THE MODELLING OF EXTREME EVENTS

By D. E. A. Sanders

[Presented to the Institute of Actuaries, 4 April 2005]

ABSTRACT

The modelling of extreme events is becoming of increased importance to actuaries. This paper outlines the various theories. It outlines the consistent theory underlying many of the differing approaches and gives examples of the analysis of models. A review of non-standard extreme events is given, and issues of public policy are outlined.

KEYWORDS

Extreme Events; Extreme Value Theory; Generalised Pareto Distribution; Power Law; Cat Models

CONTACT ADDRESS

David Sanders, Milliman Ltd, Finsbury Tower, 103-105 Bunhill Row, London EC1Y 8LZ, U.K. Tel: +44(0)20 7847 6186; Fax: +44(0)20 7847 6105; E-mail: david.sanders@milliman.com

1. Summary

1.1 This paper is intended to give a broad introductory overview of extreme events and the theory underlying their assessment. In it we explore how they can impact, and indeed have impacted, on actuarial work. This, naturally, concentrates on general insurance, for which extreme catastrophic events have clear impacts on reserving, pricing and solvency, but we also demonstrate that there are applications within life insurance, pensions and finance.

1.2 After a brief introduction, the paper goes into the mathematics of extreme events, and highlights the consistent assumptions between a number of apparently different techniques. The mathematical theory (EVT) has been well developed for quite a while. The first important theorem states that the maximum of a sequence of observations is approximately distributed as the generalised extreme value distribution. This distribution has three forms: the Gumbel, Frechet and Weibull. A second theorem states that the observations exceeding a high threshold, under very general conditions, are approximately distributed as the generalised Pareto distribution. The approach to extreme value analysis is then based on a Poisson process governing the rate of occurrence of exceedance of a high threshold; and the estimated generalised Pareto distribution for the excess over the threshold. Warning is also given to amalgamate geophysical and other knowledge in determining any model.

1.3 Some examples of extreme events are then given in Section 4. Indication is given, with examples, of the type of model used, in practise, to analyse these events. This analysis and description are necessarily limited, as there are considerable volumes written on each risk elsewhere. Section 5 deals with places where EVT models have been applied in practice. It should be noted that the examples follow a given pattern of analysis and assessment, and that the approach in Section 3 is followed in these examples.

1.4 Section 6 looks some unusual extreme events, and is intended to whet the appetite of the reader for more examples. Finally in Section 7, the issue of public policy and extreme events is outlined.

2. Introduction

2.1 When considering a series of insurance claims, it is usual to consider it as one distribution. This is because the philosophy of insurance is centred around a number of principles; in particular the principle of independence of claims, or, more specifically, a low correlation between claims. The 'law' of large numbers applies, resulting in an decrease in the coefficient of variation as the number of contracts increases, and enables insurers to sell risk management products at realistic prices. However, a claims distribution, in practice, may be considered in two possibly distinct parts:

— the central part of the distribution, which deals with the normal claims; and
— the extreme end of the distribution.

2.2 The insurer is also aware that a single event could result in multiple deaths, injuries and loss of property, for example a motorway pile up or a fire involving a number of premises. These dependent risks are met by traditional reinsurance, in particular excess insurance, where losses arising from one event exceed a specified amount. This is either at the risk level ('risk excess') or at an aggregate level ('catastrophe'). These types of reinsurance apply to life insurance, general insurance and pensions. These reinsurance contracts have specific features, namely a deductible and the maximum amount that might be paid in any one event. Many of these types of claims are analysed by an extension of the distribution of claims.

2.3 However, there are potential aggregate claims where reinsurance cover may be more difficult to be purchased, or may not exist, due to lack of capacity or unrealistic pricing. This paper deals with events which may be at this level. Examples of these types of claims include earthquakes, windstorms, terrorist acts, and so on. They will be of interest to actuaries reviewing capital requirements, risk retention, and so on. When reinsurance cover is no longer available, other financial devices, such as catastrophe bonds, are used to protect the capital of a financial institution. This paper concerns itself primarily in assessing the likelihood and cost of these extreme events.

2.4 Many catastrophic events are covered, or are capable of being covered, in the insurance and financial markets. For example, most earthquakes and windstorms should be met out of the insurance and reinsurance market, without significant fall out or insolvency in the players. However, some losses will occur which do have a detrimental effect on a number of players. These are generally termed extreme events.

2.5 Extreme events are relative to the entity; different players will have different views. A fire that burns down a person's home may be considered extreme to that person, but, as that risk is generally covered by traditional insurance, it is not extreme in the context of this paper. A hurricane in Florida may not measure up as a high value loss, but is an extreme event as far as the communities hit are concerned. However, it is not necessarily an extreme in the context of this paper. These events primarily fall into the central section of the claims distribution. Conversely, the collision with a large asteroid would have effects throughout the world. Extreme events can be seen from a number of points of view, but the tools to measure the potential and manage the risks are essentially the same.

2.6 In addition to insurance-based risks, insurance companies and pension funds also run financial risks. Sudden movements in the value of assets could adversely impact on the solvency of the insurer or pension fund, and also impact on policyholder and fund member benefits. (Re)insurance contracts, including financial products such as options and other hedging devices, will help offset some of the adverse movement, but market capacity will mean that many of the risks are uncovered. These risks may be analysed by extreme event models.

2.7 Ordinary companies are also exposed to operational risks, both at the firm level and at the whole economy level. There is the possibility of an extreme failure of computer systems, resulting in loss of critical information. Making a wrong decision or incorporating the wrong business model could readily lead to a company becoming bankrupt. These extreme operational risks have been the subject of investigations using EVT.

2.8 In certain instances the event may be small, but have a significant impact on a few local players. For example, the increasing level of local flooding and tornados in the United States Mid-West are viewed by many commentators as extreme events. These are not, in themselves, far reaching, but need to be considered by actuaries in their pricing role. At other times an event has far reaching effects. The aim of the risk manager is to ensure that his company survives such extreme events by managing his business in such a way as to predict the impact and diversify, if possible, such risk. This may not be possible, and could be the price which we have to pay for the concentration of undiversifiable risk. At the state and global level, the reliance on insurance alone to manage extreme events appears to be unsustainable, and alternative strategies will be needed.

3. Modelling of Extreme Events

3.1 There are many models which have been used to manage and understand extreme events. In practice, a combination of models is often needed. The type of models used in practice are:

— *statistical/actuarial models*, where past experience is used to estimate the consequence of future events;
— *physical models*, where, for example, the consequence of a landslide or submarine earthquake is estimated from a scaled down model in a laboratory; and
— *simulation or catastrophe models*, depending on computer simulations of events which include pre-determined parameters and physical constraints, for example in weather forecasts. Catastrophe models tend to make use of statistical and formulated physical models.

3.2 The structures of these model types may, themselves, be considered in three parts or processes, namely (and using hurricanes as a means of illustrating each):

— *diagnostic*, where, for example, post event hurricane loss assessment is made using current and historic observed data, with, perhaps, physical constraints such as topography, combined with extrapolating and interpolating the estimates for known locations to those where there have been no historic data;
— *investigative*, where, for instance, an explanation is formulated as to why hurricanes occur, and the relationship between hurricane intensity and, for example, ocean temperature, and the conditions in the western Sahara Desert; and
— *predictive*, which, for example, attempts to forecast the number of hurricanes that will occur in a season in a certain category or higher.

3.3 Extreme event models are very dependent on the parameterisation. Parameters can be derived from observed data; but, if an extreme event has not occurred within those data; then modelling, and hence prediction, might be difficult. Unlike normal statistical analysis, where outliers are ignored, such data are precisely what drives the extreme process. Indeed, it is the 'normal' data are ignored, as these may lead to mis-estimation of the key extreme event parameters. In a simulation type model, the key parameters may, themselves, vary stochastically.

3.4 Prediction of events becomes difficult, because the data are incomplete. How can you predict the one-in-200 years' storm if you have less than 25 years' data? Pure statistical models can suffer in this respect, and pure reliance on such models can lead to misdiagnosis. Data may also need adjusting to reflect seasonality and other factors, such as changes in values of property. Ideally, any model needs to be a combination of statistical and physical considerations.

Extreme Value Theory

3.5 Many authors have written on this subject. Set out below is a brief description of the basis underlying the mathematical analysis of extreme value theory. The key idea is that one can estimate extreme amounts and probabilities by fitting a 'model' to the empirical survival function of a set of data using only the extreme event data rather than all the data, thereby fitting the tail, and only the tail. This ensures that the estimation relies on the tail as opposed to data in the centre. It is not intended to set out full mathematical proofs or explanations in this paper. For more details, the reader should consult Embrechts *et al.* (1997), Coles (2001) and Reiss & Thomas (1997). A summary of software sources is given in Appendix 1.

3.6 One of the important results of classical statistics is the central limit theorem. This describes the characteristics of the sampling distribution of means of random samples from a given population. The central limit theorem consists of three statements:

— The mean of the sampling distribution of means is equal to the mean of the population from which the samples were drawn.
— The variance of the sampling distribution of means is equal to the variance of the population from which the samples were drawn, divided by the size of the samples.
— If the original population is distributed normally, the sampling distribution of means will also be normal. If the original population is not distributed normally, the sampling distribution of means will increasingly approximate a normal distribution as sample size increases (i.e. when increasingly large samples are drawn).

3.7 Consider the sum $S_n = X_1 + \ldots + X_n$. Then, the expected value of S_n is $n\mu$, and its standard deviation is $\sigma n^{1/2}$. Furthermore, informally speaking, the distribution of S_n approaches the normal distribution $N(n\mu, \sigma^2 n)$ as n approaches infinity. To avoid a degenerate limit, S_n is restandardised by rescalings, and setting:

$$Z_n = \frac{S_n - n\mu}{\sigma\sqrt{n}}.$$

The central limit theorem states that Z_n converges to N(0, 1).

3.8 Consider extreme values in a similar light. Instead of analysing the mean, we consider the behaviour of extremes, which may be considered by analysing the maximum order statistic:

$$M_n = \max\{X_1 ; X_2 ; : : : ; X_n\}.$$

Theoretically, assuming that the amounts are IIF, the probability of an event exceeding an amount z is:

$$\begin{aligned}\Pr(M_n < x) &= \Pr(X_1 < z ; \ldots ; X_n < z)\\ &= \Pr(X_1 < z) \ldots\ldots \Pr(X_n < z)\\ &= (\mathrm{F}(z))^n.\end{aligned}$$

The problem, in practice, is that the distribution function F is often unknown. Accordingly, it is possible to only give some broad bounds to the behaviour of M_n.

3.9 The problem of extreme events may be summarised in two questions. Firstly, what possible distributions can arise for the distribution of M_n as n tends to infinity? Secondly, is it possible to formulate these limit distributions into a single formulation, which is independent of F? If so, then we can estimate the distribution of M_n directly, without knowing F. These questions are identical to that solved by the central limit theorem for the mean of the distribution. A similar approach is applied to limits of the distribution of M_n. In this case, we consider the limiting distributions of:

$$\frac{M_n - b_n}{a_n}$$

where a_n and b_n are sequences of normalising coefficients, similar to those used in the central limit theorem. The solution to the range of possible limit distributions is given by the extreme types theorem. This first important theorem states that the maximum of a sequence of observations is approximately distributed as the generalised extreme value distribution. This distribution has three forms: the Gumbel, Frechet and Weibull. (See Fisher & Tippett, 1928; and Gumbel, 1958.)

The Generalised Extreme Value Distribution

3.10 The generalised extreme value (GEV) family describes the distribution of the maxima of the sets of observations. It is a pure statistical model. A typical application would be to try to estimate the largest annual claim Y under a given insurance contract (assuming the underlying risks and exposure are independent and uniform). The behaviour of the annual maximum Y (suitably scaled and translated by the normalising coefficients) is then approximately described by one of the GEV family of distributions, given by the cumulative distribution function:

$$P(Y < y) = \text{GEV}(y; \xi, \mu, \sigma) = \exp(-[1 + \xi(y - \mu)/\sigma]_+^{-1/\xi}).$$

For $\mu, \sigma > 0$ and $\{y : 1 + \xi(y - \mu)/\sigma > 0\}$, the notation $[y]_+$ denotes $\max(y, 0)$.

