THE INSTITUTE OF STATISTICS

THE CONSOLIDATED UNIVERSITY
OF NORTH CAROLINA

MULTIVARIATE THRESHOLD METHODS

by

R.L. Smith

June 1993

Mimeo Series #2304

DEPARTMENT OF STATISTICS
Chapel Hill, North Carolina

Mimeo Multivariate Threshold
Series Methods
#2304 R. L. Smith

Name	Date

Department of Statistics Library

MULTIVARIATE THRESHOLD METHODS

RICHARD L. SMITH
 Department of Statistics,
 University of North Carolina,
 Chapel Hill, N.C. 27599-3260, U.S.A.

May 11, 1993. Paper presented at the NIST/Temple University Conference on Extreme Value Theory and its Applications, Gaithersburg, May 1993.

Abstract. Statistical methods based on modeling the exceedances over a high threshold have emerged as one of the best ways to apply extreme value theory in practice. However, work on multivariate extensions – when there are several processes contributing to the extremes and we want some way of jointly modeling them – is still in a more primitive state. Nevertheless, the expansion of research on multivariate extreme value theory over the past several years (due especially to the work of J. Tawn and S. Coles) has now created a number of possibilities for developing multivariate threshold methods. In this paper, I review both univariate threshold methods and multivariate extreme value theory, and then discuss ways of putting them together. A novel point of the theory is that it suggests new ways to implement univariate threshold methods: it is possible to create models explicitly incorporating dependence by using multivariate techniques to model the joint distribution of exceedances in a time series. These results are illustrated with some examples taken from ongoing work being carried out at the National Institute of Statistical Sciences, on modeling high-level ozone exceedances as a function of time and meteorology.

KEYWORDS AND PHRASES: Generalized extreme value distribution, Generalized Pareto distribution, Maximum likelihood, Multivariate extreme value theory, Nonhomogeneous Poisson process, Ozone, Trends.

1. Threshold methods

In this paper, I review a variety of approaches to the estimation of extremal properties of a probability distribution or stochastic process. Statistical methods based on exceedances over a high threshold have gained in popularity in recent years, as compared with the much older methods based directly on the extreme value distributions. Nevertheless, there remain some critical questions about their application.

The simplest approach is based on a sequence of i.i.d. observations $X_1, \dots, X_n$, from an unknown distribution function F , in which, to confine attention to the tail, we model only the exceedances over a fixed high threshold u . That is, our information consists of N_n , the number of exceedances in n trials over the threshold u , and the excesses $Y_1, \dots, Y_{N_n}$, where

if the i 'th exceedance occurred on the j 'th trial of the original sample, the i 'th excess is defined by $Y_i = X_j - u$. It is natural to formulate the problem in terms of $\lambda = 1 - F(u)$, the probability of an exceedance over the threshold, and the conditional distribution function of the excesses,

$$F_u(y) = \frac{F(u+y) - F(u)}{1 - F(u)}. \quad (1.1)$$

It is obvious how to estimate λ (by the point estimate $\hat{\lambda} = N_n/n$ with associated confidence intervals etc.), but not at all obvious how we should deal with (1.1). Davison & Smith (1990), following earlier papers by Davison (1984), Smith (1984) and others leading back to Pickands (1975), suggested that an appropriate model for this is the *Generalized Pareto distribution* (GPD) defined by

$$G(y; \sigma, \xi) = 1 - \left(1 + \frac{\xi y}{\sigma}\right)_+^{-1/\xi}, \quad y > 0, \quad (1.2)$$

where $\sigma > 0$, ξ is any real number, and $x_+ = \max(x, 0)$. Thus the range of y is $0 < y < \infty$ for $\xi \geq 0$ and $0 < y < -\sigma/\xi$ if $\xi < 0$. The exponential distribution, $1 - e^{-y/\sigma}$, arises naturally as a limiting case when $\xi \rightarrow 0$.

The motivation for the Generalized Pareto distribution is based on Pickands' theorem that

$$\inf_{\xi} \lim_{u \uparrow \omega_F} \inf_{\sigma} \sup_{y > 0} |F_u(y) - G(y; \sigma, \xi)| = 0$$

if and only if F is in the domain of attraction of one of the standard extreme value distributions. Here $\omega_F = \sup\{x : F(x) < 1\}$ is the right-hand endpoint of F , which may be finite or infinite. In words, there exists a ξ (which does not depend on u) and a σ (which does) such that F_u is closely approximated by $G(\cdot; \sigma, \xi)$ whenever u is sufficiently close to the right-hand endpoint.

If we ignore the approximation and assume F_u is exactly GPD, then estimation by numerical maximum likelihood (MLE) is straightforward in nearly all practical cases. The asymptotic properties are regular whenever $\xi > -\frac{1}{2}$, and alternative remedies are available for $\xi \leq -\frac{1}{2}$. These results follow from general properties of nonregular estimation established by Smith (1985). One serious competitor to maximum likelihood, apart from Bayesian methods, is the probability weighted moments (PWM) method of Hosking & Wallis (1987). The PWM method was earlier developed for the three-parameter generalized extreme value distribution (Hosking, Wallis & Wood, 1985), where it was shown by simulation to be in general more efficient than MLE for the central range of ξ ($-0.2 \leq \xi \leq 0.2$) and moderate values of n (up to 100). However, PWMs are much less flexible than MLEs as a general estimation method — the reader who doubts this is invited to re-cast all the results of this paper in terms of PWMs! Pickands himself proposed a direct method of estimation based on quantiles, and de Haan and co-workers developed methods based on moments (Dekkers & de Haan 1989, Dekkers et al. 1989, de Haan & Rootzén 1993), but these also lack the flexibility of MLE to handle more complicated situations.

In practice, of course, we want to apply this method in situations much more complicated than i.i.d. observations, for example, dependent data and situations depending on covariates. Major themes are the following.

(i) *Modeling of covariates*: The idea here is that threshold exceedances may depend on other measured variables that can be incorporated into the analysis. For example, tropospheric (ground-level) ozone, a common cause on environmental health concern in large cities, is hugely affected by meteorological factors: the worst ozone days occur in hot weather when there is low wind. It is natural to construct a regression model in which the parameters λ , σ and (possibly) ξ are functions of the covariates, through additional parameters which are estimated by MLE. A direct method of doing this was described by Davison & Smith (1990). An alternative approach, based on nonhomogeneous Poisson processes, was developed by Smith (1989) and is outlined below.

(ii) *Dependence*: Most environmental time series are not independent but exhibit serial correlation. An informal and somewhat historical approach to this is to group the data into clusters of dependent exceedances and to apply the GPD to the maximum within each cluster. A more rigorous justification for such a procedure is based on the theory of extremes in dependent stationary processes (Leadbetter *et al.* 1983, Hsing 1987, Hsing *et al.* 1988, Leadbetter 1991) which shows how clustering naturally arises in the limiting distribution of the point process of threshold exceedance times under a simultaneous rescaling of time and space.

