52	26.06	1	< 0.0001
$Temp$	113.17	11.37	1	0.0008
$Temp_I$	110.20	5.43	2	0.066
$Precip$	107.72	0.45	1	0.50
$Precip_I$	109.42	3.86	2	0.14

This gave a final GPD model with parameters and log-likelihood estimated as

$$\hat{\beta}_0 = -7.72 \ (0.45)$$
$$\hat{\beta}_1 = 0.278 \ (0.05)$$
$$\hat{\beta}_2 = 0.343 \ (0.11)$$
$$\hat{\gamma} = 0.72 \ (0.37)$$
$$\text{Llh} = 125.2$$

Probability and quantile plots for this model can be seen in Figure 4.17. It is obvious that this model compared to the null model seen in Figure 4.16 was a much better fit for the higher quantiles while it still provided a good fit for the lower and more common storm damages.

4.3   Frequency Analysis

The distributions that were fitted to the annual number of storms and storms over threshold are summarised in Table 4.2[1]. In addition to this, plots showing

---

[1] The parameterisation of the negative binomial distribution in Table 4.2 differs from that presented in (2.18). μ corresponds to the mean of the distribution as defined in (2.19) and n is referred to as the dispersion parameter so that $p = n/(n + \mu)$ while the variance is calculated as $\mu + \mu^2/n$.
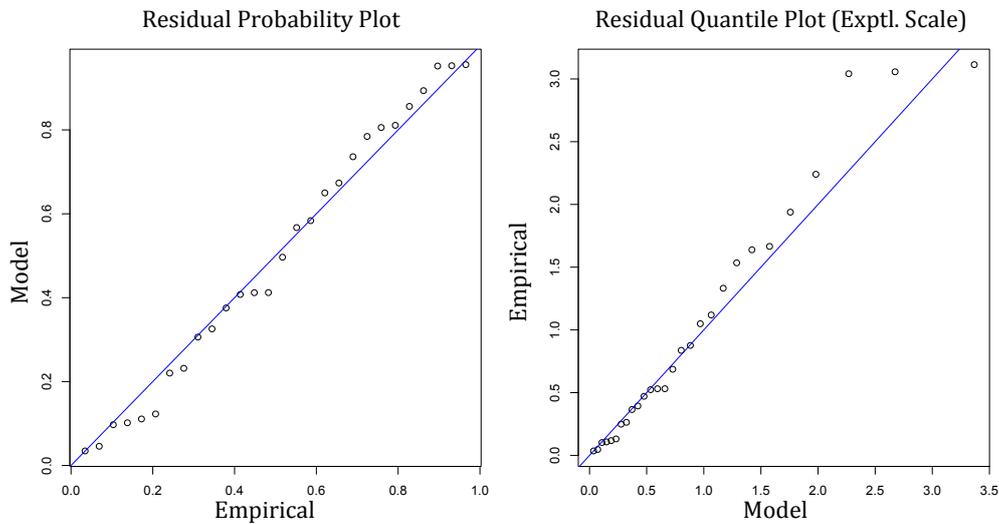
Figure 4.17: Goodness of fit diagnostics for the GPD model with counted storm gust winds and temperature in its scale parameter.

goodness of fit for respective model in terms of probability and distribution functions for the models and data can be found in Figures 4.18 to 4.21.

Starting with the complete set of storms the results gave no support for modelling with a homogeneous Poisson distribution as it was rejected on the 1% significance level by the χ^2 test. Considering that mean and variance is expected to be the same for this model the result is not surprising given that the mean for this dataset was estimated to 5,86 and the variance to 16.27. The negative binomial on the other hand showed more promise as the χ^2 test gave no support whatsoever for rejection of the distribution as a plausible model for the data. The corresponding estimated variance of the fitted negative binomial was 18.58 which agrees reasonably well with the variance of the dataset. Looking at the graphical goodness of fit for these models in Figures 4.18 and 4.19 it is clear that the negative binomial in comparison to the Poisson fits the data better although both models lacks accuracy in the upper tail. The fitted Poisson distribution is also looking likely to overestimate probabilities for years with four to eight storms.

