

How to model multivariate extremes if one must?

Thomas Mikosch*

*Laboratory of Actuarial Mathematics, University of Copenhagen,
Universitetsparken 5, DK-2100 Copenhagen, Denmark*

In this paper we discuss some approaches to modeling extremely large values in multivariate time series. In particular, we discuss the notion of multivariate regular variation as a key to modeling multivariate heavy-tailed phenomena. The latter notion has found a variety of applications in queuing theory, stochastic networks, telecommunications, insurance, finance and other areas. We contrast this approach with modeling multivariate extremes by using the multivariate student distribution and copulas.

Key Words and Phrases: multivariate regular variation, heavy-tailed distribution, extreme value distribution, copula, elliptical distribution.

1 Introduction

Over the last few years heavy-tailed phenomena have attracted a lot of attention. These include turbulence and crashes in the financial and insurance markets, but also strong deviations in weather and climate phenomena from the average. More recently, the Internet and more generally the enormous increase of computer power have led to collections of huge data sets which cannot be handled by classical statistical methods. Among others, teletraffic data (such as on/off times of computers, length and transfer times of files, etc.) exhibit not only a non-standard dependence structure which cannot be described by the methods of classical time series analysis, but these data also have clusters of unusually large data. It has been recognized early on in hydrology and meteorology, but also in insurance practice that the distributions of the classical statistical theory (such as the normal and the gamma family) are of restricted use for modeling the data at hand. Describing these data by their median, expectation, variance or by moment related quantities such as the kurtosis and skewness has rather limited value in this context. For the actuary it is not a priori of interest to know what the expectation and the variance of the data

*Thomas Mikosch's research is partially supported by MaPhySto, The Danish Research Network for Mathematical Physics and Stochastics, and DYNSTOCH, a research training network under the programme Improving Human Potential financed by The 5th Framework Programme of the European Commission. Financial support by the Danish Research Council (SNF) Grant No 21-01-0546 is gratefully acknowledged by the author.

mikosch@math.ku.dk

are, he is mainly concerned with large claims which might arise from scenarios similar to the WTC disaster. Such events are extremely rare and dangerous. It would be unwise to use the (truncated) normal or the gamma distributions to capture such an event by a mean-variance analysis. The devil sits in the tail of the distribution. It is the tail of the distribution that costs the insurance and financial industry billions of dollars (only the WTC disaster alone has cost the reinsurance industry about \$ US 20 billion up to now; see SIGMA (2003)). It is the tail of the distributions of the file sizes and transition times of files that causes the unpleasant behavior of our computer networks. Although the Internet is a traffic system which, unlike the German Autobahn, has thousands of extremely fast lanes, this system is not always able to handle the amounts of information to be transferred and, like the German Autobahn, it is subject to traffic jams caused by huge files representing movies, pictures, DVDs, CDs which make the difference. They cannot be modeled by exponential or gamma distributions, very much in contrast to classical queuing and network theory, where the exponential distribution was recognized as adequate for modeling human behavior in telephone or other consumer-service systems. Currently there is general agreement that the modern teletraffic systems are well described by distributions with tails much heavier than the exponential distribution; see e.g. WILLINGER *et al.* (1995); cf. MIKOSCH *et al.* (2003) and STEGEMAN (2002).

Early on, distributions with power law tails have been used in applications to model extremely large values. The Pareto distribution was introduced in order to describe the distribution of income in a given population (PARETO, 1896/97). Although the world has changed a lot since Pareto suggested this distribution at the end of the 19th century, it still gives a very nice fit to the world income distribution. The Pareto distribution is also a standard distribution for the purposes of reinsurance, where the largest claims of a portfolio are taken care of. The Pareto distribution in its simplest form can be written as

$$\bar{F}(x) = 1 - F(x) = (c/x)^\alpha, \quad x \geq c, \quad \text{some positive } c, \alpha.$$

As expected in general this distribution does not give a great fit to data in the center of the distribution, but it often captures the large values of the data in a convincing way. Of course, one can shift the distribution to the origin by introducing a location parameter, but the fit in the center would not improve much in this way.

The Pareto distribution appears in a completely different theoretical context, namely as the limit distribution of the excesses of an iid sequence X_1, \dots, X_n with distribution F above a high threshold. To be more precise, the only limit distribution of the excess distribution of the X_i 's is necessarily of the form (up to changes of location and scale)

$$\lim_{u \uparrow x_F, u+x < x_F} P(X_1 - u > x \mid X_1 > u) \rightarrow (1 + \xi x)_+^{-1/\xi} = \bar{G}_\xi(x), \quad (1)$$

$$x \in \mathbb{R},$$

(PICKANDS, 1975; BALKEMA and DE HAAN, 1974; cf. EMBRECHTS *et al.*, 1997, Section 3.4), where