3.11 As ξ tends to 0, the GEV converges to the Gumbel distribution:

$$P(Y, y) = F(y) = \exp(-\exp(-(Y - \mu)/\sigma)).$$

Here the key parameter is ξ, the shape parameter, and μ, σ are location and scale parameters respectively. Set out in Figure 3.1 is a Gumbel distribution, generated from the Xtremes Software package (see Appendix 1), with a location parameter of 1,000 and scale parameter of 10,000. (Note that the Gumbel distribution has no shape parameter.) Xtremes has been used because it is freely available over the internet, whereas other packages, such as Evis (see Appendix 1) or Coles (2001), require S plus.

3.12 If $\xi < 0$, the distribution ('Weibull') has a finite upper bound, indicating an absolute maximum, so one might expect to fit such a distribution to, say, the ages of a human population indicating an upper bound to possible age. Set out in Figure 3.2 is the Weibull distribution, with location 1,000 and scale 10,000 for shapes minus five and minus two, again from Xtremes. (Note that the formulation in Xtremes is slightly different from the normal formulation.)

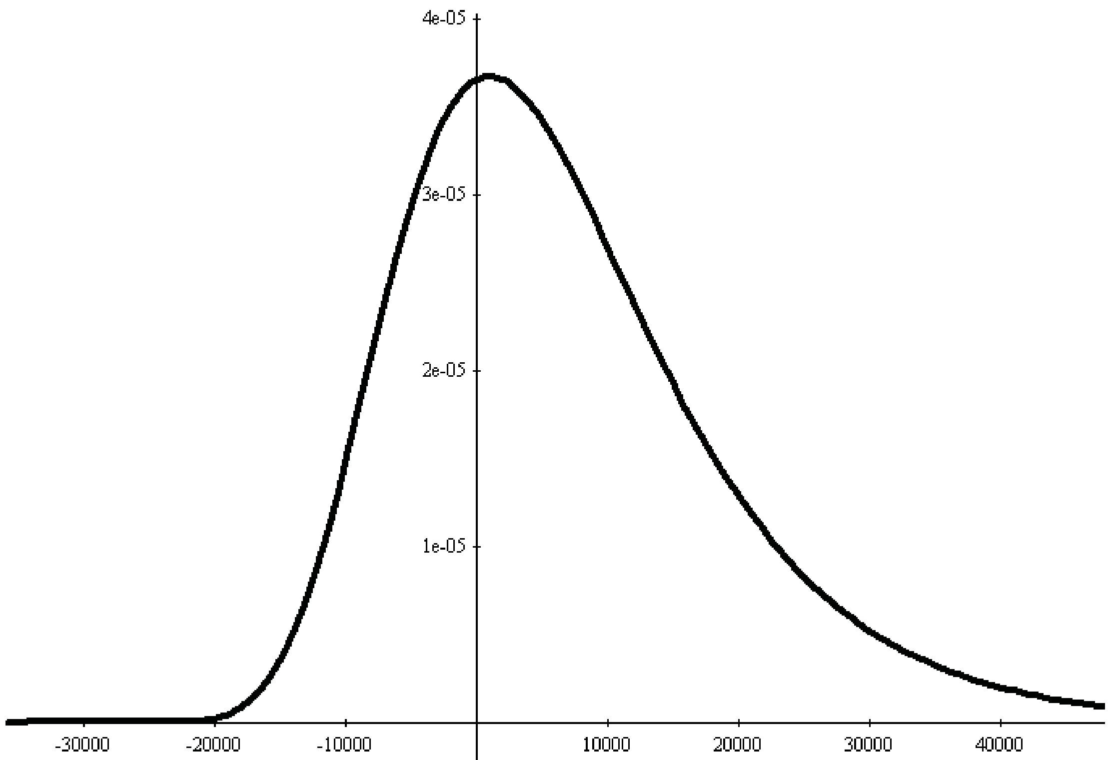

Figure 3.1. Gumbel Distribution

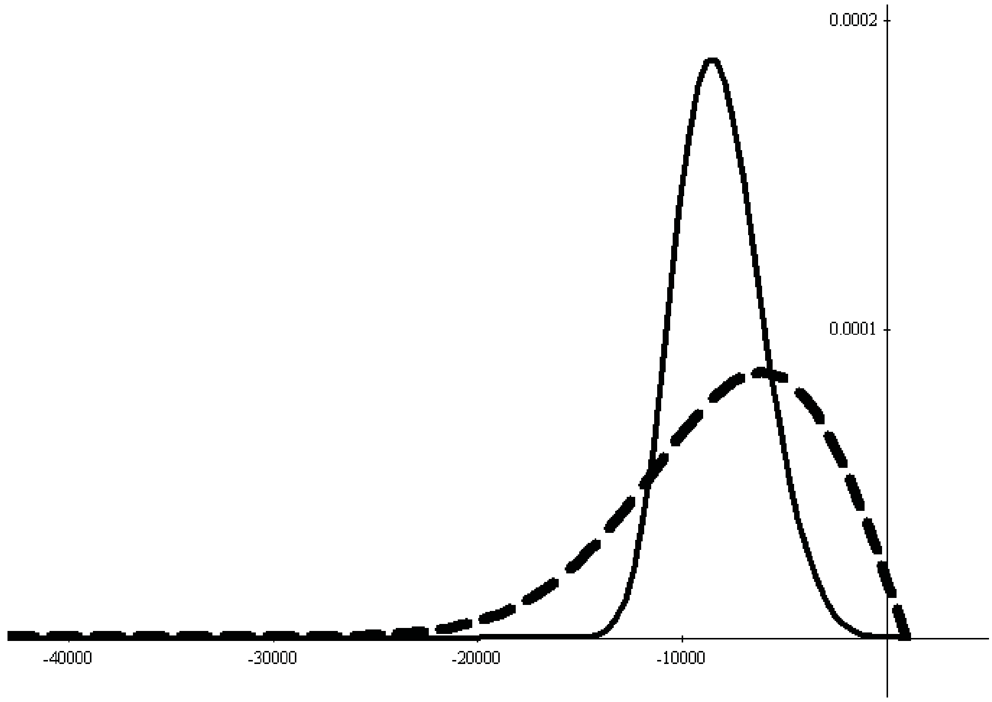

Figure 3.2. Weibull distribution

3.13 For $\xi > 0$, the tails of the distributions become heavier with infinite variance, and this is typically the subset of distributions ('Frechet' distributions) with which one would be concerned for extreme insurance events. The Frechet distribution is set out in Figure 3.3.

3.14 This result, due, in various parts, to Fisher, Tippett, von Mises and Gnedenko, is known as the 'extremal type theorem', classifying the extremal behaviour into one of the three families of distributions. See Embrechts *et al.* (1997), Coles (2001) and Reiss & Thomas (1997) for more detail.

Typical Insurance Application of GEV

3.15 A typical example would be to estimate the potential maximum annual losses in an insurance business for a particular type of event, or contract type, or even an individual contract. To apply the GEV, the data would be placed into an appropriate grouping, and then adjusted to allow for differing exposures, trends such as seasonality, and, most importantly, whether the observations was reasonably independent. If there were potentially dependent events (such as a series of windstorms hitting one location), they would need either to be amalgamated or to be removed from the series. There is considerable amount of software available to estimate the parameters of the GEV, using methods such as maximum likelihood or the method of moments.

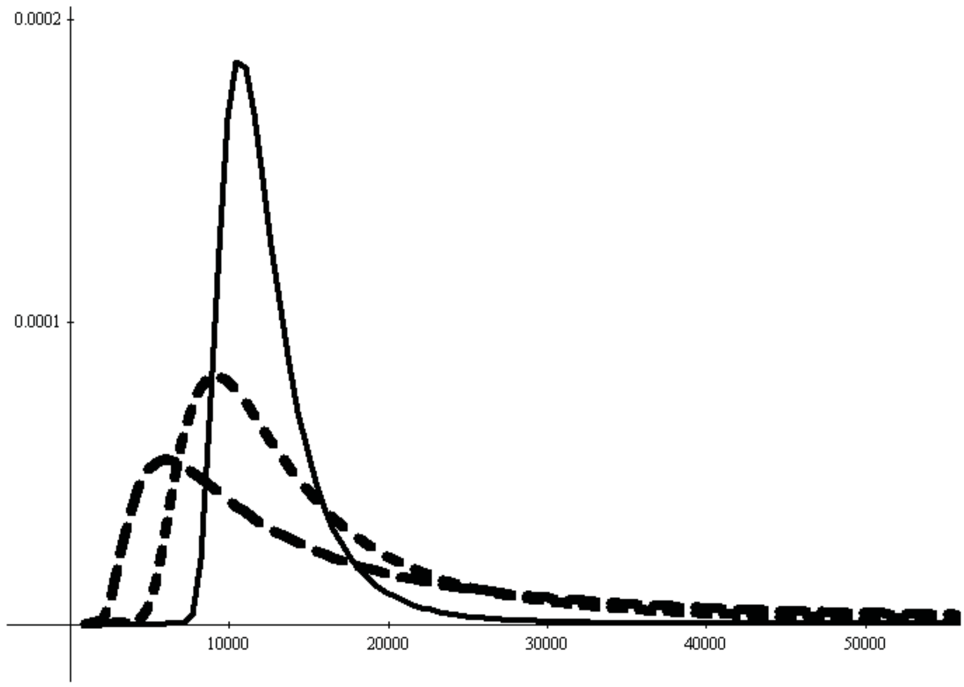

Frechet distribution with parameters 5,2 and 1

Figure 3.3. Frechet distribution

3.16 The basic process follows the diagnostic/investigative/predictive approach, as described in ¶3.2. The diagnosis is to determine which distribution, and to estimate the parameters. This estimation may be done by a number of methods (maximum likelihood, methods of moments and so on). The investigative part will determine the selected model, and the predictive element will result in, for example, a rate for the higher layer of cover.

3.17 The GEV, in itself, is only a starting point. There are times when the estimating method breaks down because of the absence of real large events in the data, but using appropriate subsets may help in the estimation. The issue of where to truncate the data is important, as using too much of the data may lead to underestimation. We are only interested in the tail. Other issues to consider when reviewing a long period of data, such as hurricane claims, even after allowing for changes in the population in cities and housing price inflation, there is always the possibility that the underlying time series may be non-stationary. An example involving this is covered in Section 4.

3.18 If the analysis was used to make an estimate of the overall claims for a portfolio, for example using a Monte Carlo simulation, then it is important to understand properly the underlying distribution. In this case, a standard truncated distribution is fitted to the bulk of the data and an extreme value distribution to the extreme claims. Simulation packages such as @Risk have the Gumbel and Weibull distributions.

The Generalised Pareto Distribution

3.19 The second important distribution is the generalised Pareto distribution (GPD). Suppose that Y is a random variable (e.g. claim size), then the tail is described by the Pickands and Balkema-De Haan theorem (see Pickands, 1975):

$$\Pr(Y > y + u | Y > u) \approx \mathrm{G}(y; u, \xi, \sigma) = [1 + \xi y/\sigma]_+^{-1/\xi}, \quad \text{for } y > 0. \qquad (1)$$

Here the distribution function $1 - \mathrm{G}$ is the GPD, and the scale parameter σ depends on the threshold u. The shape parameter ξ is the same parameter which appears in the GEV, and therefore controls the weight of the tail. The symbol $\approx$ indicates an approximation. As u increases, the approximation improves, and then, for $y > u$, large u, we get:

$$\Pr(Y \leq y) \approx 1 - \lambda_u[1 + \xi(y - u)/\sigma]_+^{-1/\xi} \qquad (2)$$

where:

$$\lambda_u = \Pr(Y > u).$$

3.20 Over some threshold, the GPD approximates most distributions. Therefore, to fit the tail of a distribution (for example, to price a high XL layer), one needs to select a threshold and then fit the GPD to the tail. For many insurance pricing purposes this is more convenient than the GEV. This approach is sometimes known as 'peaks over threshold', or POT. The main issue is the selection of the threshold, as the distribution is conditional on that amount.