As for modelling of the number of storms causing damages over the threshold the results were similar to those of all storms. The homogeneous Poisson model was again rejected on the 1% significance level by the χ^2 test albeit it fitted the data somewhat better in this case. The estimated mean and variance for this dataset was 1.41 and 2.61, which further lowered the support for the Poisson model. The support for modelling with a negative binomial was considerably lower in this case although it could still not be rejected as an appropriate model. The variance of this model was estimated to 3,05 which was higher than that of the dataset but still closer than the variance of the Poisson model. An inspection of

44

Table 4.2: Estimated parameters ($\hat{\boldsymbol{\theta}}$) and their standard errors in parenthesis together with goodness of fit for Poisson and negative binomial distributions. The variable X corresponds to the number of annual storms and includes all storms. The variable Y corresponds to the annual number of storms causing damages over the threshold only. The p-value corresponds to the probability that the data could come from the respective distribution.

Distribution	$\hat{\boldsymbol{\theta}}$	χ_k^2	p-value
$X \sim \mathrm{Po}(\lambda)$	5,86 (0,45)	39,82	< 0,0001
$X \sim \mathrm{NegBin}(\mu, n)$	5,86 (0,80)	1,42	0,70
	2,70 (1,12)		
$Y \sim \mathrm{Po}(\lambda)$	1,41 (0,22)	14,33	0,0007
$Y \sim \mathrm{NegBin}(\mu, n)$	1,41 (0,33)	1,53	0,22
	1,21 (0,74)		

the plots in Figure 4.20 shows that the Poisson model seems to overestimate the probability of years with number of storm damages over threshold close to the mean. The data however suggests that the number of such damages is associated with a higher degree of variance which is better reflected by the negative binomial in Figure 4.21.

The results from the Poisson distributions with time dependent means can be found in Table 4.3. The introduction of a linear time dependence according to (2.25) in the respective Poisson parameter yields significantly better models for both datasets than their time homogeneous counterparts. For the complete dataset the estimated parameter $\hat{\beta}_1$ gives the expected yearly increase in number of storms as 0.022 per year while the corresponding increase for storms over threshold is 0.037 per year. For further discussion of these results see Section 5.4.

Table 4.3: Trend models for annual number of storms and storms over threshold. The p-value corresponds to the probability that the difference in log-likelihood between the respective model and the corresponding time homogeneous model is due to chance.