Nevertheless, although the idea of clustering is intuitive, there are a number of approaches to identifying clusters in practice. This problem is closely tied up with *estimating the extremal index*, for which a growing literature now exists (Leadbetter *et al.* 1989, Nandagopalan 1990, Hsing 1991, Smith & Weissman 1994). The extremal index, an important parameter in its own right (Leadbetter *et al.* 1983, Leadbetter 1983) may be defined in the present context as the reciprocal of the mean cluster size in the limiting point process. It is estimated by first identifying clusters of neighboring exceedances, and then calculating the mean cluster size. (This assumes that we do not have in mind any specific model for dependence, since if we did, we might prefer a model-based approach. See section 4 for specific proposals in this case.) From this point of view, the problems of identifying clusters and estimating the extremal index are the same.

A number of procedures for identifying clusters exist. The simplest is the *blocks approach* (Leadbetter *et al.* 1989, Hsing 1991) in which the data set is divided up into blocks of consecutive observations, all the exceedances within a single block being defined to form a single cluster. An alternative is the *runs approach* in which two consecutive exceedances are defined to be in the same cluster if and only if they are less than r_n observations apart. Here r_n is a parameter to be determined. Smith & Weissman (1994) give a theoretical argument for preferring the runs approach to the blocks approach and suggest a procedure for determining r_n . In the majority of practical cases, the results are not sensitive to the precise choice of r_n , which may be based on intuitive reasoning of how

long one would expect the clusters to last, and consequently such formal procedures are not necessary.

(iii) *Seasonality*: A common feature of environmental data is that the series is not stationary, but exhibits a strong annual variability and possibly other periodic effects. (For instance, the tidal record is often decomposed into over 100 harmonic components.) Davison & Smith (1990) discussed two broad approaches to dealing with seasonality, the “prewhitening” approach in which the raw data are deseasonalized before applying the threshold, and the “seperate seasons” (or blocking) approach in which the year is divided up into blocks with a homogeneous model being assumed for each block. The main example in Smith (1989) also contains an extensive discussion of the blocking feature. The disadvantages of prewhitening are, (a) the method is unsound unless there are strong grounds, which would almost certainly have to be physical rather than statistical, for believing that the prewhitened data really are strictly stationary even in the tails, (b) it can be hard to interpret a threshold analysis of the prewhitened data as an extreme value analysis of the original data. The main disadvantage of the blocks method is that the data are unlikely to be exactly homogeneous across blocks, unless one uses very short blocks, in which case the problem becomes overparametrized. An alternative approach is to treat time of year as another covariate and use a regression approach. This is what I actually do in Section 5 below.

(iv) *Selecting the threshold*: Another practical problem in these methods is how to decide what threshold u to use. Although theoretical results on this question are available (Smith 1987, Dekkers & de Haan 1989, Dekkers *et al.* 1989), they are rather hard to use in practice because the theoretical results themselves depend on additional parameters which are unknown. A practical technique is to use the *mean residual life* plot (Davison & Smith 1990, Smith 1990). The motivation for this is the easily checked formula that for $Y \sim G(\cdot; \sigma, \xi)$ and $v > 0$ (assuming $v < -\sigma/\xi$ in the case $\xi < 0$),

$$E\{Y - v | Y > v\} = \frac{\sigma + \xi v}{1 - \xi}. \quad (1.3)$$

Therefore, an empirical plot of $E\{Y - v | Y > v\}$ against v , should be approximately a straight line. In practice, with data $X_1, \dots, X_n$, we plot

$$\frac{\sum (X_i - u) I(X_i > u)}{\sum I(X_i > u)} \quad (1.4)$$

against u , and look for the smallest u over the region in which this is a straight line. The left hand side of (1.3), or its empirical counterpart in (1.4), is known as the mean residual life in the theory of survival data, hence the name of the plot. This at least gives some empirical support to the practical choice of u , though it requires some care because the plot becomes very irregular as u approaches the upper boundary of the data. Again, from a practical point of view, the problem is often not so serious because the choice of u is either suggested on physical grounds (for example, with tropospheric ozone it is known

that different chemical processes affect the ozone levels above about 80 ppb — the current ozone standard is 120 ppb) or else the analysis is not too sensitive to the precise choice of u . Nevertheless, it must be admitted that the choice of threshold remains a major practical issue in this kind of analysis, and may still be the area where further theoretical development is most badly needed.

The Poisson process approach

An alternative approach, introduced by Smith (1989), is based on viewing the two-dimensional process of exceedance times and excesses as a point process in $\mathcal{R}^2$. Limit theorems for such processes were developed by Pickands (1971) and Resnick (1975) in the independent case, and formed a major impetus for the probabilistic work on extremes in dependent processes, reviewed in Leadbetter *et al.* (1983). However, the *statistical* application of these processes is more recent. Although at first sight more complicated than the GPD approach, it ultimately helps to simplify and unify the whole theory, especially when covariates are present.

To develop this approach, let $X_1, \dots, X_n$ be i.i.d. ($\sim F$). Suppose F is in the domain of attraction of an extreme value distribution. One way to express that property is to say that there exist constants $a_n > 0$ and b_n such that

$$n\{1 - F(a_n y + b_n)\} \rightarrow V(y) \quad (1.5)$$

for each y , and hence

$$F^n(a_n y + b_n) \rightarrow e^{-V(y)}. \quad (1.6)$$

Here we may, without loss of generality, take $V(y) = (1 + \xi y)_+^{-1/\xi}$ for fixed ξ , in which case the limiting distribution function in (1.6) is of generalized extreme value (GEV) form.

Now define a point process P_n on $[0, 1] \times \mathcal{R}$ by putting a point at each $\{i/n, (X_i - b_n)/a_n\}$, $1 \leq i \leq n$. The total number of points in this process is of course n , but the expected number in $[0, 1] \times (y, \infty)$ for any fixed y is $n\{1 - F(a_n y + b_n)\}$ which, by (1.5), converges to $V(y)$ (finite). From this it is possible to establish that P_n converges weakly, on all sets for which the second coordinate is bounded away from its lower boundary, to a nonhomogenous Poisson process P whose intensity measure Λ ($\Lambda(A)$ is the expected number of points in A , for $A \subset [0, 1] \times \mathcal{R}$) satisfies

$$\Lambda\{(t_1, t_2) \times (y, \infty)\} = (t_2 - t_1)V(y)$$

whenever $0 \leq t_1 < t_2 \leq 1$ and $y \in \mathcal{R}$ ($y > -1/\xi$ if $\xi > 0$).