3.21 The key decision in a GEV is the decision of the period over which the maximum is being assessed. Under a GPD, the decision relates to the threshold u, where the underlying data will be approximated by the conditional GPD. The trade-off is between the quality of the approximation to the GPD (good for high u) as against the level of bias in the fit which we achieve (good for low u, as there are more data to fit the GPD). Typically, one may choose u around the 90% to 95% percentile.

3.22 To fit the GPD to a dataset, the threshold u needs to be selected so that the asymptotic result (1) holds. Following this, estimates are made of the probability λ_u of exceeding this threshold, and the parameters σ and ξ. Equation (2) implies a GPD, and the mean exceedance of a threshold is linear in u, so a test is that the graph of the excess of the GPD random variable Y over the threshold u, as u varies, is linear. A simple approach is to fit a GPD to a random variable (say a claim amount) by plotting the mean exceedance of the random variable $\mathrm{E}(Y - u | Y > u)$ against u. After some threshold point, the fit will approach linearity. Above this point, fit a GPD to the data by calculating λ_u, and then use standard techniques (MLE, method of moments) to calculate ξ and σ.

3.23 Significant judgement may be required in choosing a threshold *u* to fit the GPD. As well as the trade-off between accuracy and bias, there may be difficulty in deciding at which point the mean exceedance has passed into a linear regime. The use of the mean exceedance plot is, therefore, a key diagnostic test in fitting a GPD, and requires a certain amount of judgement to choose the threshold. However, just because some level of judgement is involved does not mean that the results are invalid — all actuarial techniques require judgement to be exercised, and the use of extreme value theory highlights the role of that judgement.

3.24 In Sanders *et al.* (2002); a detailed example with a spreadsheet is given. This example is summarised below. The plot is based on data where a mean excess of 406,000 was selected. The data were, therefore, the original claim less the 406,000. For example, consider a loss of 712,119. Based on a mean excess of 406,000, the data point is 306,119. The mean excess value of 387,397 is calculated as the average of all the values from 306,119 and higher, minus 306,119.

3.25 For a true GPD, the mean excess function is linear. Therefore, if the plot becomes approximately linear in the tail, it is an indication that the GPD might be a good fit to the tail of the data. If no part of the plot is linear, the GPD is not an appropriate fit to the data. There is always the question that the higher the threshold, the less the data points used in the fit, and, accordingly, the GPD will fit a distribution at a sufficiently high threshold. The mean excess plot for the example is given in Figure 3.4.

3.26 Here, one might decide that the graph is approximately linear from

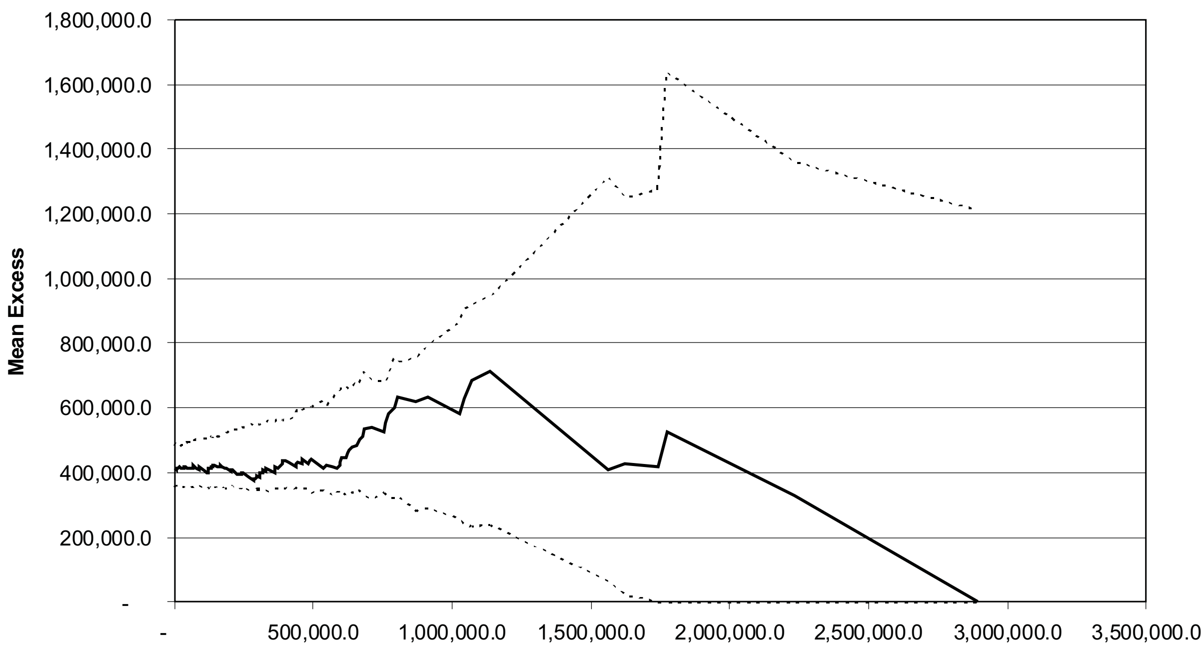

Figure 3.4. Mean excess plot

the chosen threshold (of 406,000) or at a higher threshold (of around 1,000,000, say). Both choices would fit the GPD, however, they would fit the GPD with different parameters. At the chosen threshold of 406,000, there are 123 data points included. At a threshold of 1,000,000, this reduces to 28. The trade-off between approximation to the underlying distribution (good for high threshold) and bias (good for low threshold) is particularly brutal in the world of extreme values. As a rough observation, significantly fewer than 40 observations can make the GPD fit quite variable. At a threshold of 1,500,000, there would be just seven observations used to fit the GPD. Clearly, each one of these observations can have a significant impact on the fitted distribution.

3.27 One significant observation is that, at the chosen threshold, the mean excess plot is, broadly speaking, flat. At a higher threshold, the mean excess plot decisively slopes downward. This change would point to a significant difference in the parameterisation of the fitted distribution at the different thresholds. This is fundamental, as it shows that the shape parameter of the distribution is potentially changing. Any simulation might have to factor this in.

3.28 We now need to check on the reasonableness of the estimates, and, to help, we have three visual tests.

QQ plot

3.29 This plots the data values against the equivalent percentiles for the parameterised GPD. For example, in the sample dataset there are 123 data values (when the threshold is set to 406,000). The value of 306,119 is the 62nd highest value, which is equivalent to the median (50th percentile) in the empirical distribution. This value is plotted against the 50th percentile of the GPD, which is 269,184. A 'perfect' fit would result in a line $y = x$. How close the QQ plot is to this line is a measure of goodness of fit. Drift away from this line in one direction indicates that the underlying distribution may have a heavier (or lighter tail) than the fitted distribution. The QQ plot for the example is given in Figure 3.5. For the sample dataset, the fitted distribution appears to have a heavier tail than the empirical distribution, as the higher empirical percentiles lie beneath the line $y = x$. This corresponds to the drift downward in the mean excess plot, which indicates that a higher threshold chosen would lead to a less heavy tail in the fitted distribution. This indicates, as we suspected, that claims in excess of £1m may need a different parameterisation. However, if we are viewing the cost of a layer £500k excess of £500k, then the model appears reasonable.

PP plot

3.30 This plots the percentile of the empirical distribution for each datum value, against the implied percentile of the GPD for the same value.

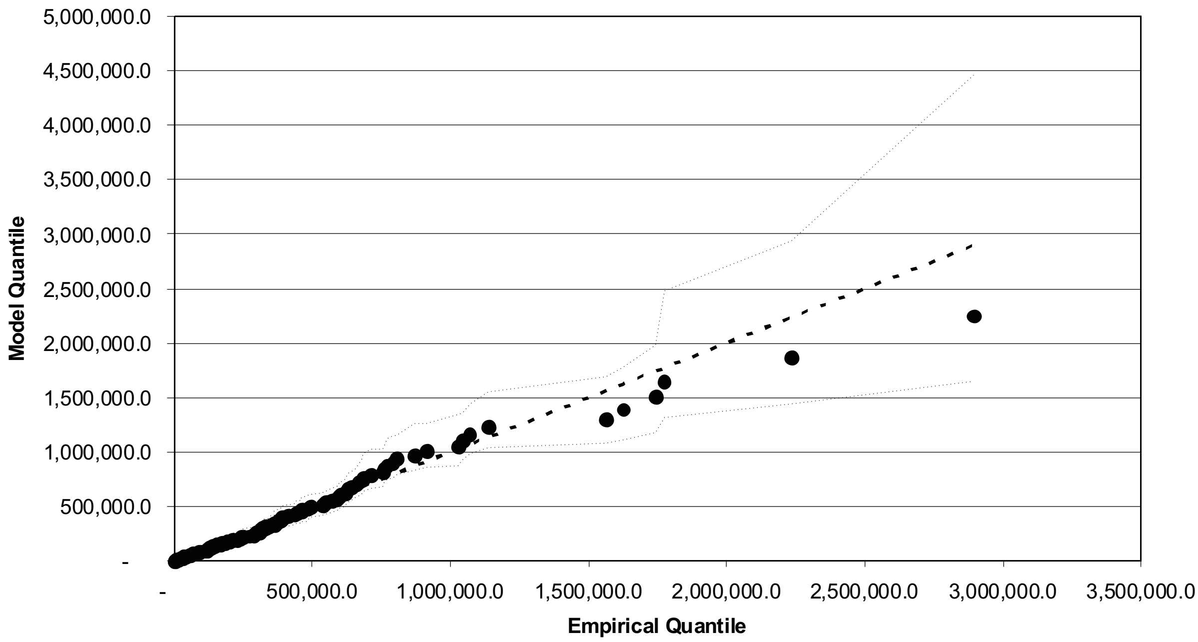

Figure 3.5. QQ plot

For example, the value 306,119 is the 50th percentile of the empirical distribution and the 54th percentile (approximately) of the GPD. As with the QQ plot, a 'perfect' fit would result in a line $y = x$. The PP plot for the example is given in Figure 3.6.

3.31 The QQ plot is better than the PP plot when assessing the goodness

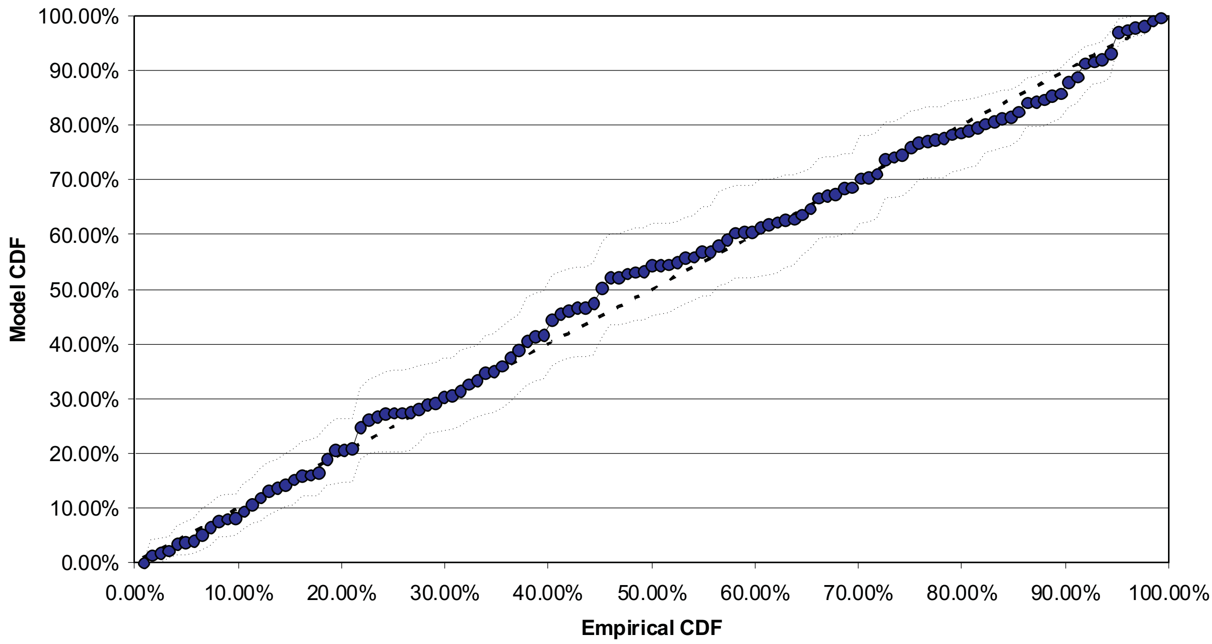

Figure 3.6. PP plot

of fit in the tail of the distribution; the values of the PP plot will tend towards 100%, so a bad fit in the tail is not always apparent, as can be seen when comparing the two outputs.