Dataset	$\hat{\beta}_0$	$\hat{\beta}_1$	D	p-value
All storms	1,75 (0,078)	0,022 (0,0093)	5,79	0,016
Storms over threshold	0,30 (0,16)	0,037 (0,019)	3,88	0,049
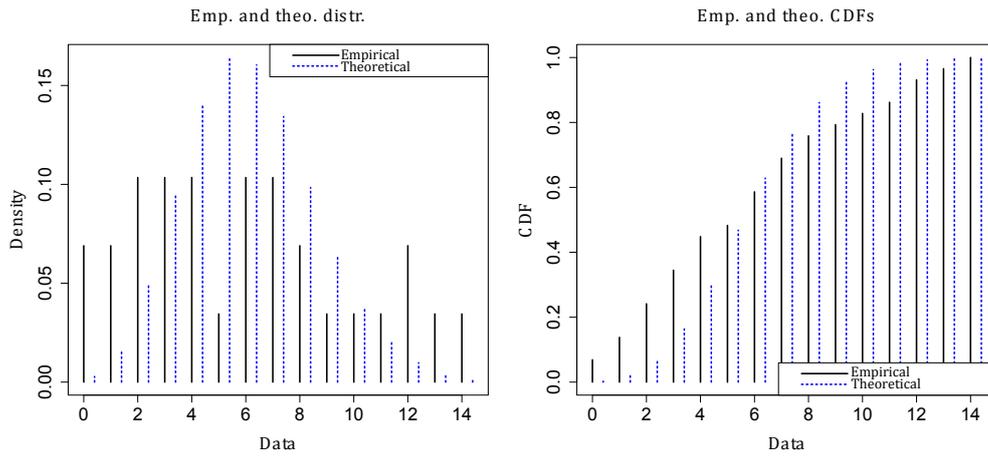
Figure 4.18: Poisson fit for annual number of storms.
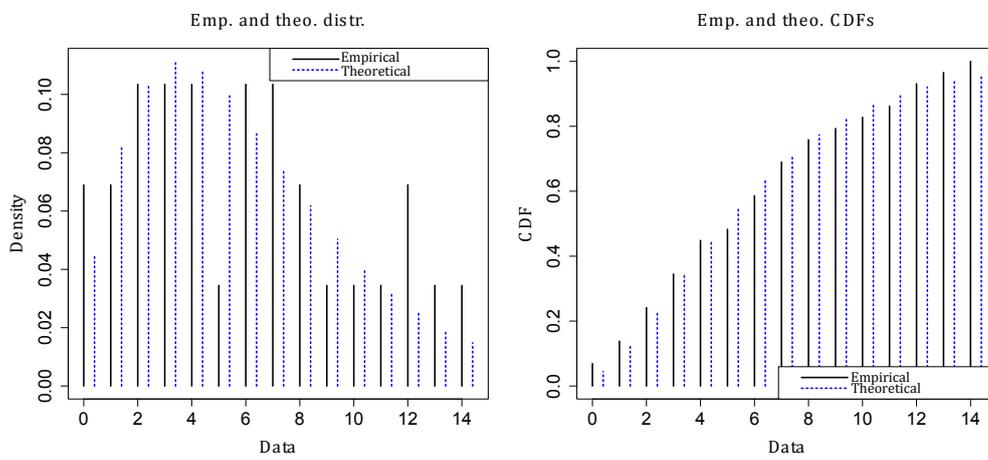
Figure 4.19: Negative binomial fit for annual number of storms.
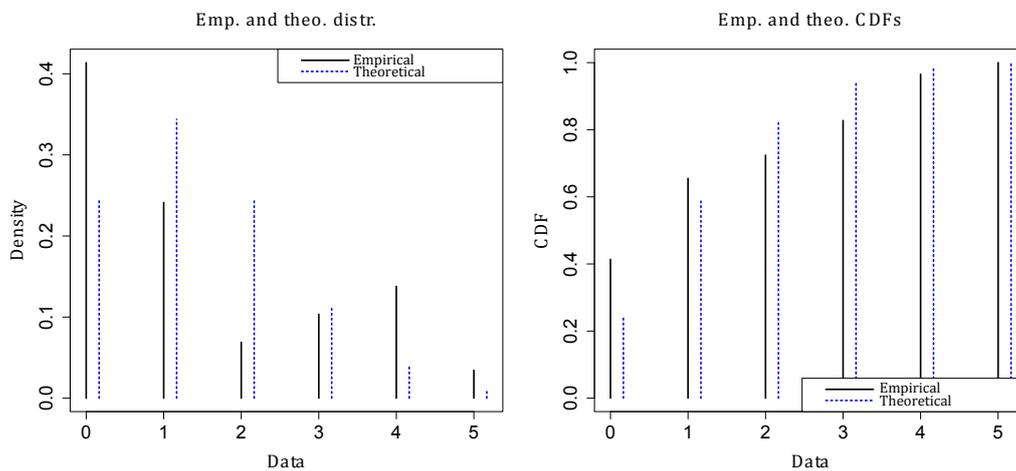
Figure 4.20: Poisson fit for annual number of storm damages over the threshold
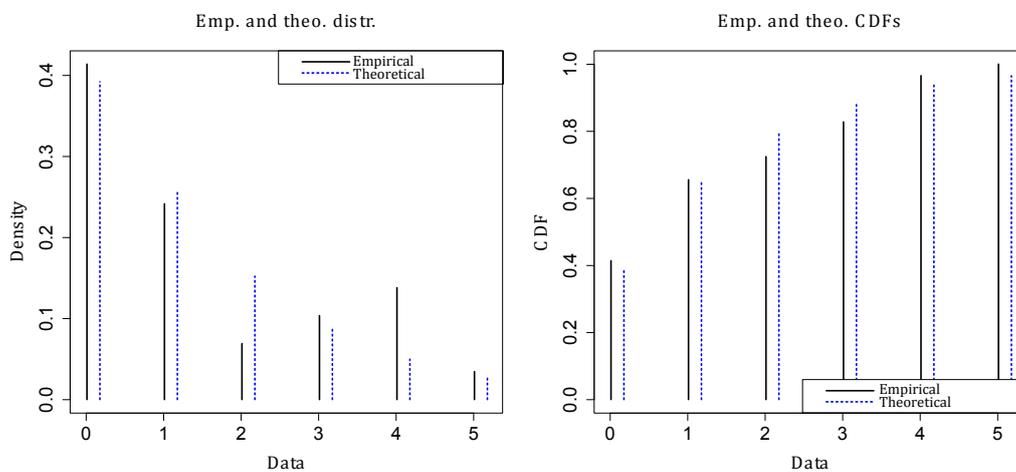
Figure 4.21: Negative binomial fit for annual number of storm damages over the threshold