In practice, rather than try to construct the limiting process with the renormalization required, a more practical approach is to work directly with a nonhomogeneous Poisson process applied to all observations over a threshold u . Thus, for any observation X_t , taken

at time t , for which $X_t > u$, one puts a point at (t, X_t) . This is treated as part of a nonhomogeneous Poisson process on $\mathcal{R} \times (u, \infty)$ whose intensity measure satisfies

$$\Lambda\{(t_1, t_2), (x, \infty)\} = (t_2 - t_1) \left(1 + \xi \frac{x - \mu}{\psi}\right)_+^{-1/\xi}, \quad t_1 \leq t_2, \quad x \geq u. \quad (1.7)$$

It can be seen, from a number of different points of view, that this is equivalent to either of the usual approaches to extreme value theory. If M_T denotes the maximum of the process over time interval $(0, T)$, assumed bigger than u , then the event $M_T \leq x$ is equivalent to saying that the set $(0, T) \times (x, \infty)$ is empty, and this has probability

$$\exp \left\{ -T \left(1 + \xi \frac{x - \mu}{\psi}\right)_+^{-1/\xi} \right\}, \quad x \geq u,$$

which is of the usual generalized extreme value form. Another fact that follows directly from (1.7) is that, given a point (T_i, X_i) for which $X_i > u$, the conditional probability that $X_i - u > y$ is

$$\frac{\left(1 + \xi \frac{u+y-\mu}{\psi}\right)_+^{-1/\xi}}{\left(1 + \xi \frac{u-\mu}{\psi}\right)_+^{-1/\xi}} = \left\{ 1 + \frac{\xi y}{\psi + \xi(u - \mu)} \right\}_+^{-1/\xi}$$

which is of GPD form (1.2) with $\sigma = \psi + \xi(u - \mu)$.

Estimation from (1.7) proceeds by constructing the appropriate likelihood for an non-homogeneous Poisson process (Cox & Lewis 1966, Section 3.3): if we re-define $V(x)$ to be

$$V(x; \mu, \psi, \xi) = \left(1 + \xi \frac{x - \mu}{\psi}\right)_+^{-1/\xi},$$

and let $v(x; \mu, \psi, \xi) = -\partial V(x; \mu, \psi, \xi)/\partial x$, then if we observe a random number N of exceedances (T_j, X_j) with $X_j > u$ over a time period $(0, T)$, the approximate likelihood function is

$$L(\mu, \psi, \xi) = \exp\{-TV(u; \mu, \psi, \xi)\} \cdot \prod_{j=1}^N v(X_j; \mu, \psi, \xi). \quad (1.8)$$

Estimation of μ, ψ and ξ proceeds by numerical maximization of (1.8).

Comparison of the GPD and Poisson process approaches

The two approaches are equivalent in the sense that any model which is expressed in terms of one of the two approaches may be rewritten as an equivalent model in terms of the other. The main advantage of the Poisson approach is that the parametrization

is more convenient for complex models involving covariates. Suppose there is a vector of covariates z_t for each time point t , and that the quantities $\mu = \mu_t$, $\psi = \psi_t$, $\xi = \xi_t$ depend on a parameter vector η through functions $\mu_t(\eta) = \mu(z_t; \eta)$, $\psi_t(\eta) = \psi(z_t; \eta)$, $\xi_t(\eta) = \xi(z_t; \eta)$. Then the likelihood, as a function of η , may be expressed in the form

$$L(\eta) = \exp \left\{ - \int_0^T V(u; \mu_t(\eta), \psi_t(\eta), \xi_t(\eta)) dt \right\} \cdot \prod_{j=1}^N v(X_j; \mu_{T_j}(\eta), \psi_{T_j}(\eta), \xi_{T_j}(\eta)), \quad (1.9)$$

extending (1.8). For example, in the not uncommon situation that the location parameter of a distribution is taken to depend on the covariates, with the scale and shape parameters held fixed, we can simplify (1.9) by setting $\psi_t(\eta)$ and $\xi_t(\eta)$ equal to constants ψ and ξ . Such a simplification was used by Smith (1989) to model the case of an additive linear trend. It is not so easy to see how to incorporate such a feature into the GPD model. In principle, then, I believe this method to be superior to the GPD approach, though for practical application it is beneficial to have both methods in hand.

2. Multivariate extremes

Multivariate extreme value theory is concerned with the joint distribution of extremes of two or more dependent random variables. The traditional approach to their definition has been via componentwise maxima (or minima): if $\{(X_{i1}, \dots, X_{ip}), i = 1, \dots, n\}$ is a sample of independent p -vectors, then the vector of componentwise maxima is $(M_{n1}, \dots, M_{np})$, where $M_{nj} = \max(X_{1j}, \dots, X_{nj})$, $j = 1, \dots, p$. The limiting joint distributions of these maxima $(M_{n1}, \dots, M_{np})$, subject to location-scale renormalization of each component, are the *multivariate extreme value distributions*. The bivariate case $p = 2$ is of particular significance, both because it was the first case to be studied in detail and because much of the theory is simpler for this than the general multivariate case.

The main contrast with the univariate case is that there is no finite-parameter family which exhausts the class of multivariate extreme value distributions, for any $p \geq 2$. Hence there are two basic approaches, namely those based on parametric subfamilies and those which are essentially nonparametric. The discussion in this paper is concerned entirely with parametric families. How severe a restriction this is, from a practical point of view, is still not clear, but there are now several different approaches to the construction of multivariate extreme value families, so this is still a very broad approach.

There are several different but equivalent characterizations of multivariate extreme value distributions, reviewed in detail in the books of Resnick (1987) and Galambos (1987). One feature of multivariate extreme value distributions is that the dependence structure is preserved under transformations of the marginal distributions, so there is no loss of generality in restricting attention to a particular univariate extreme value family. For example, de Haan & Resnick (1977) assumed unit Fréchet ($F(y) = e^{-1/y}$) margins and developed the characterization

$$G(y_1, \dots, y_p) = \exp \left\{ - \int_{T_p} \max_j \left(\frac{u_j}{y_j} \right) \mu(du_1, \dots, du_p) \right\} \quad (2.1)$$

where μ is a positive measure on the set

$$T_p = \{(u_1, \dots, u_p) : u_j \geq 0, \sum_j u_j^2 = 1\}$$

subject to the condition

$$\int_{T_p} u_j \mu(du_1, \dots, du_p) = 1 \text{ for all } j.$$

An alternative approach, due to Pickands, leads to a representation formula (with GEV margins) of the form

$$G(y_1, \dots, y_p) = \exp \left\{ - \left(\sum_{j=1}^p t_j \right) A \left(\frac{t_1}{\sum_{j=1}^p t_j}, \dots, \frac{t_p}{\sum_{j=1}^p t_j} \right) \right\}, \quad (2.2)$$

where $A(x_1, \dots, x_p)$ is a convex function over the unit simplex $S_p = \{(x_1, \dots, x_p) : x_1 \geq 0, \dots, x_p \geq 0, \sum x_j = 1\}$ which takes the value 1 at the corner points (when one x_j is 1 and the rest 0), and $t_j = \{1 + \xi_j(y_j - \mu_j)/\psi_j\}_+^{-1/\xi_j}$. In the bivariate case this reduces to a simpler form

$$G(y_1, y_2) = \exp \left\{ - (t_1 + t_2) A \left(\frac{t_2}{t_1 + t_2} \right) \right\} \quad (2.3)$$

with A a convex function on $[0, 1]$ satisfying $A(0) = 1$, $A(1) = 1$, $A(x) \geq \max(x, 1 - x)$ for $0 < x < 1$ (the case where equality is attained in the last expression being the case where both components are equal with probability 1, the extreme case of complete dependence).

