Extremes 2:4, 339–365, 1999

Dependence Measures for Extreme Value Analyses

STUART COLES, JANET HEFFERNAN, AND JONATHAN TAWN
Department of Mathematics and Statistics, Lancaster University, Lancaster, LA1 4YF, UK
E-mail: j.tawn@lancaster.ac.uk

[Received June 15, 1999; Revised January 11, 2000; Accepted January 28, 2000]

Abstract. Quantifying dependence is a central theme in probabilistic and statistical methods for multivariate extreme values. Two situations are possible: one where, in a limiting sense, the extremes are dependent; the other where, in the same sense, the extremes are independent. This paper comprises an overview of the principal issues through a unified approach which encompasses both these situations. Novel diagnostic measures for dependence are also developed which provide complementary information about different aspects of extremal dependence. The paper is written in an elementary style, with the methodology illustrated by application to theoretical examples and typical data-sets. These data-sets and the S-plus functions used for the analyses are available online.

Key words. asymptotic independence, bivariate extreme value distribution, copula, point processes

1. Introduction

Extreme value analyses are frequently applied in the context of modeling environmental data, for which the phenomenon of dependence is often intrinsic. Dependence occurs, for example, when different processes under study have a stochastic behavior that is linked, say, to common meteorological conditions. Dependence may also arise when a single process is studied at different spatial locations or studied in terms of its temporal evolution. For example, the sea-level is a combination of still-water level and waves, with both processes being driven by regional meteorology. Flooding occurs at any particular location when the combined still-water level and waves exceed a critical level. Whilst short-term breaches of coastal defences are likely to be sustainable, persistence in high levels of the process may cause severe damage. Furthermore, strong spatial dependence will lead to such conditions occurring simultaneously along entire coastal stretches, thus creating the potential for widespread regional damage. The total risk assessment therefore depends on whether the still-water level and wave processes arise independently or not, and the degrees of spatial and temporal dependence of each process. Similar environmental examples occur in applications to climatology, hydrology and pollution control. Examples from other fields are less common, though multivariate extreme value methods have been used for the analysis of reliability, finance and athletics data.

Our objective is the development of measures of extremal dependence for bivariate random variables (X,Y). Assuming, for the moment, that the marginal distributions of X and Y are identical, one natural measure is

$$\chi = \lim_{z \to z^*} \Pr(Y > z | X > z), \tag{1.1}$$

where z^* is the upper limit of the support of the common marginal distribution. Loosely stated, χ is the probability of one variable being extreme given that the other is extreme. In the case $\chi = 0$ the variables are said to be asymptotically independent. The importance of this class was recognized as far back as Geffroy (1958/59), Sibuya (1960), Tiago de Oliveira (1962/63) and Mardia (1964). Moreover, empirical analysis of real data often leads to estimates of $\chi = 0$. Despite this, standard methodology for multivariate extremes is based on distributions for which either $\chi > 0$, or the special case of exact independence, for which $\chi = 0$, suggesting a limitation in the applicability of standard methods. Applying models for which $\chi > 0$ to asymptotically independent data leads to the over-estimation of probabilities of extreme joint events, since there is a mis-placed assumption that the most extreme marginal events may occur simultaneously.

Statistical models for the general class of distributions having $\chi = 0$ have been developed only comparatively recently (Ledford and Tawn, 1996, 1997; Bruun and Tawn, 1998; Bortot and Tawn, 1998). Although all members of this class are asymptotically independent, at finite levels quite different degrees of dependence are attainable. Thus, we define a new quantity, $\bar{\chi}$, which gives a suitable measure of extremal dependence within the class, and develop associated diagnostics for estimation. We also establish the connections between χ and $\bar{\chi}$, and standard models for multivariate extreme values. For ease of exposition, the development throughout is limited to bivariate vectors, but the definitions and techniques extend naturally to vectors of arbitrary dimension.

The paper is structured in the following way. In Section 2 we introduce three data-sets that are used throughout the paper to illustrate the various principles and procedures. These examples demonstrate a range of contexts in which extremal dependence may arise in practice, and also illustrate quite different forms of extremal behavior. In Section 3 we derive, from elementary arguments, two fundamental measures of extremal dependence, and develop non-parametric diagnostic procedures for estimating each. One of these measures is linked to an established measure, but the second, obtained as a dual of the first, is novel. We argue that the complete pair of measures should be considered when summarizing extremal dependence. In Section 4 we relate the extremal dependence summaries to standard models and procedures for extreme value inference, showing that these measures should be an integral part of diagnostic and inference procedures. In Section 5, we discuss connections between the present work and general models for multivariate extremes. Access to the data, the S-plus functions used in the analyses in this paper and information on their use is available online at URL: http://www.math.lancs.ac.uk./~coless.