CDF plot

3.32 This plots the empirical claims distribution (CDF) against the GPD. The same comments on the values apply as for the PP plot, in that both the empirical and the fitted distribution must both converge to 100%, and so a poor fit in the tail may be obscured. The CDF plot for the example is given in Figure 3.7.

3.33 In this particular case, it is not surprising that the whole dataset is a good fit, as it comprises only losses over a threshold; i.e. the dataset already represents the tail of an underlying distribution. If the plots suggest that a good fit has not been obtained, then the threshold needs to be increased until, either an acceptable plot is obtained or it becomes clear that a GPD is not a good fit. As often occurs, there might be a trade-off between goodness of fit and quantity of data, in which case there is no definitive answer, and judgement must be applied. In situations with smaller datasets, the application of judgement becomes key, and a variety of possible distributions are justifiable, in the same way that significant variation may emerge through different Cat models.

Multi Dimensional Distributions

3.34 The above deals with single variate EVT in a static environment.

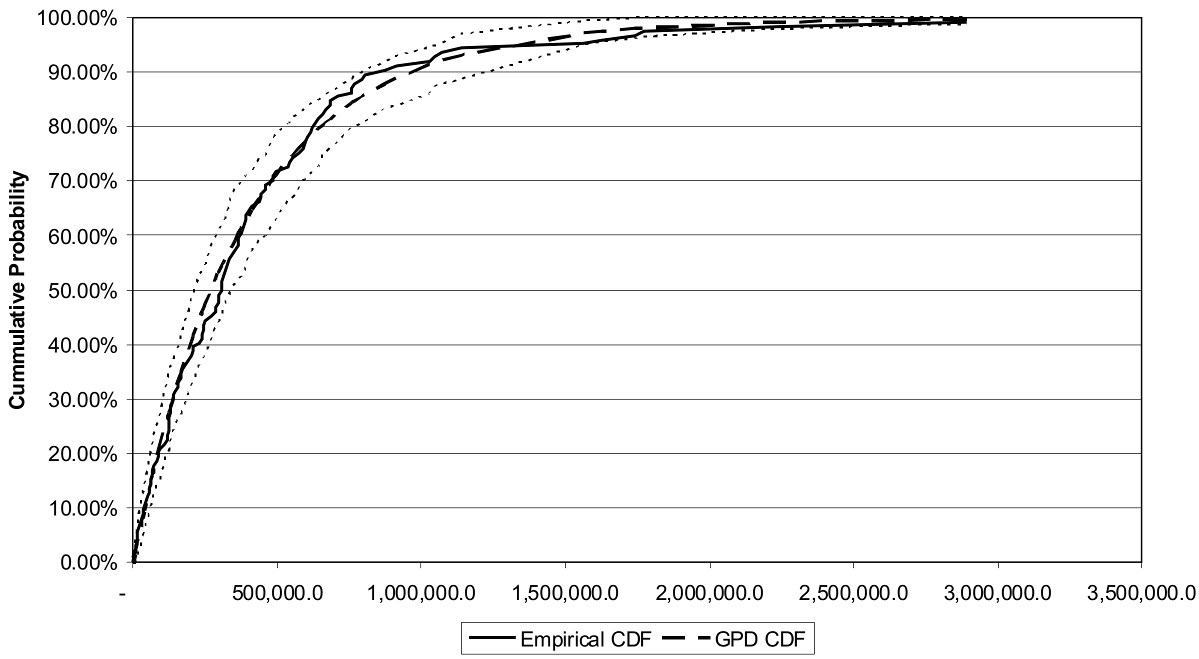

Figure 3.7. CDF plot

Recent developments involve extending the theory to multivariate extremes. Let us consider a model which has two processes. Each process may be modelled separately, and be connected by some relationship. However, these relationships mean that the models are often less fully prescribed. Furthermore, the additional dimensions make computation and validation more difficult.

3.35 Examples of multivariate distributions are found particularly in finance, when considering the relationship between stock market indices and foreign exchange, or the level of sea water over a number of sites, or even losses from windstorm from different countries.

3.36 One approach is to view a vector of component wise maxima:

$$\mathbf{M_n} = (M_{x,n}, M_{y,n})$$

where:

$$M_{x,n} = \max\{X_1; X_2; : : : ; X_n\}$$

and

$$M_{y,n} = \max\{Y_1; Y_2; : : : ; Y_n\}$$

and analysing $\mathbf{M_n}$. See Coles (2001) and Fougeres (2002). An alternative approach is to link the values through copulas.

3.37 Copulas are simply the joint distribution functions of random vectors with standard uniform marginal distributions. Their value is that they provide a way of understanding how marginal distributions of single risks are coupled together to form joint distributions of groups of risks; that is, they provide a way of understanding the idea of statistical dependence. There are two principal ways of using the copula idea. They can be considered from known multivariate distribution functions or created by joining arbitrary marginal distributions together with copulas.

3.38 The copula is a function C from the n dimensional cube to [0, 1]. In a two-dimensional copula we have:

$$C(0, t) = C(t, 0) = 0$$

and

$$C(1, t) = C(t, 1) = 1.$$

If $u1 < u2$ and $v1 < v2$, then:

$$C(u2, v2) - C(u1, v2) - C(u2, v1) + C(u1, v1) > 0.$$

More particularly, a copula is defined as follows. A copula is any function C from the n dimensional cube $[0, 1]^n$ to $[0, 1]$, that has the following three properties:

(1) $C(x_1, \ldots, x_n)$ is increasing for each component x_i.
(2) $C(1, \ldots, 1, x_i, 1, \ldots, 1) = x_i$ for all i.
(3) For all $(a_1, \ldots, a_n), (b_1, \ldots, b_n)$ with $a_i \leq b_i$, then:

$$\sum_{i1=1}^{2} \cdots \sum_{in=1}^{2} (-1)^{i1+\ldots+in} C(x_{1i1}, \ldots, x_{nin}) \geq 0$$

where $x_{j1} = a_j$ and $x_{j2} = b_j$.

3.39 Embrechts *et al.* (1999) indicated that the concept of correlation had many pitfalls, and proposed the use of copulas to quantify dependence.

3.40 The extension of multivariate to infinite dimensions leads to max-stable processes (see Smith, 2003), which are beyond the scope of this paper.

Power Laws

3.41 A power law is a function $f(x)$, where the value y is proportional to some power of the input x:

$$f(x) = y = x^{-\alpha}.$$

The power law distribution is denoted $P(> x) = x^{-\alpha}$. Because power laws usually describe systems where the larger events are more rare than smaller events (i.e. magnitude eight earthquakes happen much less often than magnitude two), α is positive. Power law models tend to relate to geophysical events. An example is the Gutenberg-Richter law. Data indicate that the number of earthquakes of magnitude M is proportional to 10^{-bM}. The graph in Figure 3.8 relates to the number of earthquakes in 1995. The fitted line gives the Gutenberg-Richter prediction with $b = 1$. The value of b seems to vary from area to area, but worldwide, it seems to be around $b = 1$. Atmospheric phenomena have been modelled on power law (Dessai & Walter, 2003), with the results set out in Table 3.1.

3.42 It is important to note that EVT methods of tail estimation rely heavily on a power law assumption that is, that the tail of the survival function is assumed to be a power law times a slowly varying function. The 'tail index' α is a parameter to be estimated.

3.43 Power laws are related to Pareto law, whose main work was on the distribution of income. Instead of asking what the nth largest income is, he asked: "How many people have an income greater than a specific amount?" Note the similarity in the question to the one asked regarding extreme events. Pareto's law states that the cumulative number of events larger than x is an inverse power of x; that is $P(> x) = x^{-\alpha}$. The GPD is just an extension of this

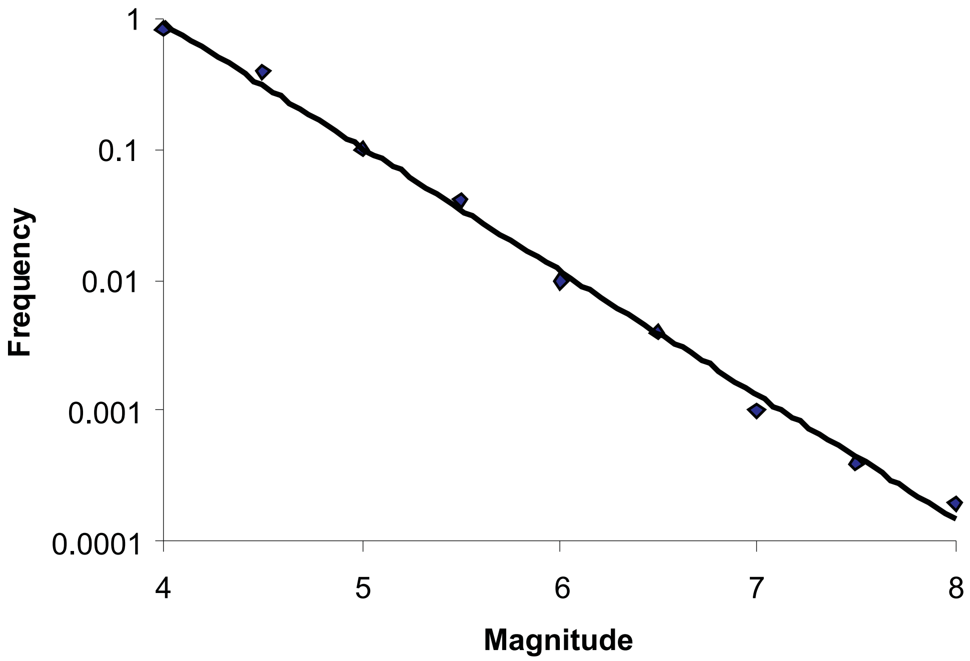

Figure 3.8. Earthquake numbers in 1995 and Gutenberg-Richter prediction

Table 3.1. Predicted power law exponents for atmospheric phenomena

	Exponent	R-squared
Atlantic hurricanes	2.2961	0.8190
Pacific hurricanes	1.0362	0.7450
Windstorms	11.873	0.9800
Tornadoes	2.7215	0,7478
Forest fires and precipitation	0.6551	0.4376

law. For a more detailed analysis of Pareto and other distributions, see Kleiber & Kotz (2003).

3.44 Given that many natural phenomena appear to follow a power law, it is not surprising that EVT, which depends on such an assumption, gives a good fit to reality. There is, thus, a strong connection between the geophysical power law models and the statistical GPD models. Thus, EVT and the power laws of geophysics appear, for the purpose of modelling, to be consistent.

3.45 Power laws have the characteristic of randomly distributed values that come from a scarce resource. They also have the characteristic of scale inversion, that is the that the distribution looks the same, regardless of the scale. This leads to other concepts, such as self organising criticality and complexity, which will be dealt with briefly in Section 6.