$$x_F = \sup\{x \in \mathbb{R} : F(x) < 1\}$$

is the right endpoint of the distribution F and $\xi \in \mathbb{R}$ is a shape parameter. For $\xi = 0$ the limit has to be interpreted as the tail of the standard exponential distribution. The so defined limit distribution G_ξ is called the *generalized Pareto distribution* (GPD). Relation (1) holds only for a restricted class of distributions F . Indeed, (1) is satisfied if and only if a limit relation of the following type holds for suitable constants $d_n \in \mathbb{R}$, $c_n > 0$, and the partial maxima $M_n = \max(X_1, \dots, X_n)$ (see EMBRECHTS *et al.*, 1997, Chapter 3):

$$P(c_n^{-1}(M_n - d_n) \leq x) \rightarrow \exp\left\{-(1 + \xi x)_+^{-1/\xi}\right\} = H_\xi(x), \quad n \rightarrow \infty, \quad x \in \mathbb{R}. \quad (2)$$

For $\xi = 0$ the distribution has to be interpreted as the Gumbel distribution $H_0(x) = e^{-e^{-x}}$. The limit distribution is called the *generalized extreme value distribution* and (up to changes of scale and location) it is the only possible non-degenerate limit distribution for centered and normalized maxima of iid sequences. We say that the underlying distribution F belongs to the *maximum domain of attraction of the extreme value distribution* H_ξ ($F \in \text{MDA}(H_\xi)$). The case $\xi > 0$ is particularly interesting for modeling extremes with unlimited values. Then the extreme value distribution H_ξ can be reparametrized and written as the so-called *Fréchet distribution* with $\alpha = \xi^{-1}$ (Figure 1):

$$\Phi_\alpha(x) = e^{-x^{-\alpha}}, \quad x > 0.$$

Every distribution $F \in \text{MDA}(\Phi_\alpha)$ is completely characterized by the relation

$$\bar{F}(x) = 1 - F(x) = \frac{L(x)}{x^\alpha}, \quad x > 0, \quad (3)$$

where L is a *slowly varying* function, i.e., L is a positive function on $(0, \infty)$ with property $L(cx)/L(x) \rightarrow 1$ as $x \rightarrow \infty$ for every $c > 0$. Notice that distributional tails




Fig. 1. Left: Plot of 9558 S&P500 daily log-returns from 2 January 1953 to 31 December 1990. The year marks indicate the beginning of the calendar year. Right: QQ-plot of the S&P500 data against the normal distribution whose mean and variance are estimated from the data. The data come from a distribution which has much heavier left and right tail than the normal distribution.

of type (3) are a generalization of distributions with pure power law tails such as the Pareto distribution. It is a semiparametric description of a large class of distributions; the slowly varying functions L represent a nuisance parameter which is not further specified. This is very much in agreement with real-life data analysis where it is hard to believe that the data come from a pure Pareto distribution. In particular, L is not specified in any finite interval which leaves the question about the form of the distribution F in its center open. Several distributions with a name have *regularly varying* right tail, i.e., (3) holds, e.g. the Pareto, Burr, log-gamma, Student, Cauchy, Fréchet and infinite variance stable distributions; see EMBRECHTS *et al.* (1997), p. 35, for definitions of these distributions.

The stable distributions consist of the only possible non-degenerate limit distributions H for the partial sums $S_n = X_1 + \cdots + X_n$ of an iid sequence (X_i) with distribution F , i.e., there exist constants $c_n > 0$, $d_n \in \mathbb{R}$, such that

$$\lim_{x \rightarrow \infty} P(c_n^{-1}(S_n - d_n) \leq x) = H(x), \quad x \in \mathbb{R}.$$

We say that F belongs to the *domain of attraction* of the stable distribution H ($F \in \text{DA}(H)$). The best known stable distribution is the normal whose domain of attraction contains all F with slowly varying truncated second moment $\int_{|y| \leq x} y^2 dF(y)$ (FELLER, 1971), i.e., it contains almost all distributions of interest in statistics. The remaining stable distributions are less known; they have infinite variance and so are the members of their domains of attraction. In particular, every infinite variance stable distribution is characterized by a shape parameter $\alpha \in (0, 2)$ which appears as the tail parameter of these distributions H_α . Moreover, it also appears in the tails of distributions $F \in \text{DA}(H_\alpha)$:

$$F(-x) \sim p \frac{L(x)}{x^\alpha} \text{ and } \bar{F}(x) \sim q \frac{L(x)}{x^\alpha}, \quad x > 0, \quad (4)$$

where L is slowly varying and $p, q \geq 0$ such that $p + q = 1$. Relation (4) is also referred to as *tail balance condition* and F is said to be *regularly varying with index* $\alpha \in (0, 2)$.