In principle, one can generate multivariate distributions by taking general formulae such as (2.1)-(2.3) and inserting specific functions for μ or A . Simple-minded attempts to do this fail to yield useful or tractable families, but there are by now a number of approaches which do yield usable parametric families. First I consider the bivariate case and then generalize to $p > 2$.

Extensive work on bivariate extremes was carried out by Tiago de Oliveira in a long series of papers (for a review see Tiago de Oliveira, 1984), but recent years have seen a number of new models introduced. Among the models currently considered are

(a) *Mixed Model* (Gumbel & Mustafi 1967):

$$A(w) = \theta w^2 - \theta w + 1, \quad \theta \in [0, 1],$$

(b) *Logistic Model*: (Gumbel 1960):

$$A(w) = \{(1 - w)^{1/\alpha} + w^{1/\alpha}\}^\alpha, \quad 0 \leq \alpha \leq 1,$$

the limits $\alpha = 1$ and $\alpha \downarrow 0$ corresponding to independence and complete dependence respectively,

(c) *Asymmetric Mixed Model* (Tawn 1988):

$$A(w) = \phi w^3 + \theta w^2 - (\theta + \phi)w + 1, \quad (\theta \geq 0, \theta + \phi \leq 1, \theta + 2\phi \leq 1, \theta + 3\phi \geq 0),$$

(d) *Asymmetric Logistic Model* (Tawn 1988):

$$A(w) = \{\theta^{1/\alpha}(1-w)^{1/\alpha} + \phi^{1/\alpha}w^{1/\alpha}\}^\alpha + (\theta - \phi)w + 1 - \theta, \quad 0 \leq \theta \leq 1, 0 \leq \phi \leq 1, 0 \leq \alpha \leq 1.$$

The asymmetric versions of the mixed and logistic models allow for non-exchangability between the two components.

(e) *Inverted Logistic Model* (Joe 1990):

$$A(w) = 1 - [(\phi_1 w)^{-\tau} + \{\phi_2(1-w)\}^{-\tau}]^{-1/\tau}, \quad \tau \geq 0, 0 \leq \phi_1 \leq 1, 0 \leq \phi_2 \leq 1.$$

The simplest (symmetric) form is when $\phi_1 = \phi_2 = 1$; in this form the mixed model with $\theta = 1$ arises when $\tau = 1$ while the limits $\tau \rightarrow 0$, $\tau \rightarrow \infty$ correspond respectively to independence and complete dependence.

(f) *Gaussian Model*:

$$A(w) = (1-w)\Phi\left\{\frac{a}{2} + \frac{1}{a}\log\left(\frac{1-w}{w}\right)\right\} + w\Phi\left\{\frac{a}{2} + \frac{1}{a}\log\left(\frac{w}{1-w}\right)\right\}$$

for $a \in [0, \infty]$, where Φ is the normal distribution function. This was introduced independently by Hüsler and Reiss (1989), who derived it as a penultimate approximation for dependent normal extremes, and Smith (1991) as an example of a max-stable process (see Stuart Coles' paper at this meeting for a discussion of max-stable processes).

(g) *Bilogistic Model* (Joe, Smith and Weissman 1992):

$$A(w) = \int_0^1 \max\{(1-\alpha)(1-w)u^{-\alpha}, (1-\beta)w(1-u)^{-\beta}\} du, \quad (0 \leq \alpha \leq 1, 0 \leq \beta \leq 1),$$

which is an alternative to (d) as an asymmetric form of the logistic model, to which it reduces when $\alpha = \beta$.

For multivariate ($p > 2$) cases, the possibilities are even more diverse. In this case an equivalent characterization of Pickands' A function (from (2.2)) is

$$A(w_1, \dots, w_p) = \int_{S_p} \max_j (w_j u_j) dH(u_1, \dots, u_p), \quad (2.4)$$

H being a positive measure on S_p satisfying

$$\int_{S_p} u_j dH(u_1, \dots, u_p) = 1, \quad j = 1, \dots, p. \quad (2.5)$$

Suppose we restrict H to be an absolutely continuous measure on the interior of S_p . This is a nontrivial restriction, since many examples require a degenerate H measure — for instance, in the independent case H gives mass 1 to each of the corner points of the simplex. We can construct families of multivariate extreme value distributions by finding measures H satisfying (2.5) and then applying (2.4). However, it is still not so easy to find suitable H , especially with the restriction (2.5).

One construction suggested by Coles & Tawn (1991) was to take an arbitrary density h^* on S_p with positive first moments, and to define

$$m_j = \int_{S_p} u_j h^*(u_1, \dots, u_p) du_1, \dots, du_p, \quad j = 1, \dots, p.$$

followed by

$$h(w_1, \dots, w_p) = \prod_{j=1}^p m_j \left(\sum_j m_j w_j \right)^{-p-1} h^* \left(\frac{m_1 w_1}{\sum_j m_j w_j}, \dots, \frac{m_p w_p}{\sum_j m_j w_j} \right).$$

Then h is the density of a valid measure H satisfying (2.5), and (2.4) may be used to construct the A function and hence the distribution function.

An example for which this construction works is the Dirichlet density,

$$h^*(w_1, \dots, w_p) = \left\{ \frac{\Gamma(\sum_j \alpha_j)}{\prod_j \Gamma(\alpha_j)} \right\} \prod_{j=1}^p w_j^{\alpha_j-1}, \quad \alpha_j > 0, \quad j = 1, \dots, p, \quad (w_1, \dots, w_p) \in S_p.$$

Then $m_j = \alpha_j / (\sum_k \alpha_k)$ and

$$h(w_1, \dots, w_p) = \prod_{j=1}^p \left\{ \frac{\alpha_j}{\Gamma(\alpha_j)} \right\} \frac{\Gamma(\sum_j \alpha_j + 1)}{(\sum_j \alpha_j w_j)^{p+1}} \prod_{j=1}^p \left(\frac{\alpha_j w_j}{\sum_k \alpha_k w_k} \right)^{\alpha_j-1}.$$

This is therefore known as the Dirichlet model for multivariate extremes.

Other approaches to the construction of multivariate extreme value distributions include a hierarchical approach introduced by Tawn (1990) and in a related way by Joe (1990), time series constructions (cf. Section 4 below) and methods based on specializations of max-stable processes.