Chapter 5

Discussion

The following sections are dedicated to an in-depth analysis of the results achieved in Section 4 and for comments and thoughts concerning the work on this thesis and its relation to other work on the same topic.

5.1 Claim data and storm definition

The fact that this thesis is more or less based on a definition of certain events derived from large payments made by LFAB due to forest damages puts the definition itself in the spotlight for a thorough review to be able to interpret the results in an appropriate way. The definition was used to identify and isolate insurance claims that could be expected to have a common cause and together summed up to substantial payments for LFAB. The only factor separating two events according to this definition was time, thus leaving the geographical aspect aside. From this premise it followed that two events could only be considered independent from each other if they were sufficiently separated in time. This however also introduced an error source in the form of claims with a possibly incorrect date stamp. In most cases such errors were probably not included after the introduction of a threshold but some events in the storm damage set could perhaps be false in the sense that they really should have been a part of another event. An example of this can be found in Figure 4.1 where the storm damage with number 21, corresponding to the eighth worst storm in terms of claims since 1996, is coupled with winds that in this context seems very modest. It is possible that this is just a matter of damages that can't be explained by high gust winds but there are a few facts that suggests that this is not the case. Firstly, the data from LFAB only contained payments due to forest damages in the form of fallen trees. Damages of this significance would therefore had to be explained by some other major factor that is not wind which in this case is hard to imagine. Secondly, we have not been able to find any kind of reports or paper articles that mentions those damages. This is strange considering the relatively large payments that LFAB had to make to cover the damages. To

further build on the suspicion of possibly wrong dated storm damages, the claims making up this storm were registered exactly one month after the notorious storm Gudrun in 2005. Similar patterns can be found for some of the other storms in the same figure and our contact at LFAB has confirmed that discrepancies between the registered date of damage and the actual date of damage probably exists to some extent as a result of human factors in the process of finding, reporting and registering the damage. It was however impossible to identify those claims in a certain way and they were therefore accepted as correctly registered. The discrepancies between the registered date of damage and the actual date of damage could also have some impact in the sense that damages with a common cause has been registered in a time span of maybe four or five days instead of one or two. Such a source of errors would mean more and less severe storm damages thus resulting in a biased dataset in the end.

The decision to not include a geographical aspect in the definition of separate events was primarily based on the fact that all payments made by LFAB was due to forest damages categorised as solely storm damages. This meant that damages caused by fire, bugs or diseases for example was not included in the data from LFAB and that the claims corresponded to the value of trees that had fell over or had stem brakes. The substantial storm damages would therefore require strong winds and as mentioned in Section 3.1.2 these are often the result of weather systems that affects large geographic areas. However, there may still be substantial local differences in wind and the other factors that could also play a role for the susceptibility of a tree to fall over. These concerns call for an analysis on smaller geographical scale. The major problem of doing this, for example on a province basis, was that of claim data availability. Using our storm definition, there wasn't enough storms for each province to do a meaningful statistical analysis using our method, see Section 5.3 below. Further, if the analysis is made on a province scale, the insured forest area for each province would need to be taken into account which of course could be done, but more data would be needed than what was included in the original dataset.