Regular variation also occurs in a surprising way in solutions to stochastic recurrence equations. We consider here the simplest one-dimensional case. Assume the *stochastic recurrence equation*

$$Y_t = A_t Y_{t-1} + B_t, \quad t \in \mathbb{Z}, \quad (5)$$

has a strictly stationary causal solution, where $((A_t, B_t))$ constitute an iid sequence of non-negative random variables. Causality refers to the fact that Y_t is a function only of (A_s, B_s) , $s \leq t$. A sufficient condition for the existence of such a solution is given by $E \log^+ A_0 < \infty$, $E \log^+ B_0 < \infty$ and $E \log A_0 < 0$. Equations of type (5) occur in the context of financial time series models. For example, the celebrated (2003 Nobel prize winning) ARCH and GARCH models of ENGLE (1982) and BOLLERSLEV (1986) can be embedded in a stochastic recurrence equation (Figure 2). We illustrate




Fig. 2. 2493 Danish fire insurance claims in Danish Kroner from the period 1980–1992. The data (left) and a QQ-plot of the data against standard exponential quantiles (right). If the data came from a standard exponential distribution, the curve should be close to the straight line in the QQ-plot. Since the graph curves down, one may conclude that the data have a much heavier tail than the exponential distribution.

this with the GARCH(1,1) model (*generalized autoregressive conditionally heteroscedastic model of order (1,1)*) which is given by the equation

$$X_t = \sigma_t Z_t, \quad \sigma_t^2 = \alpha_0 + \alpha_1 X_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \quad t \in \mathbb{Z}.$$

Here (Z_t) is an iid sequence with $EZ_1 = 0$ and $\text{var}(Z_1) = 1$ and $\alpha_0 > 0$, α_1, β_1 are non-negative parameters. Obviously, $Y_t = \sigma_t^2$ satisfies the stochastic recurrence equation (5) with $A_t = \alpha_1 Z_{t-1}^2 + \beta_1$ and $B_t = \alpha_0$. An important result by KESTEN (1973) (see also GOLDIE, 1991) says that under general conditions on the distribution of A_0 the equation

$$EA_0^\kappa = 1 \tag{6}$$

has a unique positive solution κ_1 and then for some $c > 0$,

$$P(Y_0 > x) \sim cx^{-\kappa_1}, \quad x \rightarrow \infty. \tag{7}$$

In particular, the mentioned ARCH and GARCH models have marginal distribution with regularly varying tail of type (7). For a GARCH(1,1) model, (6) turns into

$$E(\alpha_1 Z_0^2 + \beta_1)^\kappa = 1,$$

which has a solution κ_1 , e.g. when Z_0 is normally distributed (Figure 3). Hence $Y_0 = \sigma_0^2$ satisfies (7) and a standard argument on regular variation implies that $P(X_0 > x) \sim c'x^{-\kappa_1/2}$. This is a rather surprising result which says that light-tailed input (noise) can cause heavy-tailed output in a non-linear time series. Such a result is impossible for linear processes (such as ARMA processes) driven by iid noise; see e.g. EMBRECHTS *et al.* (1997), Appendix A3, or MIKOSCH and SAMORODNITSKY (2000). We refer to BASRAK *et al.* (2002a,b) as general references on GARCH and regular variation, to EMBRECHTS *et al.* (1997), Section 8.4, for an introduction to stochastic recurrence equations and the tails of their solutions. See also MIKOSCH




Fig. 3. A time series of measured file sizes handled by a webserver. The data (left) and a QQ-plot of the data against standard exponential quantiles (right). The data have tail much heavier than the exponential distribution.

(2003) for a survey paper on financial time series models, their extremes and regular variation.

The conclusion of this introduction should be that one-dimensional distributions with regularly varying tails are very natural for modeling extremal events when large values are involved. Since there exists some theoretical background as to why these distributions occur in different contexts, it is only consequent to fit them to real-life data. On the other hand, distributions such as the gamma, the exponential or the normal distributions are less appropriate for fitting extremes when large values occur.

The next section discusses multivariate regular variation as a suitable tool for modeling multivariate extremes. We continue in Section 3 by discussing some of the alternative approaches such as the multivariate Student distribution and copulas for modeling multivariate data with very large values.

2 Multivariate regularly varying distributions

Over the last few years there has been some search for multivariate distributions which might be appropriate for modeling very large values such as present in financial or insurance portfolios. The multivariate Gaussian distribution is not an appropriate tool in this context since there is strong evidence that heavy-tailed marginal distributions are present. Nevertheless it has become standard in risk management to apply the multivariate Gaussian distribution, e.g. for calculating the Value at Risk (VaR). A major reason for the use of the normal distribution is its “simplicity”: linear combinations of the components of a normally distributed vector \mathbf{X} are normal and the distribution of \mathbf{X} is completely determined by its mean and covariance structure.