Catastrophe Models

3.46 Catastrophe (or 'Cat') models are used to assess the cost of a hypothecated catastrophe on a book of business, or to estimate the actual likely cost of a known catastrophe. By means of using a number of hypothecated catastrophes, estimated costs of various layers of excess of loss reinsurance can be assessed and compared. They often use models derived from data using EVT analysis to estimate the cost.

3.47 Cat models are essentially different from EVT and geophysical models, in that they are more complex, and rely on a substantial number of assumptions and great computer power. They are analysing the losses over the whole range of catastrophic events, rather than concentrating on the extreme events.

3.48 Each insurer has a 'loss profile', giving a graph of the estimated loss against the return period of storms, By comparing these profiles, a reinsurer can assess whether a particular insurer is a 'better risk' than another.

3.49 A Cat model consists of a number of modules:

— *the stochastic module*, which is used to randomly generate the catastrophic events;
— *the hazard module*, which is used to determine the geographical effect of a catastrophic event brought about by variations in topography, e.g. regional flood depths, lower wind speeds in sheltered areas, and so on;
— *the vulnerability module*, which is used to calculate damage to buildings, contents and business turnover, based on a number of factors, including, for example:
  — building type: residential, commercial or agricultural;
  — building design and construction quality;
  — building height/number of floors; and
  — occupancy type, and so on; and

— *the financial module*, which quantifies the financial loss to the insurer.

3.50 Cat models tend to underestimate the impact of extreme events, because the underlying parameters are based on normal events. The extreme catastrophes are not modelled, due to the lack of really large losses in the data needed to adjust the parameterisation of the model. Large catastrophes have, historically, produced losses greater than predicted by such models, but the modelling is improving.

3.51 Cat models have the advantage that they meet the insurer's need. In particular, they are based on scientifically/statistically derived models of relevant perils, and can take account of changes in the insurer's exposure, underlying policy details, building codes, the physical environment, and changes such as flood defences. Losses are calculated in grids, and these grids can be refined for more detailed assessment. As such, Cat models are similar to those used to assess weather forecasts.

3.52 Having said this, Cat models have many perceived disadvantages. The suppliers of the models naturally treat the models as proprietary, and do not wish to release their inner details. They are also extremely complex, and hence difficult to understand, and different results are obtained from different models.

3.53 The quantity of data required to run these models is enormous, and may be difficult and costly to obtain, and modelling time is long and the cost expensive. As such, up-to-date information may not be available to reinsurance underwriters. Furthermore, there are often approximations/compromises in the models, where interpolation and extrapolation are used. Most important, there is, invariably, calibration error, particularly in the more extreme catastrophes, which leads to underestimation of losses. This occurred both in Hurricane Andrew (1992) and Northridge Earthquake (1995). The latter earthquake occurred at an 'unknown', and hence unmodelled, fault, and initial estimates were out by a factor of ten on some models.

3.54 Cat modelling has been used for a number of reasons, in practice, other than for extreme events. The major extreme event management uses are as follows:

- assessing return periods of storms to assess Cat XL capacity: typically, one is looking at return periods of 100 to 250 years; however, storms, like the United Kingdom windstorm of January 1990, may typically have longer return periods (see Section 4);
- assessing capital requirement, which is often driven largely by Cat exposure;
- the Cat model results can be used as input to the stochastic DFA models; and
- Cat models have been used to asses the extreme risks in quantifying risk and costs in securitisations and catastrophe bonds.

3.55 In conclusion, the statistical/actuarial models used in practice for extreme events are based on consistent assumptions, such as EVT, power law or Pareto's law. For more sophisticated models a simulation approach is used, which should incorporate these type of models for determining the impact of extreme events.

4. Types of Extreme Events

4.1 As part of their professional responsibilities, actuaries are interested in extreme events for a number of reasons. They may not realise that they are dealing with extreme value theories, primarily because most of the work deals with numbers which are assumed, implicitly or explicitly, to be normally distributed. Some types of events where extreme value theories are, or can be, used in practice include:

- natural catastrophes, such as windstorm, volcanic eruptions, earthquake, tsunami;
- 'man made' catastrophes, such as aviation crashes, explosions, oil spillage; and
- financial events, such as sudden drops in the stock market and even increasing longevity.

Fuller details, with references, may be found in Sanders *et al.* (2002) and Sanders (2003).

4.2 There is often a reliance on the underlying mathematical modelling to predict such events, and to estimate the likely cost to society and the insurance industry. At this stage, it is worth considering such events in this context. Other examples of a more unusual, and possible controversial, nature are given in Section 6.

Natural Catastrophes

Weather related events

4.3 It is appreciated that the earth's weather is a hugely complex system, with multiple interacting elements. Because they defy simple expression, meteorologists rely on a 'bottom up' approach to modelling the climate. They divide the world into boxes, model the forces on the atmosphere and oceans in each box, and then estimate the overall effect. The approach used is similar to that used in catastrophe modelling to assess the cost of weather events. However, in both cases, to make the estimates work, one often needs to include a number of 'fudge factors' to make the models fit with known events. Weather models tend not to foretell catastrophic storms, as predictions tend towards the central assessment.

4.4 Conversely, there is also the top down approach, where the driving entity is the sun. Solar energy warms the earth, and it does so disproportionately, heating the tropics much more than the poles. This creates a temperature gradient, with the effect that heat goes from low to high latitudes and drives the climate. This heat flow produces mechanical work in the atmosphere and oceans. The process takes place at a pace which is favourable. If the process were too fast, then temperature throughout the world would be the same, and if it were too slow, then Northern Europe would be in an ice age.

4.5 The randomness of storm events required by EVT has been 'disproved' by a number of recent events. Changes in the weather pattern appear to force storms along particular routes, making multi hits. In January 1990, the boundary between the warm and cold atmospheres, the jet stream, speeded up and moved eastwards. A consequence of this was that storms that would normally pass between Iceland and Scotland were now coming up the English Channel, with devastating consequences.

4.6 The process was repeated in late 1999, when Europe (particularly

France) was being hit by storms — these storms have been estimated from catastrophe models as being 'as bad as it can get'.

4.7 During 2004, Florida and the southern states were hit by four consecutive hurricanes, producing the largest cumulative storm loss in the region in any quarter. Models which ignore the potential for patterns of losses will, thus, tend to underestimate the problem. This also illustrates the fact that, when modelling for insurance purposes, it is not the single extreme event which will, necessarily, create a realistic worst case scenario.

4.8 An example of how trends and cycles impact on the assessment may be found in a preliminary analysis of a data set of North Atlantic storms, constructed by the Benfield Greig Hazard Group ('BGHG'), and is summarised in Smith (2003). The objective of this analysis is to study long-term trends in insurance claims. The research was trying to isolate the impact of climate change on insurance claims. Claims values increase for a number of reasons, including value inflation, increasing population in certain areas, and so on. In an attempt to understand the impact of climate change, BGHG undertook an analysis to try to estimate the impact of 57 historic storms, which occurred between 1920 and 1995, if the storms had occurred in 1998. This re-indexing approach is quite normal, as U.S. hurricane data are also revised in such a way. It is important to do this re-indexing, to try to obtain IID data, so that we can apply extreme value modelling techniques.

4.9 In undertaking this analysis, BGHG factored in seasonal effects, trends and dependence on climatic indicators, such as SOI and NAO. The Southern Oscillation Index (SOI) is the pressure difference between Darwin and Tahiti, and is used as a measure of the strength of the El Nino effect. The North Atlantic Oscillation (NAO) is the pressure difference between Iceland and the Azores, and is, therefore, more relevant to British meteorology.

4.10 The problem facing the U.K. insurance industry is the likelihood of large losses exceeding the largest value in the present data set, which was the January 1990 windstorm (Cat 90A), which caused damage across a very large area of the U.K. This was more expensive than the storm of October 1987 (87J), which involved stronger winds, but was concentrated over a much smaller geographic area.

4.11 The natural question was to estimate the return period associated with the 1990 event, calculated under various models. This is summarised in Table 4.1. The first row of the table assumes the simple Pareto distribution. This distribution did not fit the data at all, since the true distribution of the tail was much closer to exponential than the Pareto. Fitting the GPD instead of the Pareto, and then adding a seasonal effect, had a small impact on reducing the estimated return period. The potential impact of the NAO is much clearer. This example is a statistical model which was combined with some physical modelling. All the components of ¶3.2 may be found in this example. This also indicates the need to make sure that the data are IID by

Table 4.1. Return period for Storm 90A based on differing assumptions

Model	Return period (years)
Pareto	18.7
Exponential	487
GPD	333
GPD with seasonal fluctuations	348
GPD with high NAO	187
GPD with low NAO	1,106
GPD with random NAO	432

removing trends, cycles, and season fluctuations before estimating the parameters of the GPD.

Geophysical Events

4.12 Geological events differ from weather events. We know that earthquakes occur at tectonic boundaries, we know where volcanoes are, and tsunami tend to impact on coasts. However, meteorite or comet collisions can occur anywhere. Set out in Section 5 is an analysis on meteorite collision. Unlike windstorms, where modern weather tracking gives many days warning, geophysical events (with the possible exception of volcanoes) often occur with little or no true predictive warning.

4.13 Earthquakes are known to follow a power curve between intensity and return period (see ¶3.41). For example, in San Francisco there are a considerable number of small quakes every day, although it takes many years between quakes of high intensity to cause disruptive damage. Following a power curve model to its conclusion means that it is possible to have earthquakes of moment magnitude intensity seven, eight, nine or even 12! The latter would, of course, be very rare. However, the largest recorded earthquake is only 9.6, off the coast of Chile in 1960. The recent earthquake in the Indian Ocean was rated a 9.0. Accordingly, there are physical restraints to consider in the model.

4.14 Furthermore, it is unlikely that an earthquake in San Francisco will exceed 8.2. The reason for this is geophysical constraints. As indicated above, the number of earthquakes with a magnitude $M > m$ is given by the Gutenberg-Richter law. However, the real modelling is more complex. There are short range temporal correlations between earthquakes which are expressed by Omori's law, which states that, immediately following a main earthquake, there is a sequence of aftershocks whose frequency decays with time as T^{-a}. A good estimate in most cases is $a = 1$. Furthermore, reviewing the data at particular earthquake regions indicates that the epicentres tend to be discrete, and that local topography can amplify or suppress the impact. Finally, different types of faults behave in different ways, some boundary faults can get 'locked', and give rise to significant earthquakes, while, in other faults, plates move smoothly past each other without giving rise to major earthquakes.

4.15 Contrast this with the value of loss following an earthquake. Although there may well be a maximum size constrained by geophysical considerations, the economic value potential loss may be seen as having a very high limit that is unlikely to be breached. Furthermore, the economic value at risk varies from time to time, and tends to increase due to increase in value and demographic movements. The matter is also complicated by the fact that earthquake waves may be focused by geophysical features, causing significantly more loss than their apparent magnitude (Mexico), or the effect may be spread over a considerable area, due to a tsunami carrying the energy from the location of the earthquake to coastlands many thousand of miles away (Indian Ocean, December 2004).

4.16 Cat models have been used with some success in making earthquake estimates. One of the critical elements is the building code of the property and its ability to withstand the ground shake. An example of an earthquake model dealing with is given in Kuzak *et al.* (2004).

Man made catastrophes

4.17 The term 'man made' catastrophes is used for other physical losses. More detail is given in Sanders *et al.* (2002). These include:

— *Explosions and fires* occur, of which the most damaging are chemical fertiliser or petroleum based, such as the Halifax Explosion (1917), the Texas City Explosion (1947), Piper Alpha (1987) and Toulouse (2001). These large losses have a return period of about 30 to 40 years, with potential loss values as high as the September 11 loss. As such, they are covered by traditional reinsurance. Fitting a GPD is the usual approach to analysing this loss.

— *Major motor losses* tend to be of two types: either one vehicle causes considerable damage (for example, the Selby rail crash), or there are other events which have severe consequence, for example recent Alpine tunnel fires. These losses are covered by insurance and reinsurers, although reinsurers are becoming more reluctant to cover the larger potential losses.