5.2   Covariate analysis

5.2.1   Wind

Considering that the payments made by LFAB covered the value of fallen trees it was hardly surprising that different measures of wind, among other covariates, were most correlated to storm damages. However, we wanted to see which measure or, in other words, which winds that were best correlated to the extreme damages. If the correlation plots in Figure 4.1 are studied some differences in correlation between the measures and damages can be seen. The mean and the 95% quantile measures are seemingly more scattered around the diagonal than the maximum gust measure and are for some of the highest records of wind and storm damages giving low support for correlation to each other. For example,

the third worst storm damage is coupled with a relatively low mean gust wind and inversely is the second highest recorded 95% quantile gust wind coupled with a storm damage that, in this context, must be considered as mediocre. These two are the most extreme cases of non-correlation and the second case could even perhaps be a false storm as described in 5.1. There are however several other points that also greatly deviate from the diagonal and highlights the shortcomings of these measures as predictors of storm damage. Maximum gust wind seems to correlate better with storm damages and it looks like the correlation gets stronger for gust winds exceeding speeds around 25 m/s and 625 $(m/s)^2$ in the respective plot. Figure 4.3(a) supports this as it tells that gust winds over 24,5 m/s were recorded during the seven worst storm damages and in comparison to Figure 4.3(b) it is showing more consistency in correlation to the worst damages. The storm wind count measure however contains more information than its maximum gust wind counterpart as it also says something about the length and/or the spatial extent of storm gusts during a period when substantial damages has been observed. Figure 4.4 shows this in greater detail as it presents how many measures of storm gusts that each station recorded during meteorological storms between 1996 and 2011. In most of these cases it was only one or two stations that recorded such winds and during no storm did all stations report of storm gusts. It is thus obvious why the maximum mean gust wind measure, that also was expected to carry some information about the spatial extent of a storm weather, gave comparatively bad correlation as this measure cut off the most important parts when all stations were averaged. With this in mind it also explains the relatively bad correlation of the 95% quantile measure as this measure became biased when all records during a storm were merged together. As a final remark concerning the most extreme winds it should be noted that the wind data delivered by SMHI in some cases was missing records of gust winds for some stations that would have been of interest. As a possible consequence of some of the worst storms those stations have temporarily been shut down which could mean that the number of storm gusts recorded during some storms, and possibly maximum gust winds, are underestimates of the true values. We do however believe that it is unlikely that this would have any significant impact on the analysis as a whole.

As for storm damages below the top seven and especially below the threshold, it seems that gust winds to a lesser extent decides the size of damages. All correlation plots in Figure 4.10 are characterised by a clump of storm damages that show great variability in their correlation with gust winds. This could imply that these damages are more sensitive to other factors that on their own or in combination with lower gust winds or other factors are able to cause damage. It is also well worth noting that only one storm damage below the threshold is associated with gust winds over 24,5 m/s which strengthens the picture of strong gust winds being exclusively connected to the worst damages.

5.2.2   Soil stability

The correlation between the parameters used for describing soil stability and storm damages were not expected to be as strong as for the different indices of gust winds. This was mostly due to the fact that soil stability had to be considered as a secondary factor to gust winds to cause trees to fall over. Strong gust winds would be needed for that to happen but the soil stability was expected to influence how strong those gust winds would have to be. Starting with the temperature's contribution to soil stability it was expected that ground frost would serve as extra support for the root system of a tree. If Figure 4.6(a) is studied it is easy to see that the vast majority of the damages over the threshold are indeed coupled with temperatures over the freezing point for which the soil stabilising effect of ground frost can be expected to be low. It is however also obvious that even though temperature may have some impact on the resulting damage it still plays a secondary role in this context as some of the highest temperatures are associated with low damages. The box plot in Figure 4.5(a) also shows that although the worst storm damages occurred during relatively high temperatures there is still a great variability within each of the temperature categories. The boxes are however somewhat shifted towards higher damages for increasing temperatures which indicates that possible correlation cannot be ruled out.

Precipitation was the other parameter that was expected to play an important role for soil stability. As can be seen in Figure 4.6(b) the amount of precipitation during and the days before a storm was however not very correlated to storm damages. Figure 4.5(b) tells a similar story as the corresponding temperature plot but even though the worst damages are associated with higher levels of precipitation the boxes are not giving further support for correlation. The precipitation data did not indicate wether the precipitation had come in the form of rain or snow and this could have had some influence on the resulting plots.