It is the aim of this section to introduce multivariate distributions which have, in a sense, power law tails in all directions. We will also indicate that these distributions appear in a natural way, e.g. as domain of attraction conditions for weakly converging partial sums and componentwise maxima of iid vectors.

To start with, we rewrite the defining property of a one-dimensional regularly varying distribution F (see (4)) as follows: for every $t > 0$,

$$\lim_{x \rightarrow \infty} \frac{P(x^{-1}X_1 \in (t, \infty])}{P(|X_1| > x)} = qt^{-\alpha} = \mu(t, \infty], \quad (8)$$

$$\lim_{x \rightarrow \infty} \frac{P(x^{-1}X_1 \in [-\infty, -t])}{P(|X_1| > x)} = pt^{-\alpha} = \mu[-\infty, -t]. \quad (9)$$

These relations are immediate from the properties of regularly varying functions; see e.g. BINGHAM *et al.* (1987). Notice that the right hand expressions can be interpreted as the μ -measure of the sets $(t, \infty]$ and $[-\infty, -t]$, where μ is defined on the Borel sets of $\mathbb{R} \setminus \{0\}$:

$$d\mu(x) = \alpha[p|x|^{-\alpha-1}I_{[-\infty, 0)}(x) + qx^{-\alpha-1}I_{(0, \infty)}(x)]dx.$$

Relations (8) and (9) can be understood as convergence of measures:

$$\mu_x(\cdot) = \frac{P(x^{-1}X_1 \in \cdot)}{P(|X_1| > x)} \xrightarrow{v} \mu(\cdot), \quad x \rightarrow \infty, \quad (10)$$

where \xrightarrow{v} refers to *vague convergence* on the Borel σ -field of $\mathbb{R} \setminus \{0\}$. This simply means in our context that $\mu_x(A) \rightarrow \mu(A)$ for every Borel set $A \subset \mathbb{R} \setminus \{0\}$ which is bounded away from zero and satisfies $\mu(\partial A) = 0$. We refer to KALLENBERG (1983) or RESNICK (1987) for the definition and properties of vaguely converging measures.

Relation (10) allows one to extend the notion of regular variation to Euclidean space. Indeed, we say that the vector \mathbf{X} with values in \mathbb{R}^d and its distribution are *regularly varying with limiting measure μ* if the relation

$$\mu_x(\cdot) = \frac{P(x^{-1}\mathbf{X} \in \cdot)}{P(|\mathbf{X}| > x)} \xrightarrow{v} \mu(\cdot), \quad x \rightarrow \infty, \quad (11)$$

holds for a non-null measure μ on the Borel σ -field of $\mathbb{R}^d \setminus \{\mathbf{0}\}$. Again, this relation means nothing but $\mu_x(A) \rightarrow \mu(A)$ for any set $A \subset \mathbb{R}^d \setminus \{\mathbf{0}\}$ which is bounded away from zero and satisfies $\mu(\partial A) = 0$. Regular variation of \mathbf{X} implies regular variation of $|\mathbf{X}|$ with a positive index α and therefore

$$\mu_x(tA) = \frac{P(x^{-1}\mathbf{X} \in tA) P(|\mathbf{X}| > tx)}{P(|\mathbf{X}| > tx) P(|\mathbf{X}| > x)} \rightarrow \mu(A)t^{-\alpha}.$$

This means that μ satisfies the *homogeneity property* $\mu(tA) = t^{-\alpha}\mu(A)$, and we therefore also say that \mathbf{X} is *regularly varying with index α* .

Now define the sets

$$A(t, S) = \{\mathbf{x} : |\mathbf{x}| > t, \tilde{\mathbf{x}} \in S\},$$

where $t > 0$ and $S \subset \mathbb{S}^{d-1}$, the unit sphere of \mathbb{R}^d with respect to a given norm, and

$$\tilde{\mathbf{X}} = \mathbf{X}/|\mathbf{X}|.$$

The defining property of regular variation also implies that

$$\mu_x(A(t, S)) = \frac{P(|\mathbf{X}| > tx, \tilde{\mathbf{X}} \in S)}{P(|\mathbf{X}| > x)} \rightarrow \mu(A(t, S)) = t^{-\alpha} \mu(A(1, S)), \quad (12)$$

where we assume that $\mu(\partial A(1, S)) = 0$. Since sets of the form $A(t, S)$ generate vague convergence in $\bar{\mathbb{R}}^d \setminus \{\mathbf{0}\}$, (11) and (12) both define regular variation of \mathbf{X} . The family of the values $\mu(A(1, S))$, for any Borel set $S \subset \mathbb{S}^{d-1}$, defines a probability measure $P(\Theta \in \cdot)$ on the Borel σ -field of \mathbb{S}^{d-1} , the so-called *spectral measure* of μ . The spectral measure $P(\Theta \in \cdot)$ and the index $\alpha > 0$ completely determine the measure μ .