— *Aviation disasters* are all well documented. On average, 700 lives are lost in commercial aviation losses each year.

— *Terrorist acts* are the subject of considerable debate as to whether the costs should be borne by the reinsurers or by government.

— *Marine disasters*, where, excluding ferry losses in the Philippines and World War Two losses in the Baltic, the major marine disasters involving loss of life are the Titanic and the Estonia. There are in excess of 380 vessels sinking (or being total losses) every year. Claims excluding loss of life include pollution spills. The most expensive loss was Exxon Valdez (36,000 tons). However, this ranks only in the mid 30s in terms of the largest spill.

Extreme man made losses tend to be analysed using extreme event theory analysis, as they tend to be random events.

5. Applications of Extreme Value Theory and Other Models

5.1 The assessment or analysis of extreme (and other large loss) events consist of three elements. These are:
- identifying all such events and the range of likely losses at any given level of event;
- estimating their likelihoods (e.g. using mean annual frequencies); and
- quantifying the uncertainty in these estimates.

All of these elements give rise to challenges, most particularly because of the incompleteness of data. It is almost certainly necessary to incorporate some form of extrapolation, or information from other models or sources, to simulate completeness.

Excess of Loss Pricing

5.2 Extreme event theory has been used, in practice, by actuaries for a number of years. Beard (1962) and others used it to explore the largest property fire claims, to provide estimates for the price for excess of loss reinsurance contractors. There is a set of Danish fire data that has been analysed extensively (McNeil, 1997) and others.

5.3 Set out in Figure 5.1 is the empirical distribution function of the

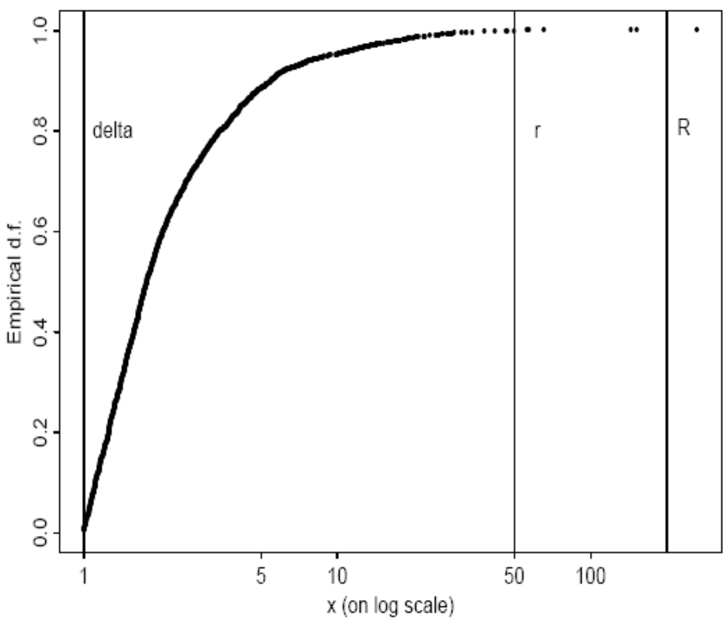

Figure 5.1. Danish fire data (from McNeil)

Danish fire loss data, evaluated at each of the data points. What is needed is an estimation of the distribution beyond the data. The Danish fire data comprised 2,157 losses in excess of 1 million Danish Krone. In his paper, McNeil placed estimates on a high excess layer running from 50 to 200, where there are only six observed losses. A reinsurer needs to be able to obtain a good estimate of the severity distribution in the tail.

5.4 In the analysis, McNeil identified the extreme losses, and also whether there is evidence of clustering of large losses, which might cast doubt on the assumption of independent identically distributed data. This did not seem to be the case. McNeil then followed the same process as described in Section 3, namely estimating a mean excess function and the Q-Q plot. The result are summarised in Table 5.1, with graphic interpretation in Figure 5.2. A similar example may be found in the Xtremes package in respect of Norwegian fire data (see Reiss, 1997).

5.5 Gary Patrik made a presentation at the CAS Spring Meeting in 1998, where he illustrated the application to long tail liability business based on the ISO claim severity model. He found that the right tail of the ISO U.S. liability claim severity model fitted to claims, generally up to a limit of \$1,000,000, is a U.S. Pareto (or Burr distribution with x-exponent $=1$). He identified that the shape parameter of the GPD was approximately -1, and that this applied over several lines of business. He defined $\xi = -1/a$, and estimated the parameter 'a'. Examples of his estimates of 'a', by class of business, are shown in Table 5.2. He also indicated that an analysis of very large claims by a large global reinsurance company has shown that their distribution could be approximated with a GDP with shape parameter $\xi = -1$. The final comments in his presentation were: "P&C actuaries should test their large claims data for convergence to GPD and for k

Table 5.1. Results of Danish fire analysis

Model	Threshold	Excess	Shape	s.e.	.999th	.999th	P
GDP	3	532	0.67	0.07	129	603	0.21
GDP	4	362	0.72	0.09	147	770	0.24
GDP	5	254	0.63	0.10	122	524	0.19
GDP	10	109	0.58	0.14	95	306	0.13
GDP	20	36	0.68	0.28	103	477	0.15
		Models fitted to whole data set					
GDP	All	Data	0.60	0.04	101	410	0.15
Pareto	All	Data			235	1,453	0.10
Lognormal	All	Data			82	239	0.41
		Scenario models					
GPD	10	108	0.39	0.13	77	201	0.09
GPD	10	106	0.17	0.11	54	96	0.02
GPD	10	110	0.60	0.15	118	469	0.19

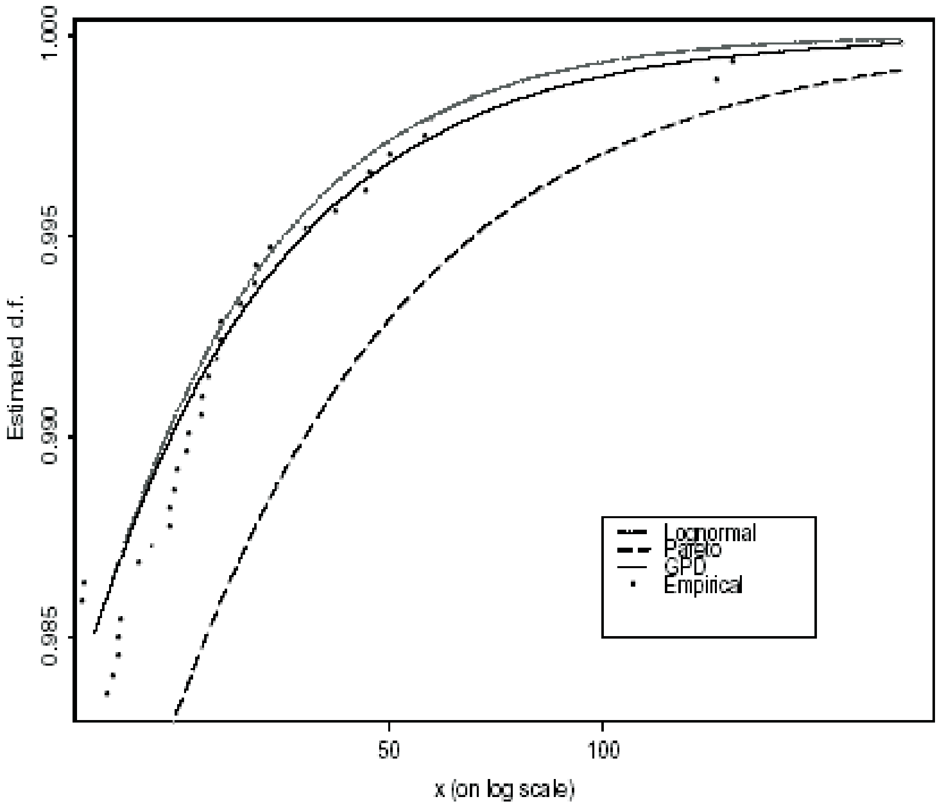

Figure 5.2. Graphic interpretation of results (from McNeil)

Table 5.2. Estimates of Patrik's *a* parameters

Class	Low *a*	High *a*
Commercial auto	0.95	1.45
Premises/operations	1.10	1.70
Products	1.10	1.20
Medical professional liability	0.75	1.40

approximately equal to −1. P&C actuaries should use the GPD to model excess losses, or make some lame excuse". Extreme value distributions are a necessary part of the actuarial pricing toolbox, and the parameter −1 is a good initial estimate based on Patrik's work.

PML Estimation

5.6 A problem is estimating the likely maximum loss for a particular class of business, for example cargo. It is difficult for cargo insurers to know where the actual risk is at all times, be it on board a vessel or somewhere onshore, for example a warehouse. It could be on the high seas, where each vessel carries as much as $2bn insured loss, or it could be in a warehouse, where an earthquake could cause considerable loss (e.g. Kobe). The idea follows precisely the same approach as a normal estimation of extreme events:

— first, take the known claims over a period, and adjust to current values, taking account of underwriting trends and other changes;
— estimate the threshold, and derive the EVT parameter;
— test the estimate; and
— using an extreme value function, estimate the tail.

One can then derive the PML for reinsurance purposes from the calculated distribution.

Catastrophe Bonds

5.7 The market is still young, with the first securitisation of catastrophic risks taking place in 1997. Up to 2002, the market was small, at about $1 billion per year. The perils underwritten were earthquakes and wind storms. The market was limited to a few participants, with the risk analysis being reviewed by Cat models. Basically, the bond paid a higher rate to allow for the possibility that a catastrophe would cost more than a specified amount. If the cost of a catastrophe exceeded a specific amount, then the bondholders would first lose the interest on the bond and then the capital.

5.8 Catastrophe, or ‘Cat’, bonds are generally priced using Cat models. The offer documents include an analysis of the range of outcomes from various losses and the associated costs to the bondholder.

5.9 Cat bonds have become more popular recently, and have been extended to new perils, such as extreme mortality and World Cup cancellation. There has been some criticism of the use of Cat models, in particular the fact that, historically, they have underestimated the cost, particularly at the extreme end of the range. Cat bonds are not restricted to insurance companies, one major bond being issued by Eurodisney. Possibly the greatest challenge is in respect of a bond based on terrorist losses. A more recent development is contingent debit facilities (also known as contingent capital). Under these, if a specific level of loss becomes payable, then a guaranteed loan facility is triggered.

5.10 The evaluation of risk, using a combination of Cat modelling and EVT, is an important process in assessing a Cat bond. If techniques are used that do not relate to EVT, then the bond is in danger of being underpriced.

Swiss Re Mortality Bond

5.11 Swiss Re is one of the world’s largest provider of life and health reinsurance. As such, it is exposed to very large amounts of mortality risk that cannot effectively be transferred to other life reinsurers, that already bear similar risks. Therefore, it has turned to the capital markets to hedge its mortality exposure. Swiss Re set up a special purpose vehicle, dubbed Vita

Capital Ltd, that issued $400 million in mortality indexed notes. The proceeds of the issue have been placed in a collateral account, and invested in high quality securities, whose returns are swapped into a Libor-based return. The notes are expected to mature on 1 January 2007, at par. However, if mortality in five countries increases to pre-defined levels, investors in the notes may receive less than par or even zero at maturity. Note holders will receive quarterly interest payments of Libor plus 1.35% (annualised), in return for bearing extreme mortality risk.

5.12 The bond was priced by a bootstrap stochastic method, as was described in the offer document. Beelders & Colarossi (2004) reviewed the methodology using EVT. The results and stress tests showed higher results than the model used, although well within the parameters offered. They concluded: "Is the 135bps spread sufficient compensation for the risk? According to the market, it is!"