5.2.3   Forestry

Our results couldn't show any correlation between storm damages and forest stand type or age even though this was suggested in the literature. It's possible there still is a correlation between forest age and/or stand type, but with data as approximative as ours none was found. The approximate nature of our data has several explanations. The data gathered from SLU is approximative to begin with as it is based on spot checks that they use to create yearly data. This yearly data however is a three year gliding mean calculated from those spot checks. In our analysis we continue to make the data even more approximative by calculating a mean of the SLU data to go from province spatial scale to the whole area of Götaland. All these steps decrease the variance in the data to a level that we believe is too low to differ the forest age and/or stand type on an meaningful level for the different storm events.

5.3 GPD fitting analysis

The storm damages data that we wanted to model with a GPD had a very high variance that made it difficult to model without taking any covariates into account. This was obvious when the goodness of fit in Figure 4.16 on page 42 was studied. The damages that were caused during the storms Gudrun and Per above all, was according to that model unexpected as the deviation of those damages compared to the rest were on a scale that was very unlikely. It however turned out that factors coupled to wind and temperature to some extent could explain the extent of these damages. A comparison of the log-likelihood of the null model and the model using covariates showed that wind and temperature explained about 33% of the variance that could not be explained by the null model but it still left a majority of the variance unexplained. A part of this variance could perhaps be explained if the spatial resolution could be enhanced so that for example only a single province is studied. This would mean more exact estimation of values for different covariates for each damage which could prove to explain more of the variation in data. We did try to do this by only studying the provinces Kronoberg and Jönköping but it turned out that the data was not sufficient as less than ten storm damages remained after the threshold was applied. Longer data series for claims and covariates would thus be needed for this to be possible and it would also be interesting in the sense that it would give a better view of long term trends in storm damages. An alternative way to improve the spatial resolution could perhaps be to give more importance to the composition of claims inside each storm damage. This would mean that all claims that constitute a storm damage would have to be grouped and summed for each province to see their respective contribution to the damage. The covariate values could then be calculated so that they are weighted according to the contribution of each province to give better estimates of their values when the damage occurred.

Another way to explain more variance could be to reconsider how some of the covariates are quantified. The soil stability is an example of this as it perhaps could have explained more if it was represented in another way. We did this by using individual measures of temperature and precipitation but it seems that their values are not sufficient to give a prediction of soil stability. This could be because inadequate methods of measuring those quantities were used but it is perhaps more likely that soil stability is indeed the result of a complex process that requires the knowledge of more parameters than temperature and precipitation, and furthermore a proper expression that gives a single measure of this quantity.

The final GPD model used measures of wind and temperature as parameters in the scale parameter. Both parameters were estimated as positive values which implied that higher values of wind and temperature yields a higher probability for large scale damages. The shape parameter was estimated to a value less than one although the standard error suggested that a value more than one could not be ruled out. The resulting GPD was however heavy tailed as the value of the shape parameter indicated which was expected and in line with similar studies

to this such as Rootzén & Tajvidi (1997). A value less than one means that the distribution has a finite value of the mean and variance, whereas these properties are infinite for higher values. Since the final GPD model is a function of other parameters it is not possible to directly tell what to expect from future storm damages. By studying extensive historical data of wind and temperature it is however possible to create good statistical models of these parameters. Monte Carlo simulation then offers a possibility to generate values from our GPD model where the covariate values are randomly picked from the respective distribution for wind and temperature. Numerical calculation can then be used to obtain probabilities for different storm damages and return levels.

5.4   Frequency analysis

The results from the frequency analysis gave no support to the industry standard of assuming that storm damages occurs according to a homogeneous Poisson process. Instead, our analysis suggested that a negative binomial distribution could be a better model of the real process. The underlying reason for this was the great deal of variance of yearly storm damage counts that could not be represented by the Poisson distribution. The source for this variance could perhaps partly be deduced from clustering of events of this type as suggested by Mailier et al (2006), implying that some years or periods of a couple of years will bring several severe storm damages while other pass unnoticed in these terms. It is however important to remember that the events that were counted were defined according to our definition and that similar sources of errors that have been mentioned in Section 5.1 also could play a role in explaining this variance.

The frequency analysis also included a test for a possible trend in the mean of the Poisson process to see if any evidence for a increasing number of yearly storm damages could be found. The results indicated that such a trend was likely both in the case of damages in common and for damages over the threshold. We do however believe that it is not necessarily so that this is an actual trend of a growing number of annual storms but perhaps a question of increasing variance in this number, especially for recent years and perhaps due to clustering.