The notion of multivariate regular variation is a natural extension of one-dimensional regular variation. Indeed, regular variation of iid \mathbf{X}_i 's with index $\alpha \in (0, 2)$ is equivalent to the property that centered and normalized partial sums $\mathbf{X}_1 + \dots + \mathbf{X}_n$ converge in distribution to an α -stable random vector; see Rvačeva (1962). Moreover, the properly normalized and centered componentwise partial maxima of the \mathbf{X}_i 's, i.e.,

$$(\max_{i=1, \dots, n} X_i^{(1)}, \dots, \max_{i=1, \dots, n} X_i^{(d)})$$

converge in distribution to a multivariate extreme value distribution whose marginals are of the type Fréchet Φ_α for some $\alpha > 0$ if and only if the vector of the componentwise positive parts of \mathbf{X}_1 is regularly varying with index α ; see DE HAAN and RESNICK (1977), RESNICK (1987). Moreover, KESTEN (1973) proves that, under general conditions, a unique strictly stationary solution to the d -dimensional stochastic recurrence equation

$$\mathbf{Y}_t = \mathbf{A}_t \mathbf{Y}_{t-1} + \mathbf{B}_t, \quad t \in \mathbb{Z},$$

exists and satisfies

$$P((\mathbf{x}, \mathbf{Y}_0) > x) \sim c(\mathbf{x})x^{-\kappa_1}, \quad x \rightarrow \infty, \quad (13)$$

for some positive $c(\mathbf{x})$, κ_1 , provided $\mathbf{x} \neq \mathbf{0}$ has non-negative components. Here $((\mathbf{A}_t, \mathbf{B}_t))$ is an iid sequence, where \mathbf{A}_t are $d \times d$ matrices and \mathbf{B}_t are d -dimensional vectors, both with non-negative entries. Kesten's result implies in particular, that the one-dimensional marginal distribution of a GARCH process is regularly varying; see the discussion in Section 1. Unfortunately, the definition of multivariate regular variation of a vector \mathbf{Y}_0 in the sense of (11) or (12) and the definition via linear combinations of the components of \mathbf{Y}_0 in the sense of (13) are in general not known to be equivalent; see BASRAK *et al.* (2002a), HULT (2003).

Well known multivariate regularly varying distributions include the multivariate Student, Cauchy and F -distributions as well as the extreme value distributions with Fréchet marginals and the multivariate α -stable distributions with $\alpha \in (0, 2)$. In the

context of extreme value theory for multivariate data the spherical representation of the limiting measure μ (see (12)) has a nice interpretation. Indeed, we see that \mathbf{X} is regularly varying with index α and spectral measure $P(\boldsymbol{\Theta} \in \cdot)$ if and only if for sets $S \subset \mathbb{S}^{d-1}$ with $P(\boldsymbol{\Theta} \in \partial S) = 0$ and $t > 0$,

$$\lim_{x \rightarrow \infty} \frac{P(|\mathbf{X}| > tx)}{P(|\mathbf{X}| > x)} = t^{-\alpha} \text{ and } \lim_{x \rightarrow \infty} P(x^{-1}\mathbf{X} \in S \mid |\mathbf{X}| > x) = P(\boldsymbol{\Theta} \in S).$$

This means that the radial and the spherical parts of a regularly varying vector \mathbf{X} become “independent” for values $|\mathbf{X}|$ far away from the origin. For an iid sample of multivariate regularly varying \mathbf{X}_i ’s the spectral measure tells us about the likelihood of the directions of those \mathbf{X}_i ’s which are farthest away from zero.

The spectral measure can be estimated from data; we refer to DE HAAN and RESNICK (1993), DE HAAN and DE RONDE (1998) and EINMAHL *et al.* (2001). For two-dimensional vectors \mathbf{X}_t we can write $\boldsymbol{\Theta} = (\cos(\Phi), \sin(\Phi))$ and estimate the distribution of Φ on $[-\pi, \pi]$. Assuming a density of Φ , one can estimate this *spectral density*; see e.g. Figures 4 and 5 for some attempts. Figure 4 shows the typical shape of a spectral measure when the components of \mathbf{X}_t are independent. Then the spectral measure is concentrated at the intersection of the unit sphere with the axes, i.e., the spectral measure is discrete. This is in contrast to the case of extremal dependence when the spectral measure is concentrated not only at the intersection with the axes. Figure 5 gives a typical shape of a spectral density where the components of the vectors \mathbf{X}_t exhibit extremal dependence. The two valleys at the angles $-\pi/4$ and $3\pi/4$ show that X_{t-1} and X_t are not both extreme when they have different signs. However, the peak at $-3\pi/4$ shows that X_{t-1} and X_t are quite likely to be extreme and negative. There is a peak at zero indicating that X_{t-1} can be large and positive, whereas X_t is “less extreme” the next day.