Coles & Tawn (1991) reviewed these approaches in more detail. Two models that have been known for a very long time are the p -dimensional generalization of Gumbel's logistic model (Gumbel 1960), $A(w_1, \dots, w_p) = \left(\sum w_j^{1/\alpha} \right)^\alpha$, and the family introduced by Marshall & Olkin (1967). However, Gumbel's model has the severe restriction that it forces all p components to be exchangeable, while the Marshall-Olkin model can never be absolutely continuous, a fact which restricts its range of applicability in the areas in which we are interested. With the above classes of models, there are now many more possibilities.

Estimation

The estimation of multivariate extreme value distributions has long been known to be problematic because all of them become nonregular as the independence case is approached. That is, standard maximum likelihood asymptotic theory breaks down, as the Fisher information about the dependence parameter tends to ∞ . This problem was, however, essentially solved by Tawn (1988, 1990), who developed stable asymptotic distributions for the score statistic at independence, and used this property to characterize the limiting properties of maximum likelihood in this case. For interior points of the parameter space (i.e. bounded away from independence or any of the other boundary conditions) these difficulties do not arise and maximum likelihood seems to behave perfectly well. In any case, our own preference is now to use standard maximum likelihood theory for all calculations except that of testing independence, for which we know that the standard χ^2 distribution for the likelihood ratio statistic breaks down. Even then, maximum likelihood may not be easy to implement in practice, as it requires being able to evaluate the density, and in several of the cases described above (for instance, the Dirichlet model) this is not available in closed form. Computer algebra is one of the techniques that has been used to get round this problem.

3. Threshold methods for multivariate extremes

The first attempts to construct threshold-based methods of statistical inference in the multivariate cases were those of Coles & Tawn (1991) and Joe *et al.* (1992). Both of these are based on a generalization of the point process approach described in Section 1. If we have multivariate data $\{X_{i1}, \dots, X_{ip}\}$ and renormalize each component so that $(M_{nj} - b_{nj})/a_{nj}$ tends to a limiting univariate extreme value distribution, then we can form a point process P_n on $[0, 1] \times \mathcal{R}^p$ by putting a point at each $\{i/n, (X_{i1} - b_{i1})/a_{i1}, \dots, (X_{ip} - b_{ip})/a_{ip}\}$, $1 \leq i \leq n$. In this case the number of points in any rectangle of the form $(t_1, t_2) \times (y_1, \infty) \times \dots \times (y_p, \infty)$ converges to a limiting form, which we may write in the form $(t_2 - t_1)V(y_1, \dots, y_p)$, as $n \rightarrow \infty$. In fact the limiting extreme value distribution, in the sense defined in Section 2, is just $\exp\{-V(y_1, \dots, y_p)\}$ and this result combined with representations such as (2.2) may be used to obtain point-process methods for inference about multivariate extremes. For applications of this technique, I refer to Coles & Tawn (1994) and Jonathan Tawn's paper at this meeting.

For the remainder of this section I discuss an alternative approach, currently being developed by Jonathan Tawn, Stuart Coles and myself (Smith *et al.* 1993), in which we work directly with the exceedances over a threshold, analogous to the GPD approach in the univariate case.

Suppose $(X_1, \dots, X_p)$ is a typical p -variate data point but that, in the spirit of threshold methods, for each j between 1 and p , we observe not X_j but $Z_j = \max(X_j, u_j)$, together with an indicator δ_j ($\delta_j = 1$ if $X_j > u_j$, 0 otherwise). Here u_j is a fixed threshold, which may be different for each component j . We want to derive an approximation for the joint distribution of $\{(\delta_j, Z_j), j = 1, \dots, p\}$. The underlying assumption is that $(X_1, \dots, X_p)$ lies in the domain of attraction of a multivariate extreme value distribution. One characterization of that, described in detail by Resnick (1987), is first to transform each margin to unit Fréchet form by means of a probability integral transformation, and then to define the distribution function $F_*(v_1, \dots, v_p)$ for the transformed distribution. Resnick's Proposition 5.15 shows that the distribution is in the domain of attraction of a multivariate extreme value family if and only if, for all $(v_1, \dots, v_p)$ in which each $v_j > 0$,

$$\lim_{t \rightarrow \infty} \frac{1 - F_*(tv_1, \dots, tv_p)}{1 - F_*(t, \dots, t)} = \frac{-\log G_*(v_1, \dots, v_p)}{-\log G_*(1, \dots, 1)} = \frac{V(v_1, \dots, v_p)}{V(1, \dots, 1)}. \quad (3.1)$$

where G_* is a multivariate extreme value distribution function with unit Fréchet margins.

By analogy with threshold methods for univariate extremes, in which the approximate generalized Pareto distribution is in effect treated as exact for sufficiently high thresholds, we may consider a scheme in which (3.1) is treated as exact, for sufficiently large t and $v_1, \dots, v_p$.

The assumption we make about each marginal component is that for sufficiently large threshold u_j , the marginal distribution of $X_j - u_j$ given $X_j > u_j$ is GPD. This allows us to write the j 'th marginal distribution function in the form

$$F_j(x) = 1 - \lambda_j \{1 + \xi_j(x - u_j)/\sigma_j\}_+^{-1/\xi_j}, \quad x \geq u_j,$$

where $\lambda_j = 1 - F_j(u_j)$. By analogy, we also take the limiting result (3.1) as an identity for sufficiently large t . In fact, it is more convenient to treat (3.1) as an identity for some fixed $t = t_c$, provided the v_j are sufficiently large; this is clearly an equivalent interpretation. Specifically, taking $t_c = 1$ and $v_j \geq \lambda_j^{-1}$, $j = 1, \dots, p$ ensures that the GPD approximation is applicable for each marginal component. Our additional assumption is that these levels are also sufficiently high for the asymptotic dependence structure to be a valid approximation through (3.1).

By assuming this form of F_j , transforming to unit Fréchet (in the tails) and applying (3.1), the form of multivariate joint distribution that arises is

$$F(x_1, \dots, x_p) = 1 - V \left\{ \lambda_1^{-1} \left(1 + \xi_1 \frac{x_1 - u_1}{\sigma_1} \right)_+^{1/\xi_1}, \dots, \lambda_p^{-1} \left(1 + \xi_p \frac{x_p - u_p}{\sigma_p} \right)_+^{1/\xi_p} \right\}, \quad (3.2)$$

valid for $x_1 \geq u_1, \dots, x_p \geq u_p$. For example, in the $p = 2$ case with the logistic model of Section 2, we have $V(y_1, y_2) = (y_1^{-1/\alpha} + y_2^{-1/\alpha})^\alpha$, so equation (3.2) becomes, for $x_1 \geq u_1, x_2 \geq u_2$,

$$F(x_1, x_2) = 1 - \left\{ \lambda_1^{1/\alpha} \left(1 + \xi_1 \frac{x_1 - u_1}{\sigma_1} \right)_+^{-1/(\alpha\xi_1)} + \lambda_2^{1/\alpha} \left(1 + \xi_2 \frac{x_2 - u_2}{\sigma_2} \right)_+^{-1/(\alpha\xi_2)} \right\}^\alpha. \quad (3.3)$$