The theory on clustering would be interesting to examine if longer records of claim data were available. We could in our data see a pattern of more annual storm damages during the 2000s than during earlier periods. The variance between different years however also varied a lot as some years during this decade brought no storm damages at all. It could therefore be interesting to see if there are other periods in time that show a similar behaviour and in that case identify common conditions that could explain the clustering effect and the period between such periods. A long term decreasing wind climate has been observed in Sweden since the early decades of the 20th century (Wern & Bärring, 2009) but it is still possible that shorter periods in the size of a couple of years with windier climate could occur occasionally and explain a part of the variance.

5.5 Topics for further research

During the work on this thesis some areas that show promise for further research were not included due to time and scope issues. One of these is a direct continuation of this thesis, while the others are suggestions on how to do similar studies more thoroughly or with a different approach.

The most obvious continuation of this thesis would be to do the Monte Carlo simulation as suggested in Section 5.3. By using the model for storm frequency to simulate the annual number of storm damages over the threshold and then simulate the extent of each storm's damage with values from the GPD model, it would be possible to generate values of annual maximum storm damages for an arbitrary choice of years. A generalized extreme value distribution (GEV) could then be fitted to the maxima so that a prediction of yearly storm damage return levels with confidence intervals would be possible.

Other suggestions follows:

- Finding other significant covariates. Wind direction, soil acidification and nitrogen levels and topography are all possible covariates that weren't studied in this thesis. The wind direction was excluded due to the seemingly stable wind climate in Götaland as a whole, with a big majority of high wind speeds coming from the west. Soil acidification and nitrogen levels as well as topography were excluded due to not having any proper data available. Wind direction could possibly have more variation locally making it an interesting covariate on a smaller regional scale or in a region with more variable wind directions. Soil acidification and nitrogen levels could possibly be combined with precipitation and temperature data to make a better index for soil stability as discussed in Section 5.2.2.

- Doing the analysis on a different region, in Sweden or elsewhere to see if the results are consistent.

- Doing the analysis on a smaller region or weighting covariates after regional fraction of claim size per storm as discussed in Section 5.3. It would be possible to increase the spatial resolution of the analysis with more claim data which perhaps could result in a more accurate model. This could be possible by combining different insurance companies' data and/or using data that stretches further back in time. Also, every day the insurance databases keep growing.

Bibliography

Alexandersson H., Tuomenvirta H., Schmith T. and Iden K., 2000.
  Trends of storms in NW Europe derived from an updated pressure data set,
  Climate Research, 2000:14, pp. 71-73

Bengtsson A. and Nilsson C., 2007.
  Extreme value modelling of storm damage in Swedish forests, *Natural Hazards
  and Earth Systems Sciences*, 2007:7, pp. 515–521

Blennow K., Andersson M., Bergh J., Sallnäs O. and Olofsson E., 2010.
  Potential climate change impacts on the probability of wind damage in a south
  Swedish forest, *Climatic Change*, 99:2010, pp. 261–278

Blom G., Enger J., Englund G., Grandell J. and Holst L., 2009.
  Sannolikhetsteori och statistikteori med tillämpningar, Femte upplagan, Lund:
  Studentlitteratur

Bonazzi A., Cusack S., Mitas C., and Jewson S., 2012
  The spatial structure of European wind storms as characterized by bivariate
  extreme-value copulas, *Natural Hazards and Earth Systems Sciences*, 2012:12,
  pp. 1769–1782

Braun S., Schindler C., Volz R. and Flückiger W., 2003.
  Forest damages by the storm 'Lothar' in permanent observation plots in
  Switzerland: The Significance of soil acidification and nitrogen deposition,
  Water, Air, and Soil Pollution, 2003:142, pp. 327–340

Brodin E. and Rootzén H., 2009.
  Univariate and bivariate GPD methods for predicting extreme wind storm
  losses, *Insurance: Mathematics and Economics*, 44:3, pp. 345–356

Coles S., 2001.
  An Introduction to Statistical Modelling of Extreme Values, London: Springer
  Verlag

Della-Marta P.M., Mathis H., Frei C., Liniger M.A., Appenzeller C., 2007.
  Extreme wind storms over Europe: Statistical Analyses of ERA-40, *Arbeits-
  berichte der MeteoSchweiz*, 216, 79 pp.