The aim of this section was to explain that the notion of regularly varying distribution is very natural in the context of extreme value theory. As the notion of one-dimensional regular variation, it arises as a domain of attraction condition




Fig. 4. Left: Scatterplot $\mathbf{X}_t = (X_{t-1}, X_t)$ of the teletraffic data from Figure 3. Right: Estimated spectral density on $[0, \pi/2]$. The density is estimated from the values \mathbf{X}_t with $|\mathbf{X}_t| > 80000$. The density has two clear peaks at the angles 0 and $\pi/2$, indicating that the components of \mathbf{X}_t behave like iid components far away from 0.




Fig. 5. Left: Scatterplot (X_{t-1}, X_t) of the S&P500 data from Figure 1. Right: Estimated spectral density on $[-\pi, \pi]$. The vertical lines indicate multiples of $\pi/4$.

for partial maxima and partial sums of iid vectors \mathbf{X}_t . There exist, however, enormous statistical problems if one wants to fit regularly varying distributions. Among others, one needs large sample sizes (thousands of data, say) in order to come up with reasonable statistical answers. Successful applications of multivariate extreme value theory have been conducted in dimensions two and three; see e.g. the survey paper by DE HAAN and DE RONDE (1998) for hydrological applications. The limitations of the method are due to the fact that one has to estimate multivariate measures.

In the one-dimensional case, the generalized extreme value distribution (see (2)) is the only limit distribution for normalized and centered partial maxima of iid data. This is equivalent to the weak convergence of the excesses to the generalized Pareto distribution. Modern extreme value statistics focuses mostly on fitting the generalized Pareto distribution (GPD) via the excesses (so-called POT – peaks over threshold method; see EMBRECHTS *et al.* (1997), Chapter 6). Although desirable, in the multivariate case a general result of BALKEMA–DE HAAN–PICKANDS type (see (1)) is not available as yet, but see TAJVIDI (1996) and BALKEMA and EMBRECHTS (2004) for some approaches.

We mention in passing that the notion of multivariate regular variation also allows for extensions to function spaces. DE HAAN and LIN (2002) applied regular variation in the space $\mathbb{D}[0, 1]$ of càdlàg functions on $[0, 1]$ in order to describe the extremal behavior of continuous stochastic processes. HULT and LINDSKOG (2004) extended these results to càdlàg jump functions with the aim of describing the extremal behavior of, among others, Lévy processes with regularly varying Lévy measure.

3 A discussion of some alternative approaches to multivariate extremes

Various other classes of multivariate distributions have been proposed in the context of risk management with the aim of finding a realistic “heavy-tailed” distribution.

Those include the *multivariate student distribution*; see e.g. GLASSERMAN (2004), Section 9.3. As mentioned above, these distributions are regularly varying. Their spectral measure is completely determined by their covariance structure; any Student distribution has representation in law

$$\mathbf{X} \stackrel{d}{=} \frac{1}{\sqrt{\chi_d^2/d}} A \tilde{\mathbf{Z}},$$

where $AA' = \Sigma$ is a covariance matrix (for $d > 2$ the covariance matrix of \mathbf{X} exists and is given by $d\Sigma/(d-2)$), χ_d^2 is χ^2 -distributed with d degrees of freedom and \mathbf{Z} consists of iid standard normal random variables. Moreover, \mathbf{Z} and χ_d^2 are independent. This may be attractive as regards the statistical properties of such distributions, but it is questionable whether the extremal dependence structure of financial data is determined by covariances. The Student distribution has a rather limited flexibility as regards modeling the directions of the extremes.

A second approach using *copulas* has been suggested. For simplicity, assume in the sequel that the vector $\mathbf{X} = (X, Y) \stackrel{d}{=} \mathbf{X}_t$ is two-dimensional and its components have continuous distributions F_X and F_Y , respectively. Define for any distribution function G its *quantile function* by

$$G^-(t) = \inf\{x \in \mathbb{R} : G(x) \geq t\}, \quad t \in (0, 1).$$

Then

$$P(X \leq F_X^-(x), Y \leq F_Y^-(y)) = C(x, y),$$

is a distribution on $(0, 1)^2$, referred to as the copula of (X, Y) . Of course, $F_X(X)$ and $F_Y(Y)$ are uniformly distributed on $(0, 1)$ and the dependence between X and Y sits in the copula function C . Copulas have been used in extreme value statistics for several decades; see e.g. the survey paper by DE HAAN and DE RONDE (1998) or the monograph by GALAMBOS (1987). The purpose of copulas in extreme value statistics is to transform the marginals of the vectors \mathbf{X}_t to distributions with comparable size; otherwise the extremal behavior of an iid sequence (\mathbf{X}_t) would be determined only by the extremes of one dominating component. (Another standard method is to transform the marginals of \mathbf{X}_t to standard Fréchet (Φ_1) marginals.) Thus the transformation of the data to equal marginals makes statistical sense if we cannot be sure that the marginals have comparable tails.