5.13 The pricing of this type of bond also needs to take account of mortality trends, due to factors other than ageing. The recent 2003 heat wave in France suggested that there were up to 10,000 excess deaths. Old people's homes were operating with a reduced level of staff during the August holiday, hospitals closed wards for the month and were unable to offer even basic treatment. A number of people died in their homes from the effects of dehydration and other heat related problems, while neighbours and relatives were away.

Cat Swap (Derivative)

5.14 In August 2003, Mitsui-Sumitomo Insurance (MSI) swapped with Swiss Re $100m of windstorm and flood risk in Japan against $100m of windstorm risk in the U.S.A. and Europe. Japanese insurance companies tend to underwrite concentrating on the risks located in Japan, and this allowed both MSI and Swiss Re to diversify and spread their risks.

6. Unusual Extreme Events and Other Ideas

6.1 In this section we review some unusual events to which we can apply the theory.

What are Really Large Events?

6.2 In Sanders *et al.* (2002) was an estimate of the average asteroid collisions cost. The results are set out in Table 6.1, and follow, for example, similar techniques applied to reinsurance pricing in Sanders (1995a). The estimated number of deaths is 3,900 p.a. For comparison, the average annual death toll from earthquakes is about 10,000, although, in many cases, the impact can exceed more than 100,000. The number of deaths in commercial airliner crashes is about 700 p.a.

Table 6.1. Estimate of death toll from various types of impact

Asteroid diameter (m)	Area devastated (sq km)	'Typical' direct fatalities	Ratio of indirect/direct fatalities	Total fatalities	Annual chance for inhabited regions 1 in ...	Equivalent annual death toll
50	1,900	200,000	4	1 million	900	1,100
100	7,200	650,000	4	3 million	8,000	400
200	29,000	2,000,000	6	14 million	30,000	500
500	70,000	4,000,000	8	35 million	180,000	200
1 km	200,000	7,000,000	8	63 million	290,000	200
2 km	–	–	–	1.5 billion	1,000,000	1,500
All					800	3,900

6.3 There are, of course, other extremes. One of most dangerous realistic potential extreme events is a giant caldera explosion. Geological research implies that this sort of event occurs about once every 70,000 to 100,0000 years. The last was 70,000 years ago, and genetic studies indicate that it almost led to our extinction! An example of such a caldera is Yellowstone Park, where the caldera is not obvious and is several kilometres wide. A second example of a realistic potential extreme is a tsunami, following a landslide on the Canary Islands. Tsunami generated by submarine earthquakes tend to have a maximum height measured in low tens of meters. Those from landslides could be as much as ten times as great. These different characteristics are determined from physical simulation models.

6.4 These results raise the fundamental question of the role of insurance companies in extreme events. Currently, all these losses are covered under the contracts issued. We know that aviation and satellite insurers struggle to make a profit in aggregate. Earthquake deaths are mainly in developing countries, but, in the near future, these lives may be insured at significantly higher values. This issue is covered more in Section 7.

What is the Maximum Age to which We can Live?

6.5 Benjamin Gompertz attempted, over 100 years ago, to find a mathematical explanation of the relationship between age and death rates. The relationship is described by the Gompertz curve. However, the Gompertz curve simply carries on forever, suggesting that there is no limit to the human life span.

6.6 The Gompertz distribution is a special case of an extreme value distribution. Based on analysis in Sweden and the Netherlands, it was found that, although the Gompertz curve may have been applicable for most ages, it failed when dealing with extremes. Computer analysis showed that the data for the 'oldest old' fit a quite different extreme value distribution. Furthermore, the data indicated that there is, after all, a maximum life span for humans. However, the data are not complete, and estimates suggest that a reasonable choice for the confidence bounds is 113 to 124.

Banks and Operational Risk

6.7 A number of papers have been written on Basel II and operational risk. (Embrechts *et al.*, 2001, is a good example.) The issue is that value at risk is used as the measure. This measurement of risk has been criticised, as the models do not take into account operational risks (extreme or not), and they make various assumptions about normality, so exclude extreme and rare events (including major social or political events). The basis of an extreme event analysis of operational risk is to identify a threshold and model a GPD on events in excess of that threshold, in much the same way as we have analysed reinsurance pricing.

6.8 The need for this type of approach has been recognised by central banks:

> "A natural consequence of the existence of a lender of last resort is that there will be some sort of allocation of burden of risk of extreme outcomes. Thus, central banks are led to provide what essentially amounts to catastrophic insurance coverage ... From the point of view of the risk manager, inappropriate use of the normal distribution can lead to an understatement of risk, which must be balanced against the significant advantage of simplification. From the central bank's corner, the consequences are even more serious because we often need to concentrate on the left tail of the distribution in formulating lender-of-last-resort policies. Improving the characterisation of the distribution of extreme values is of paramount importance."
>
> Alan Greenspan, Joint Central Bank Research Conference, 1995

6.9 Applying an EVT to determine the operational risk cost in insurance companies is more difficult, as the main operational risks tend to relate to underwriting and claims management, and the losses arising from poor controls in these areas are only recognised well after the event, if at all. However, examples of 'operational risk' exist as benchmarks, the largest being the mortgage guarantee problems of the early 1990s, the catastrophe spiral of the late 1980s and the Enron and E&O/D&O losses of recent years.

Self-Organising Criticality

6.10 The concept of 'self-organised criticality' (SOC) was introduced by Bak (1996) to explain the occurrence of power laws in nature. Under a number of systems, such as those explained by Thom's catastrophe models (Thom, 1989), a critical point is reached at which a system changes radically, for example, from solid to liquid, or from a bull to a bear market. Here, there are parameters which can be varied or controlled to get this radical change. It is virtually impossible to tell when it will happen from such models.

6.11 Self-organised critical phenomena, by contrast, are exhibited by driven systems which reach a critical state by their intrinsic dynamics, independently of the value of any control parameter.

6.12 The example used by Bak is a sand pile. Sand is slowly dropped onto a surface, forming a pile. The pile grows, and as it grows natural sand avalanches occur, which carry sand from the top to the bottom of the pile. The slope of the pile becomes independent of the rate at which the system is driven by dropping sand. This is the (self-organised) critical slope. By measuring the distribution of avalanche sizes, a power law arises. As mentioned above, power laws exist in EVT, such as the Gutenberg-Richter relationship, which refers to the distribution between frequency and magnitude for earthquakes and in hurricanes and windstorms. More importantly, SOC systems are signatures of deterministic chaos, which have been considered to be a new way of viewing nature and natural catastrophic events.

Chaos and Complexity Theory

6.13 'Chaos theory' deals with physical processes that appear superficially to be random because of their variability, but are patterned in hidden ways. Chaos theory is used to find order in presumably random processes. Although determined, chaotic processes are not predictable beyond a short period, because the slightest variation in initial conditions changes the outcome.

6.14 'Complexity theory' relates to dynamical systems that generate paths of how a system evolves. These include the self organising systems mentioned above.

6.15 Chaos theory is also of interest because of the introduction of stable functions as explanations (Cauchy and Levy-stable), giving wider tails. These distributions arise out of power law considerations, and are being utilised in some investment simulation models.

6.16 What might be surprising is the close relationship between all these apparently different mathematical fields. The statistical theory is, for independent events, in excellent harmony with our understanding or modelling of the physical processes, and links directly to chaos and complexity theory, which are the 'new' sciences.

6.17 Under power laws, extreme events are much more common than under Gaussian laws. This is why members of management, who tend to be used to a 'bell shaped' expectation of events, are taken aback by the frequency of events that cause significant problems. Extreme value theory is used to explain stock exchange models and other financial events, but, like geophysical models, they need additional modelling components to be considered before putting into practice. Option pricing theory relies on Brownian motion which follows a lognormal distribution, and, therefore, does not exhibit power law type behaviour without some additional assumptions.

Financial Extremes

6.18 Sanders (2003) describes a number of examples of financial extremes. The main characteristics of many of the highlighted extreme events are:

— the creation of a bubble which has distinct economic assumptions from the rest of the market — this bubble philosophy creates values substantially higher than expected;
— the realisation that the economic assumptions are either flawed or unsustainable; and
— an overreaction, resulting in loss of value.

The initial distinct economic assumptions may be encouraged by individuals for personal gain with misinformation.

6.19 The characteristics of these extreme financial events seem to be out of place with the traditional theory of extreme events. However, attempts are being made to bring them into EVT. The reason for the demise of Long Term Capital Management has been put down by Dunbar (2000) to two theories. The first theory is the perfect storm. According to this theory, the events of summer 1998 were similar to a one-in-250 years hurricane. However, financial market crises are not random geophysical events, but are created by human behaviour. LTCM might be considered as an attempt to control the weather which went badly wrong, as opposed to a luckless ship being tossed by a gigantic waves, although some commentators believe that the problem was model mis-specification. If LCTM management had allowed for the divergence of sovereign debt and domestic debt as two asset classes, then it may have been less at risk. It should also be noted that any model that is based on Brownian motion, by itself, does not satisfy the Pareto assumption, and, therefore, might not be most appropriate for estimating the impact of extreme events.

Epidemics

6.20 The last major epidemic was the influenza epidemic of 1919. With global travel the potential for epidemics is considerable, for example SARS and other diseases transmitted between hosts by insects, ticks, or other invertebrae. These vector borne disease epidemics often seem to occur in conjunction with extreme climatic or weather events For example, there appears to be a correlation between El Nino and malaria transmission, A severe drought in the Western U.S.A. may have provided suitable conditions for the spread of the West Nile virus.

6.21 Population movement is another cause of epidemics. Amongst European populations, certain infectious diseases were held in check by acquired or inherited immunity. European exploration and colonisation of other lands led to the disastrous transmission of these diseases to non-resistant populations elsewhere. In Europe, such diseases may have been chronic problems, but in the New World of the Americas they became extreme events. AIDS is now following a similar epidemic path. An example of a non-human epidemic was the recent foot and mouth outbreak in the U.K.

6.22 It is unusual for extreme event models to be used in assessing epidemics, as the models depend on transmission rates. Such models are similar to Cat models. However, increasingly, actuaries are becoming interested in variations in mortality and morbidity, and EVT modelling has been used in these circumstances. This has even been extended to modelling the transmission of computer viruses over the Internet.

7. Extreme Events and Public Policy

7.1 Extreme events have undetermined consequences on society and on public policy. The profession has a key role in help formulate understanding of the issues and potential solutions. Jean Jacques Rousseau, on hearing of the destruction of Lisbon following earthquake and tsunami, asked: "Is this the price mankind must pay for civilisation?" Politicians are now aware of the consequences of extreme events and global warming. This concluding section is intended to give an overview. For more details see, for example, Gurenko (2004).

7.2 Man is increasingly building larger cities which are near to the ocean, were they are subject to windstorm exposure, on earthquake faults or close to volcanoes. Houses are being built on flood plains. As the Red Cross stated:

> "These challenges emerge in the face of societal trends that are converging to increase the likelihood, magnitude, and diversity of disasters. Growing population, migration of population to coasts and to cities, increased economic and technological interdependence, and increased environmental degradation are just a few of the interacting factors that underlie the mounting threat of disasters."
>
> Red Cross, 1999; Centre for Research on the Epidemiology of Disasters, 2001

7.3 To assess how dangerous an insurance portfolio is, it is often convenient to assess the normal Pareto parameter. The 80:20 rule states that 20% of the claims in a particular portfolio are responsible for more than 80% of the total portfolio claim amount. This is the Pareto formulation of EVT, given above. Thus, using extreme value theory, we can readily assess the Pareto parameter, and estimate how dangerous an insurance portfolio is, by revealing the presence of areas where a single event could spell financial ruin. The 80:20 rule works pretty well for many insurance segments, such as fire. However, for other segments it does not. Hurricane data in the Caribbean indicate that insurers can be making profit for a number of years, and then find themselves hit by a one-in-1,000 years hurricane, which swallows up 95% of the total cover in one go. When Hugo (1988) hit the U.S. Virgin Isles, losses were reported to be in excess of 1,000 times premiums; yet this was not a one-in-1,000 years event.