This formula still has one disadvantage: when $\alpha = 1$ it does not reduce exactly to the independent GPD case. To get around this problem, Jonathan Tawn and Anthony Ledford have suggested an alternative construction, starting from replacing $1 - F_\star$ with $-\log F_\star$ in (3.1), which leads to the formula

$$F(x_1, \dots, x_p) = \exp \left[-V \left\{ \left(-\log \left(1 - \lambda_1 \left(1 + \xi_1 \frac{x_1 - u_1}{\sigma_1} \right)_+^{-1/\xi_1} \right) \right)^{-1}, \dots, \right. \right. \\ \left. \left. \left(-\log \left(1 - \lambda_p \left(1 + \xi_p \frac{x_p - u_p}{\sigma_p} \right)_+^{-1/\xi_p} \right) \right)^{-1} \right\} \right] \quad (3.4)$$

which shares the same asymptotic properties as (3.2) while avoiding this one undesirable feature. In this case, (3.3) becomes

$$F(x_1, x_2) = \exp \left[- \left\{ \left(-\log \left(1 - \lambda_1 \left(1 + \xi_1 \frac{x_1 - u_1}{\sigma_1} \right)_+^{-1/\xi_1} \right) \right)^{1/\alpha} \right. \right. \\ \left. \left. + \left(-\log \left(1 - \lambda_2 \left(1 + \xi_2 \frac{x_2 - u_2}{\sigma_2} \right)_+^{-1/\xi_2} \right) \right)^{1/\alpha} \right\}^\alpha \right]. \quad (3.5)$$

The main complication in the estimation of these distributions is the form of the likelihood. The contribution to the likelihood from a vector $\{(\delta_j, Z_j)\}$ is constructed by differentiating the joint distribution function with respect to the observed components. As an example, suppose $p = 3$ and $\delta_1 = \delta_3 = 1, \delta_2 = 0$ (observations 1 and 3 above the threshold, 2 below). The contribution to the likelihood function from this triple is

$$\frac{\partial^2 F(x_1, x_2, x_3)}{\partial x_1 \partial x_3} \Big|_{x_1=X_1, x_2=u_2, x_3=X_3}.$$

Maximum likelihood estimation in these cases is therefore straightforward in principle. In terms of its asymptotic properties, maximum likelihood is regular at interior points of the parameter space, but there are still difficulties in testing independence which require the construction of special tests (work by Ledford & Tawn, in progress).

4. Markov chains

One consequence of the methodology of Section 3 is that it gives us a new way to think about extreme values in univariate time series. The work described here is based on Smith *et al.* (1993).

The threshold approach in Section 2 is based on identifying clusters of neighboring exceedances, and then using the GPD to model the largest value over the threshold within each cluster. This is open to two major objections, first that the method of identifying clusters is arbitrary, and second that it seems inefficient to model only the largest value within each cluster when all of them contain useful information. An alternative approach would be to assume some model for the dependence in the time series. However, the two major classes of models used in time series analysis are Gaussian models and linear models. Gaussian models are not very satisfactory in extreme value analysis because, subject to a very weak long-range independence condition, they always have extremal index 1, and therefore necessarily fail to capture the features we are trying to model. Linear models are in principle a more general class, but the required extremal computations are not easy, and there remains the question of how to verify the correctness of the linearity assumption, which could be crucial to any calculation of extreme value properties.

The approach suggested here is to model the short-term dependence by assuming the process to be a k -th order Markov chain, i.e. the distribution of X_n given X_m , $m < n$ is determined completely by a transition density depending only on $X_{n-k}, \dots, X_{n-1}$. We also assume the process to be stationary. Such a process is, of course, completely specified by its consecutive $(k+1)$ 'st order joint distributions, and if we assume these are in the domain of attraction of a $(k+1)$ -dimensional Markov chain, an attractive theory results.

Suppose the consecutive j -dimensional joint densities are denoted by $f_j(x_1, \dots, x_j)$. Then the joint density of a realization $(X_1, \dots, X_n)$ for $n > k$ is given by

$$\frac{\prod_{j=1}^{n-k} f_{k+1}(X_j, \dots, X_{j+k})}{\prod_{j=2}^{n-k} f_k(X_j, \dots, X_{j+k-1})}. \quad (4.1)$$

Thus, given parametric families for f_k and f_{k+1} , the likelihood function is a ratio of two terms, each of which is constructed in the same way as a likelihood for independent multivariate data.

Now suppose the parametric family is not completely specified, but only determined in the tails, as in Section 3. Then we may approximate each of the numerator and denominator in (4.1) by an appropriate limiting form. Our specific proposal is to use the approximations developed in Section 3 to do this.

One consequence of adopting this approach is that it is still necessary to calculate extremal properties of the process for the fitted Markov chain. There has been considerable recent progress in this area. For example, Smith (1992) developed a method of

calculating the extremal index for a Markov chain under the main assumption that the consecutive bivariate distributions lie in the domain of attraction of a bivariate extreme value distribution. Yun (1993) has extended this calculation to k 'th order Markov chains. Smith *et al.* (1993) discuss a variety of other functions such as the sum of exceedances within a cluster. Anderson & Dancy (1992) refer to this as the severity of an extreme event, and develop a characterization of its asymptotic form in terms of the limiting point process of crossings of the threshold. With the exception of Smith (1992), which reduced to the numerical solution of a Wiener-Hopf integral equation, all of these solutions rely on simulation to calculate limiting distributions of functionals within a cluster, but such methods seem entirely within the scope of available computing resources.

Smith *et al.* (1993) applied these ideas (with $k = 1$) to model seasonal minima in a long-term temperature series, and a follow-up paper by Coles *et al.* (1993) will develop this analysis further.

5. An application: High-level exceedances of tropospheric ozone

The following discussion is based on work currently being carried out at the National Institute of Statistical Sciences, as part of a cooperative agreement with the United States Environmental Protection Agency (EPA).

Ozone is produced by complex chemical processes in the lower atmosphere as a result of the emissions into the atmosphere of hydrocarbons and nitrogen oxides. Ozone is not itself emitted from a source, but it does have a direct effect on human health and therefore is monitored by the EPA as a measure of general pollution in the atmosphere. Specifically, the national ambient air quality standard (NAAQS) specifies that, at any individual monitoring site, the number of exceedances by daily ozone maxima of the level of 120 parts per billion (ppb) should not exceed 3 in any 3-year period. Regions which violate this are deemed non-attainment regions. These non-attainment regions are generally in major cities such as New York, Chicago, Houston and Los Angeles.