It is, however, wishful thinking if one believes that copulas help one to simplify the statistical analysis of multivariate extremes. The main problem in multivariate extreme value statistics for data with very large values is the estimation of the spectral measure (which gives one the likelihood of the directions of extremes) and the index of regular variation (which gives one the likelihood of the distance from the origin where extremes occur) or, equivalently, the estimation of the measure μ . It is an illusion to believe that one can estimate these quantities “in a simpler way” by introducing copulas.

We list here various problems which arise from the use of copulas.

1. *What is a reasonable choice of copula?* If one accepts that the extreme value theory outlined above makes some sense, one should search for copulas which correspond to multivariate extreme value distributions or multivariate regularly varying distribution. Since copulas stand for any dependence structure between X and Y it is not a priori clear which copulas obey this property. The choice of some ad hoc copula such as the popular arithmetic copula is completely arbitrary.
2. *Should one use extreme value copulas?* So-called *extreme value copulas* have been suggested as possible candidates for copulas for modeling extremes, such as the popular Gumbel copula. These copulas are obtained by transforming the marginals of some very specific parametric multivariate extreme value distribution to the unit cube, imposing some very specific parametric dependence structure on the extremes, very much in the spirit of assuming a multivariate student distribution as discussed above. The fit of an extreme value copula needs to be justified by verifying that the data come from an extreme value generating mechanism. For example, if we consider the annual maximal heights of sea waves at different sites along the Dutch coast, it can be reasonable to fit a multivariate extreme value distribution to these multivariate data. Then the data are extremes themselves. In other cases it is questionable to apply extreme value copulas.
3. *How do we transform the marginals to the unit cube?* Since we do not know F_X^- or F_Y^- we would have to take surrogates. The empirical quantile functions are possible candidates. However, the empirical distribution function has bounded support: its tail is zero for arguments larger than the maximum of the data. This means we would not be able to capture extremal behavior outside the range of the sample. Any other approach, for example by fitting GPDs to the marginals and inverting them to the unit cube can go wrong as well, as long as we have not got any theoretical justification for the approach. Moreover, if we are interested in a practical statistical problem we also have to back-transform the marginals from the unit cube to the original problem. There we make another error. In some cases a theoretical justification for such an approach has been given in the context of two- or three-dimensional extreme value statistics; see again the survey paper by DE HAAN and DE RONDE (1998). In the latter paper it is also mentioned that the same kind of problems arises if one wants to transform the marginals to other distributions such as the standard Fréchet distribution.
4. *Do copulas overcome the curse of dimensionality?* One can fit any parametric copula to any high-dimensional data set. As explained above, in general, one will fit a distribution which has nothing in common with the extremal structure of the data. But, given the copula is in agreement with the extremal dependence structure of the data, copula fitting faces the same problems as multivariate

extreme value statistics which so far can give honest answers for two- or three-dimensional problems.

The discussion about which multivariate non-Gaussian distributions are reasonable and mathematically tractable models for extremes is not finished. The aim of this paper was to recall that there exists a probabilistic theory for multivariate extremes that can serve as a basis for honest statistical techniques. Whether one gains by other ad hoc methods is an open question.

Acknowledgements

I would like to thank the two referees for constructive remarks and careful reading. The final version of this paper was written at the Mittag Leffler Institute in Stockholm in October 2004. I am very grateful for the hospitality I have experienced at this wonderful place.

A personal remark

From August 1994 to December 2000 I worked at the Department of Mathematics of the University of Groningen together with Willem Schaafsma and Herold Dehling. I enjoyed my time in Groningen not only because I had plenty of time for research and our few students were really very pleasant to work with, but also because Willem and Herold were very kind colleagues. Willem liked to discuss problems of statistics. Discussions with Willem were always at a philosophical level which, I admit, I sometimes could not follow. Willem's understanding of statistics was certainly different from mine which led him to call me a "rationalist". Whatever a rationalist might be, I do not feel bad about being one. The discussions with Willem were always enjoyable, even if they got heated. The day after a quarrel one could continue the same discussion. Willem has a warm personality. He was very much liked by our students. He was not always rational in situations when one should be rational, but this made one feel very comfortable in his presence. I wish Willem many more years of creative thinking and good discussions.

References

- BALKEMA, A. A. and DE HAAN, L. (1974), Residual lifetime at great age, *Annals of Probability* **2**, 792–804.
- BALKEMA, G. and P. EMBRECHTS (2004), Multivariate excess distributions, ETH Zurich, Department of Mathematic, Technical Report.
- BASRAK, B., R. A. DAVIS and T. MIKOSCH (2002a), A characterization of multivariate regular variation, *Annals of Applied Probability* **12**, 908–920.