Delignette-Muller, M. L., Pouillot, R., Denis J.-B., and Dutang C., 2012.

fitdistrplus: help to fit of a parametric distribution to non-censored or censored data

European Forest Institute, 2010.
Destructive Storms in European Forests: Past and Forthcoming Impacts, *European Forest Institute*, [Accessed online January 23, 2013], Available at: `http://www.efiatlantic.efi.int/files/attachments/efiatlantic/2010-storm/storms_final_report_main_text.pdf`

Gilleland E. and Katz R.W., 2004.
Extremes toolkit (extRemes): weather and climate applications of extreme value statistics, *National Center for Atmospheric Research (NCAR)*

Klugman S.A., Panjer H.H., Wilmot G.E., 2004.
Loss models, from data to decision, Second edition, New York: Wiley-Interscience

Leadbetter M.R., Lindgren G. and Rootzén H., 1983.
Extremes and Related Properties of Random Sequences and Series, New York: Springer Verlag

The MathWorks, 2012.
MATLAB and Statistics Toolbox Release 2012b, Inc., Natick, Massachusetts, United States. `http://www.mathworks.se/products/matlab/`

Mailier J.P., Stephenson B.D. and Ferro A.T.C., 2006.
Serial Clustering of Extratropical Cyclones, *Monthly Weather Review*, 2006:134, pp. 2224–2240

Nelder, J.A., and Wedderburn, R.W., 1972.
Generalized Linear Models, *Journal of the Royal Statistical Society. Series A (General)*, Vol. 135, 1972:3, pp. 370-384

Nilsson C., Stjernquist I., Bärring L., Schlyter P., Jonsson A-M. and Samuelsson H., 2001.
Recorded storm damage in Swedish forests 1901–2000, *Forest Ecology and Management*, 199:2004, pp. 165–173

R Core Team.
R: A language and environment for statistical computing. *R Foundation for Statistical Computing, Vienna, Austria* URL `http://www.R-project.org/`.

Rootzén H. and Tajvidi N., 1997.
Extreme value statistics and wind storm losses: A case study, *Scandinavian Actuarial Journal*, 1997:1, pp. 70-94

Rootzén H. and Tajvidi N., 2001.
Can Losses Caused by Wind Storms be Predicted from Meteorological Obeservations?, *Scandinavian Actuarial Journal*, 2001:2, pp. 162–175

Schelhaas M-J., Nabuurs G-J. and Schuck A., 2003.

Natural disturbances in the European forests in the 19th and 20th centuries, *Global Change Biology*, 2003:9, pp. 1620–1633

Schmoeckel, J. and Kottmeier, C., 2008.
Storm damage in the Black Forest caused by the winter storm "Lothar" – Part 1: Airborne damage assessment, *Natural Hazards and Earth Systems Sciences*, 2008:8, pp. 795–803

SMHI, 2011.
Gudrun - Januaristormen 2005,
Swedish Meteorological and Hydrological Institute, [Accessed online March 13, 2013], available at: `http://www.smhi.se/nyhetsarkiv/adventsstormen-jamford-med-stormarna-gudrun-och-per-1.18870`

Swiss Re, 2000.
Storm over Europe - An underestimated risk, Zürich: Swiss Re.

Usbeck T., Wohlhemuth T., Dobbertin M., Pfister C., Bürgi A. and Rebetez M., 2010.
Increasing storm damage to forests in Switzerland from 1858 to 2007, *Agricultural and Forest Meteorology*, 150:2010, pp.47–55

Wasa Group, 1998.
Changing waves and storms in the northeast Atlantic?, *Bulletin Of The American Meteorological Society*, 79:5, pp. 741-760, Academic Search Complete, EBSCOhost, viewed January 31, 2013

Wern L. and Bärring L., 2009.
Sveriges vindklimat 1901-2008, Analys av förändring i geostrofisk vind, *Meteorologi*, 198:2009, [Accessed online January 9, 2013], available at: `http://www.smhi.se/polopoly_fs/1.7843!meteorologi_138.pdf`