Although ozone is produced as the result of precursor gases emitted from sources, there are many other factors determining the actual ozone level on any given day. Principal among these other factors is meteorology. Ozone levels are particularly high on hot, windless days in summer.

The purpose of the present study is to examine trends in measured ozone data, particularly taking meteorology into account. The ultimate objectives of this are to assist EPA in determining a more suitable ozone standard, which may include such questions as whether daily maximum is the most appropriate indicator of the ozone level and whether some attempt should be made to compensate for meteorological conditions in defining the standard. A more specific question, which we shall focus on here, is whether there is a trend in the data after taking meteorological factors into account. The hope is that we will

find a downward trend that would indicate the success of EPA's efforts to reduce ozone over the last 15 years.

Although the problem is not concerned solely with extreme values, the fact that the ozone standard is defined in terms of crossings of a high threshold suggests that an appropriate analysis is to study trends in the rate of crossings of a high threshold, and from that point of view, threshold methods of the sort that have been described in this paper seem relevant. The analysis that follows is based on preliminary results for a single ozone station to the north-west of Chicago.

The data analyzed were daily maximum ozone levels for 1981-1991. A total of 3323 daily values were available, which represents a high proportion (83%) of all possible days in that period. Days when no value is available are treated as missing data. In addition, hourly meteorological data are available for the Chicago area, and from these the noon values of temperature, windspeed and a number of other variables are extracted as covariates. There appears to be a seasonal factor over and above what is explainable in terms of meteorology, and this is modeled by introducing covariates $\cos(2\pi d/366)$, $\sin(2\pi d/366)$, where d is calendar day within year ($1 \leq d \leq 366$). Finally, the year itself is treated as a covariate in the analysis.

In earlier work as part of the same study, Bloomfield *et al.* (1993) used principal components analysis to determine suitable weighted averages based on a network of stations in the Chicago area, and applied nonlinear regression techniques to determine a suitable form of dependence. Their findings showed a significant negative trend when meteorological effects are included as covariates, but their analysis was not specifically confined to high-level exceedances. The following analysis focusses on that aspect.

A natural starting point is to consider just the crossing probability of a single threshold and to look for trends in that. For this purpose, the analysis was based not on the threshold 120 ppb, for which there are too few exceedances to allow a sensible analysis, but on 80 ppb. This threshold is low enough to yield considerable data (115 exceedances over 11 years for the specific site being considered here) but high enough still to be relevant for consideration of extreme ozone levels.

With this in mind, I analyzed crossings of 80 ppb as a binary data set, recording a 1 if the threshold is exceeded and a 0 otherwise. A natural model for binary data is the Bernoulli model with a logistic link function:

$$\log \left\{ \frac{\lambda(t)}{1 - \lambda(t)} \right\} = \sum_j \beta_j z_j(t) \quad (5.1)$$

where $\lambda(t)$ is the probability of exceeding the threshold on day t , and $z_j(t)$, $j = 1, 2, \dots$, are known covariates and $\beta_1, \beta_2, \dots$ unknown parameters to be estimated.

In the present case, fitting (5.1) by stepwise selection of variables resulted in a regression model including year and the seasonal variables as covariates, as well as a number

of meteorological variables: temperature, windspeed factored into directional components, pressure, visibility, relative humidity, dewpoint and a temperature-windspeed interaction term. Significance of the yearly trend may be assessed either from the standard error of the estimated coefficient, or from comparisons of log likelihoods (or deviances) of models with and without the fitted trend.

As an example of the analysis, with the covariates just mentioned, we obtain the following negative log likelihoods:

With yearly trend: 241.5,
Without yearly trend: 250.2,

for a deviance statistic of $2 \times (250.2 - 241.5) = 17.4$, clearly highly significant against a nominal χ^2_1 distribution under the null hypothesis of no trend. Since the coefficient is negative (-0.16), this does indicate a downward trend.

This analysis does not allow for time-series dependence in the data. The simplest way to bring that in is to introduce an additional covariate which takes the value 1 if the previous day was an exceedance (of the same level 80), 0 otherwise. This turns our Bernoulli process into a first-order Markov chain, but likelihood-based methods are still valid and produced the results

With yearly trend: 238.1,
Without yearly trend: 245.7,

showing that the effect of dependence is clearly significant. Note, however, that the deviance statistic for the trend of $2 \times (245.7 - 238.1) = 15.2$, though still significant, is somewhat reduced from the first analysis, reflecting the fact that dependence in the data increases the effective standard error and so reduces the level of significance of any test statistic. Another feature of including the new covariate is that it does somewhat affect the trend estimate itself (now -0.20). Thus, we do prefer to include the new covariate in the analysis. Adding a second covariate based on 2-day lags did not produce any further significant reduction in negative log likelihood and so suggests that a first-order Markov model is adequate.

This analysis, however, is still based only on the binary variable of whether the threshold 80 is crossed. Since our real interest is in higher ozone levels than this, we would like to extend the analysis to include the distribution of excesses over the threshold. For this problem, if we go back to assuming independent days, the GPD of equation (1.2) is a natural model. Current results of the analysis suggest that a logarithmic link function for σ (i.e. we assume that $\log \sigma$ is a linear function of the covariates), with constant ξ , is an appropriate model for the data. The covariates included are temperature and pressure, which for this particular station seem to be the only significant ones, as well as year. As an example of the effect of including year in this analysis, we cite the negative log likelihoods

With yearly trend: 404.4,
Without yearly trend: 404.8,

so that the yearly trend, for the excesses taken on their own, is not significant. However, when we put the two parts of the analysis together to obtain a combined log likelihood ($241.5+404.4=645.9$ with the trend, $250.2+404.8=655.0$ without), the trend is still significant. In the light of our earlier analysis in which there was a significant dependence, however, this raises the question of how to include dependence in the combined analysis of threshold crossing times and excess values.

I have analyzed this question by assuming a first-order Markov structure with (3.5) to model the joint distribution of consecutive days. Here λ_1 and λ_2 are the threshold crossing probabilities, modeled as a logit linear function of the same structure as (5.1), and we assume $\xi_1 = \xi_2$ independent of covariates, $\log \sigma_1$ and $\log \sigma_2$ as linear functions of the covariates on the respective days. Fitting this model, with the same covariates as in the previous analyses, produced a reduction in negative log likelihoods, from 645.9 and 655.0 in the independent case, to 631.3 and 639.6 with α included. The estimate was $\hat{\alpha} = 0.79$ with an approximate standard error of 0.05. Thus for this analysis at least, it appears that dependence is a significant factor and should be taken into account in any attempt to fit models to exceedances on consecutive days.

Acknowledgement. This work was also partially supported by the NSF. I thank Jonathan Tawn for discussion leading to the equations (3.4) and (3.5).