- BASRAK, B., R. A. DAVIS and T. MIKOSCH (2002b), Regular variation of GARCH processes, *Stochastic Processes and their Applications* **99**, 95–116.
- BINGHAM, N. H., C. M. GOLDIE and J. L. TEUGELS (1987), *Regular variation*, Cambridge University Press, Cambridge.
- BOLLERSLEV, T. (1986), Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* **31**, 307–327.
- EINMAHL, J. H. J., L. DE HAAN and V. I. PITERBARG (2001), Nonparametric estimation of the spectral measure of an extreme value distribution, *Annals of Statistics* **29**, 1401–1423.
- EMBRECHTS, P., C. KLÜPPELBERG and T. MIKOSCH (1997), *Modelling extremal events for insurance and finance*, Springer, Berlin.
- ENGLE, R. F. (1982), Autoregressive conditional heteroscedastic models with estimates of the variance of United Kingdom inflation, *Econometrica* **50**, 987–1007.
- FELLER, W. (1971), *An introduction to probability theory and its applications*, vol. II, second edition, Wiley, New York.
- GALAMBOS, J. (1987), *Asymptotic theory of extreme order statistics*, 2nd edition, Krieger, Malabar (Florida).
- GLASSERMAN, P. (2004), *Monte-Carlo methods in financial engineering*, Springer, New York.
- GOLDIE, C. M. (1991), Implicit renewal theory and tails of solutions of random equations, *Annals of Applied Probability* **1**, 126–166.
- DE HAAN, L. and T. LIN (2002), On convergence toward an extreme value limit in $C[0, 1]$, *Annals of Probability* **29**, 467–483.
- DE HAAN, L. and S. I. RESNICK (1977), Limit theory for multivariate sample extremes, *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **40**, 317–337.
- DE HAAN, L. and S. I. RESNICK (1993), Estimating the limiting distribution of multivariate extremes, *Communications in Statistics: Stochastic Models* **9**, 275–309.
- DE HAAN, L. and J. DE RONDE (1998), Sea and wind: multivariate extremes at work, *Extremes* **1**, 7–45.
- HULT, H. (2003), Topics on fractional Brownian motion and regular variation for stochastic processes, PhD Thesis, KTH Stockholm, Department of Mathematics.
- HULT, H. and F. LINDSKOG (2004), Extremal behavior for regularly varying stochastic processes, *Stochastic Processes and their Applications* **115**, 249–274.
- KALLENBERG, O. (1983), *Random measures*, 3rd edition, Akademieverlag, Berlin.
- KESTEN, H. (1973), Random difference equations and renewal theory for products of random matrices, *Acta Mathematica* **131**, 207–248.
- MIKOSCH, T. and G. SAMORODNITSKY (2000), The supremum of a negative drift random walk with dependent heavy-tailed steps, *Annals of Applied Probability* **10**, 1025–1064.
- MIKOSCH, T., A. STEGEMAN, S. RESNICK and H. ROOTZÉN (2003), Is network traffic approximated by stable Lévy motion or fractional Brownian motion? *Annals of Applied Probability* **12**, 23–68.
- MIKOSCH, T. (2003), Modeling dependence and tails of financial time series, in: B. FINKENSTÄDT and H. ROOTZÉN (eds), *Extreme values in finance, telecommunications, and the environment*, Chapman & Hall, Boca Raton, pp. 185–286.
- PARETO, V. (1896/97), *Cours d'économie politique professé à l'université de Lausanne*.
- PICKANDS III, J. (1975), Statistical inference using extreme order statistics, *Annals of Statistics* **3**, 119–131.
- RESNICK, S. I. (1987), *Extreme values, regular variation, and point processes*, Springer, New York.
- RVAČEVA, E. L. (1962), On domains of attraction of multi-dimensional distributions, *Selected Translations in Mathematical Statistics and Probability Theory*, American Mathematical Society, Providence, R. I., **2**, 183–205.
- SIGMA (2003), Tables on the major losses 1970–2002, *Sigma publications No 2*, Swiss Re, Zürich, p. 34, The publication is available under <http://www.swissre.com>.

- STEGEMAN, A. (2002), *The nature of the beast*, PhD Thesis, University of Groningen, Department of Mathematics and Information Technology.
- TAJVIDI, N. (1996), *Characterization and some statistical aspects of univariate and multivariate generalised Pareto distributions*, PhD Thesis, Chalmers University Gothenburg, Department of Mathematics.
- WILLINGER, W., M. S. TAQQU, M. LELAND and D. WILSON (1995), Self-similarity through high variability: statistical analysis of Ethernet LAN traffic at the source level, *Computer Communications Review* **25**, 100–113, Proceedings of the ACM/SIGCOMM'95, Cambridge, MA.

Received: July 2004. Revised: October 2004.