7.4 The regulators of the insurance industry generally target a one-in-100 years to a one-in-250 years insolvency level. They do not cater for the one-in-1,000 years event. Typical solvency levels for major developed markets are of the order of three to six times the cost of a once-in-a-century event.

7.5 The amount of insurance capital is limited, and the whole amount is potentially at risk, due to a single mega event. The insurance industry may be fooling itself if it believes that it can pay these types of loss. Three losses in excess of $100bn would suffice to create real problems. The CEO of AIG once commented that he did not wish to see 40 years of business destroyed by one minute of ground movement; and many CEOs will agree. One single event nightmare scenario is a supercyclone (one with sustained wind speeds in excess of 240 kilometres an hour), smashing into either Manhattan or Tokyo. This could generate claims approaching the total funds available for reinsurers,

7.6 Once an extreme event occurs, the (re)insurance companies have a limited number of options. The general first reaction is that virtually all insurance companies announce the first estimate of their portion of the loss, and, more importantly, have sufficient reserves and capital to pay all expected claims. The insurer then has the following options:

- to withdraw from the market;
- to write the policy, but with exclusions, deductibles and limits in the event of an major loss;
- to apply a (substantial) extra premium to buy additional cover; or
- to seek protection from government, through a guarantee pool, to cover such a loss, or future loss (for example Pool Re).

7.7 This has lead, in many countries, to where the costs of extreme events are not met by the insurance industry, but by Government. These are dealt with in Gurenko (2004), and some of the arrangements are summarised below.

Government Schemes

7.8 The U.S. Government has a complex system for dealing with natural disasters through a 'partnership' with state and local governments, non-governmental organisations, and the private sector, and involves 50 plus laws. Government may also undercut the goals of risk mitigation, by sponsoring development and redevelopment in areas of known risks, including flood plains or brownfield sites. The government schemes in the U.S.A. are as follows.

7.8.1 *Earthquake insurance in California*

The Government of California requires private companies doing business in California to offer earthquake insurance and to contribute to the funding of the California Earthquake Authority (CEA), which underwrites these policies.

7.8.2 *Homeowners' insurance in Florida*

The Government of Florida has required private companies to continue writing homeowners' policies in the state, and to participate in various residual market mechanisms, as a way of making hurricane coverage available.

7.8.3 *Flood insurance*

The Federal Government offers flood insurance through the National Flood Insurance Program (NFIP). Property owners with existing structures inside the flood plain are charged 'non-actuarial' rates, which create an implicit subsidy.

7.8.4 *Crop insurance*

The Federal Government offers farmers subsidised crop insurance, which can be triggered by natural disasters, such as flooding.

7.8.5 *Direct aid*

Emergency aid comes from government agencies and government employees at the time of the disaster and immediately following it. There are federal funding to repair state and local government facilities, loans and grants from the small business administration, and grants to individuals from FEMA and occasional assistance to flooded-out farmers, whether or not they purchased crop insurance.

7.9 Other countries have also have systems in place.

7.9.1 *Spain (Consorcio; est. 1954*)

There is cover from a typical windstorm, earthquake, volcanic eruption, a typical flood, meteorite, terrorism, rebellion, sedition, riot, civil commotion and actions of security forces in peacetime.

7.9.2 *U.K.*

Terrorism is covered via Pool Re, established after the St Mary Axe bomb.

7.9.3 *France*

There are the National Fund for Agricultural Disasters (1964 law), the GAREAT pool for terrorism risks, and the natural catastrophe scheme (1982 law).

7.10 The main reasons why governments are involved is that catastrophe risk, by its nature, is a highly correlated risk, resulting in many people having claims. The pooling of correlated risk increases the variability of risk, which is exactly opposite the fundamental premise of insurance, namely the law of large numbers. Thus, the advantage of private insurance is lost. Techniques have been, and are being, devised to mitigate correlated risk, but the costs of private solutions to some events are high. Thus, some form of government involvement is needed to keep the cost manageable.

7.11 Against the creation of government schemes, the potential political payoff from an investment in a national catastrophe risk management programme may come too late, due to the long return periods of catastrophic events. In the absence of a major recent event, creating such a programme may prove highly expensive and unpopular. "It will never happen in my lifetime!" However, once a major event occurs, then it is often a trigger to do something. The recent Indian Ocean tsunami is a good example.

7.12 Building a national catastrophe risk management programme will divert Government funds from investments in other potentially more economically productive projects. Furthermore, in developing countries, this may undermine the receipt of more aid, which would even further diminish the return on an investment in a national Cat fund. Why pay for protecting against something remote and abstract, when there are other more urgent spending priorities?!

7.13 Public policy is often seen as a linear cause and effect process. However, there is a complex dynamic underlying the process. This complexity increases the challenge to the scientific understanding and anticipation of extreme events and their consequences. We, as actuaries, have a clear role to play in helping public policy decisions, both as to predictability and to understanding. There are events that are both predictable and understandable, such as floods and hurricane landfalls. The models which we use are generally good in most circumstances. On the other hand, the development of epidemics, such as AIDS, tend to be predictable, but the real mechanics of disease transmission are not well understood. Earthquakes are understood, being described by power laws, yet are not predictable. These are susceptible to actuarial techniques for handling the uncertainty.

7.14 Certain extreme events are neither understood or predictable. Climate change is currently one example, with conflicting models and theories. Many systems are much too complex to model, and have implicit chaotic structures. A recent headline started: "London could be among the 'first cities to go' if global warming causes the planet's ice to melt, the U.K. government's chief scientific adviser has warned." In addition to London, cities such as New York and Tokyo are also vulnerable, yet the reaction of Government appears minimal. The most common outcome of these various forces at play is a political and social deadlock, over when and how to move forward towards a more sustainable risk management system.

7.15 It is clear that, in order that meaningful and constructive decisions can be made regarding extreme events, the types of process and the uncertainty need to be understood. Science, itself, is restricted to the development of predictions and generalisations, supported by quantitative data and formal deductive methods, and judged through a scientific peer review process. Actuarial science adds a further important layer, with the concept of uncertainty on the predictions, and placing a cost value on that uncertainty.

7.16 We are moving to a more complex society, with mega cities and mega organisations. This is increasing the possibility of increased costs to society of extreme events, due to increased value at risk and increased frequency, due to the larger organisational risk.

7.17 In terms of extreme events, these decisions will almost certainly involve public policy statements. The profession's views on extreme events and their impact on society need to be seen and heard, as part of the profession's contribution to making financial sense out of the future.

Acknowledgements

I would like to thank a number of people who contributed to the earlier GIRO papers, Sanders *et al.* (2002) and Sanders (2003), and to Professor Paul Embrechts, who first introduced me to the theory many years ago.

References

The following list includes, not only the works referred to in the paper, but also details of other publications which might prove to be helpful to the reader.

Beard, R.E. (1962). Statistical theory of extreme values and an application to excess of loss reinsurance. (mimeographed MS, 1962)

Bak, P. (1996). *How nature works: the science of self-organised criticality.* Copernicus, New York.

Bak, P., Tang, C. & Wiesenfeld, K. (1987). Self-organized criticality: an explanation of 1/f noise. *Phys. Rev. Lett.*, **59**, 381-384.

Beelders, M. & Colarossi, D. (2004). Modelling mortality risk with extreme value theory. *GARP*, Issue 19.

Coles, S.G. (2001). *An introduction to statistical modelling of extreme values.* Springer Verlag, New York.

Dessai, S. & Walter, M.E. (2003). Self-organised criticality and the atmospheric sciences: selected review, new findings and future directions.

Dunbar, N. (2001). *Inventing money.* Wiley.

EMBRECHTS, P., DANIELLSON, J., GOODHART, C., KEATING, C., MUENNING, F., RENUALT, O. & SHIN, H.S. (2001). An academic response to Basel II. Financial Markets Group, London School of Economics.

EMBRECHTS, P., KLUPPELBERG, C. & MIKOSCH, T. (1997). *Modelling extremal events for insurance and finance.* Springer, New York.

EMBRECHTS, P., LINDSKOG, F. & MCNEIL, A.J. (2003). Modelling dependence with copulas and applications to risk management. In *Handbook of heavy tailed distributions in finance*. Edited by Rachev, S.T. Elsevier/North-Holland, Amsterdam.

EMBRECHTS, P., MCNEIL, A.J. & STRAUMANN, D. (1999). Correlation: pitfalls and alternatives. *Risk*, May 1999, 69-71.

FALK, M., HUSLER, J. & REISS, R. (1994). *Laws of small numbers: extremes and rare events.* Birkhauser, Basel.

FINKENSTADT, B. & ROOTZEN, H. (2003). *Extreme values in finance, telecommunications, and the environment.* Chapman and Hall.

FISHER, R.A. & TIPPETT, L.H.C. (1928). Limiting forms of the frequency distributions of the largest or smallest member of a sample. *Proceedings of the Cambridge Philosophical Society*, **24**, 180-190.

FOUGERES, A.-L. (2002). Multivariate extremes. In Finkenstadt & Rootzen, 373-388.

GNEDENKO, B.V. (1943). Sur la distribution limite du terme maximum d'une serie aleatoire. *Annals of Mathematics*, **44**, 423-453.

GUMBEL, E.J. (1958). *Statistics of extremes.* Columbia University Press.

GURENKO, E.N. (ed.) (2004). *Catastrophe risk and reinsurance.* Risk Books.

HOGG, R. & KLUGMAN, S. (1984). *Loss distributions.* Wiley, New York.

KLEIBER, K. & KOTZ, S. (2003). *Statistical size distributions in economics and actuarial sciences.* John Wiley, New Jersey.

MCNEIL, A. (1997). Estimating the tails of loss severity distributions using extreme value theory, *ASTIN Bulletin*, **27**, 117-137.

PICKANDS, J. (1975). Statistical inference using extreme order statistics. *Annals of Statistics*, **3**, 119-131.

REISS, R. & THOMAS, M. (1997). Statistical analysis of extreme values. Birkhauser (includes documentation for Xtremes software package).

RESNICK, S. (1987). *Extreme values, point processes and regular variation.* Springer Verlag, New York.

SANDERS, D.E.A. (1995). When the wind blows, an introduction to catastrophe excess of loss insurance. C.A.S. Fall Forum 1995.

SANDERS, D.E.A. (2003). Extreme events part 2. Financial catastrophes. The overthrow of modern financial theory. GIRO 2003.

SANDERS, D.E.A. *et al.* (1995). Pricing in the London Market. GIRO 1995.

SANDERS, D.E.A. *et al.* (2002). The management of losses arising from extreme events. GIRO 2002.

SMITH, R.L. (1986). Extreme value theory based on the largest annual events. *Journal of Hydrology*, **86**, 27-43.

SMITH, R.L. (2003). Statistics of extremes with applications in environment, insurance and finance. In Finkenstadt & Rootzen, 1-68.

SMITH, R.L. & GOODMAN, D. (2000). Bayesian risk analysis. Chapter 17 of *Extremes and integrated risk management.* Edited by P. Embrechts. Risk Books, London.

THOM, R. (1989). *Structural stability and morphogenesis.* Addison Wesley Publishing Company.

WOO, G. (1999). *The mathematics of natural catastrophes.* Imperial College Press.

APPENDIX 1

SOFTWARE AND SPREADSHEETS

The spreadsheet accompanying the 2002 GISG paper may be downloaded from http://www.actuaries.org.uk/files/giro2002/Sanders.zip

An academic version of Xtremes which accompanies the examples in Reiss & Thomas (1997) may be downloaded from http://www.xtremes.de/xtremes/

This software contains actuarial elements other than extreme events.

S-Plus is commercial software. As part of Coles (2001), there are accompanying sub-routines which may be incorporated in S-plus. These may be found at http://www.maths.bris.ac.uk/~masgc/ismev/functions.html

Alexander McNeil has also written Extreme software for S-Plus. This is called EVIs, and may be found at http://www.math.ethz.ch/~mcneil/software.html

Printed by Bell & Bain Ltd., Glasgow