References

- Anderson, C.W. & Dancy, G. (1992), The severity of extreme events. Preprint, Sheffield University.
- Bloomfield, P., Royle, A. & Yang, Q. (1993), Accounting for meteorological effects in measuring urban ozone levels and trends. National Institute of Statistical Sciences, draft report.
- Coles, S.G. & Tawn, J.A. (1991), Modelling extreme multivariate events. *J. R. Statist. Soc. B* **53**, 377-392.
- Coles, S.G. & Tawn, J.A. (1994), Statistical methods for multivariate extremes: An application to structural design (with discussion). *Applied Statistics*, to appear.
- Coles, S.G., Tawn, J.A. & Smith, R.L. (1993), A seasonal Markov model for extremely low temperatures. In preparation.
- Cox, D.R. & Lewis, P.A.W. (1966), *The Statistical Analysis of Series of Events*. Methuen, London.
- Davison, A.C. (1984), Modelling excesses over high thresholds, with an application. In *Statistical Extremes and Applications*, J. Tiago de Oliveira (ed.), Reidel, Dordrecht, 461-482.
- Davison, A.C. & Smith, R.L. (1990), Models for exceedances over high thresholds (with discussion). *J.R. Statist. Soc. B* **52**, 393-442.

Dekkers, A.L.M. & Haan, L. de (1989), On the estimation of the extreme-value index and large quantile estimation. *Ann. Statist.* **17**, 1795-1832.

Dekkers, A.L.M., Einmahl, J.H.J. & Haan, L. de (1989), A moment estimator for the index of an extreme-value distribution. *Ann. Statist.* **17**, 1833-1855.

Galambos, J. (1987), *The Asymptotic Theory of Extreme Order Statistics* (2nd. edn.). Krieger, Melbourne, Fl. (First edn. published 1978 by John Wiley, New York.)

Gumbel, E.J. (1960), Distributions des valeurs extremes en plusieurs dimensions. *Publ. Inst. Statist. Univ. Paris* **9**, 171-173.

Gumbel, E.J. & Mustafi, C.K. (1967), Some analytic properties of bivariate extremal distributions. *J. Amer. Statist. Assoc.* **62**, 569-588.

Haan, L. de & Resnick, S.I. (1977), Limit theory for multivariate sample extremes. *Z. Wahrscheinlichkeitstheorie v. Geb.* **40**, 317-337.

Haan, L. de & Rootzén, H. (1993), On the estimation of high quantiles. *Journal of Statistical Planning and Inference* **35**, 1-13.

Hosking, J.R.M., & Wallis, J.R. (1987), Parameter and quantile estimation for the generalized Pareto distribution. *Technometrics* **29**, 339-349.

Hosking, J.R.M., Wallis, J.R. & Wood, E.F. (1985), Estimation of the generalised extreme-value distribution by the method of probability-weighted moments. *Technometrics* **27**, 251-261.

Hsing, T. (1987), On the characterization of certain point processes. *Stoch. Proc. Appl.* **26**, 297-316.

Hsing, T., Hüsler, J. & Leadbetter, M.R. (1988), On the exceedance point process for a stationary sequence. *Probability Theory and Related Fields* **78**, 97-112.

Hsing, T. (1991), Estimating the parameters of rare events. *Stoch. Proc. Appl.* **37**, 117-139.

Hüsler, J. & Reiss, R.-D. (1989), Maxima of normal random vectors: Between independence and complete dependence. *Statistics and Probability Letters* **7**, 283-286.

Joe, H. (1990), Families of min-stable multivariate exponential and multivariate extreme value distributions. *Statistics and Probability Letters* **9**, 75-81.

Joe, H., Smith, R.L. & Weissman, I. (1992), Bivariate threshold methods for extremes. *J.R. Statist. Soc. B.*, **54**, 171-183.

Leadbetter, M.R. (1983), Extremes and local dependence in stationary sequences. *Z. Wahrsch. v. Geb.* **65**, 291-306.

Leadbetter, M.R. (1991), On a basis for 'peaks over thresholds' modeling. *Statistics and Probability Letters* **12**, 357-362.

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

Leadbetter, M.R., Weissman, I., Haan, L. de & Rootzén, H. (1989), On clustering of high levels in statistically stationary series. *Proceedings of the Fourth International Meeting on Statistical Climatology*, ed. John Sansom. New Zealand Meteorological Service, P.O. Box 722, Wellington, New Zealand.

Marshall, A.W. & Olkin, I. (1967), A multivariate extremal distribution. *J. Amer. Statist. Assoc.* **62**, 30-44.

- Nandagopalan, S. (1990), Multivariate extremes and the estimation of the extremal index. Ph.D. dissertation available as Technical Report Number 315, Center for Stochastic Processes, Department of Statistics, University of North Carolina.
- Pickands, J. (1971), The two-dimensional Poisson process and extremal processes. *J. Appl. Prob.* **8**, 745-756.
- Pickands, J. (1975), Statistical inference using extreme order statistics. *Ann. Statist.* **3**, 119-131.
- Resnick, S. (1975), Weak convergence to extremal processes. *Ann. Probab.* **3**, 951-960.
- Resnick, S. (1987), *Extreme Values, Point Processes and Regular Variation*. Springer Verlag, New York.
- Smith, R.L. (1984), Threshold methods for sample extremes. In *Statistical Extremes and Applications*, J. Tiago de Oliveira (ed.), Reidel Dordrecht, 621-638.
- Smith, R.L. (1985), Maximum likelihood estimation in a class of nonregular cases. *Biometrika*, **72**, 67-90.
- Smith, R.L. (1987), Estimating tails of probability distributions. *Ann. Statist.*, **15**, 1174-1207.
- Smith, R.L. (1989), Extreme value analysis of environmental time series: An application to trend detection in ground-level ozone (with discussion). *Statistical Science* **4**, 367-393.
- Smith, R.L. (1990), Extreme value theory. Chapter 14 of *Handbook of Applicable Mathematics*, W. Ledermann (ed.). John Wiley, Chichester, pp. 437-472.
- Smith, R.L. (1991), Max-stable processes and spatial extremes. Preprint under revision.
- Smith, R.L. (1992), The extremal index for a Markov chain. *J. Applied Probability* **29**, 37-45.
- Smith, R.L., Tawn, J.A. & Coles, S.G. (1993), Markov chain models for threshold exceedances. In preparation.
- Smith, R.L. & Weissman, I. (1994), Estimating the extremal index. To appear, *J.R. Statist. Soc. B*.
- Tawn, J.A. (1988), Bivariate extreme value theory - models and estimation. *Biometrika* **75**, 397-415.
- Tawn, J.A. (1990), Modelling multivariate extreme value distributions. *Biometrika* **77**, 245-253.
- Tiago de Oliveira, J. (1984), Bivariate models for extremes. In *Statistical Extremes and Applications*, J. Tiago de Oliveira (ed.), Reidel, Dordrecht, 131-153.
- Yun, S. (1993), The extremal index of a higher-order stationary Markov chain. Manuscript in preparation.