2. Dependence in extreme value data

In this section, we introduce three data-sets which are used for illustration throughout the paper. Respectively, the data-sets give examples of situations in which dependence at extreme levels is a consequence of proximity in space, proximity in time and dependence on a common covariate.

2.1. Oxford and Worthing temperature series

Figure 1 plots the annual maximum temperatures at Oxford and Worthing, both in southern England, for the period 1901–1980. The plot shows the data after the addition of a small amount of noise to reduce the effect of the data being rounded to the nearest degree Fahrenheit. The data demonstrate an apparent tendency for large maxima at one location to coincide with large maxima at the other, though this interpretation requires some care since the values themselves may not have arisen simultaneously. Nonetheless, the apparent dependence in the data suggests that spatial cohesion in the temperature process induces dependence in the series across locations. Thus, the probability of both locations experiencing a particularly high annual maximum is greater than would be expected in the case of independence; the extent to which this is so requires a measure of extremal dependence. Smith (1990) considered this issue and also examined the marginal distributions, which are found to be similar apart from a 5°F shift due to the cooler climate of Worthing which is coastal.

2.2. Rainfall time series

Figure 2 shows a 54-year time series, and associated empirical marginal distribution, of daily rainfall aggregates recorded at a single location in the south-west of England. The marginal extremal properties of these data were studied by Coles and Tawn (1996) and their temporal structure by Coles (1994). Persistence in meteorological conditions is likely to induce short term dependence in such series, and if such a phenomenon were to manifest itself at the most extreme levels, then flooding would be likely as a consequence of rainfall aggregation through time. As the rainfall distribution is long-tailed, even slight dependence could lead to substantially greater estimates of flood levels than those

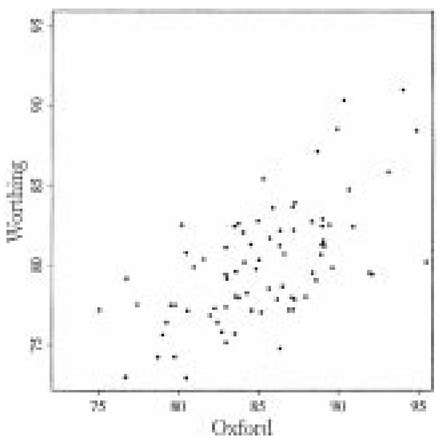

Figure 1. Oxford and Worthing annual maximum temperatures measured in degrees Fahrenheit.

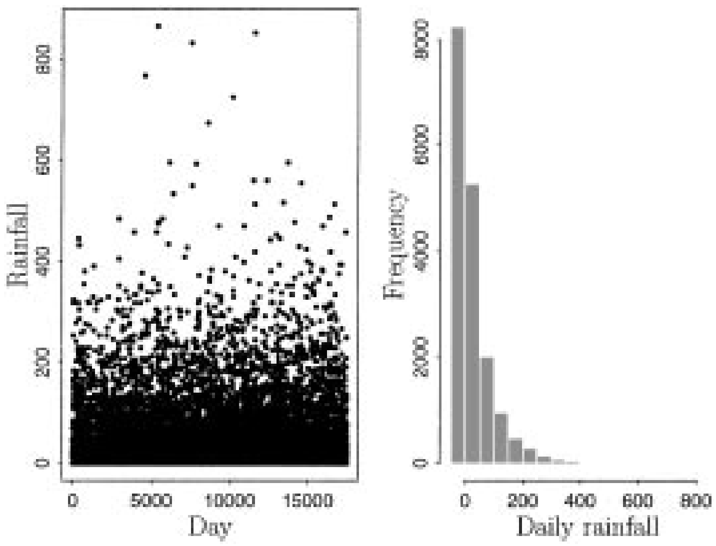

Figure 2. Time series and histogram of daily rainfall aggregates. The data are recorded in tenths of millimeters.

obtained under the assumption of independence. From Figure 3, which shows pairs of lagged values in the series, there is little apparent evidence for the most extreme levels of rainfall to occur on consecutive days. Again, though, this needs to be verified through a more formal analysis.

2.3. Wave–surge levels

Figure 4 shows a filtered series of 3-hourly measurements of the surge and wave heights at Newlyn, a coastal town in the south-west of England. The surge variable is the meteorologically-induced non-tidal component of the still-water level of the sea. Hence, both the surge and wave processes are driven by common meteorological conditions and dependence at extreme levels is likely. This phenomenon seems to be borne out by Figure 4. The motivation for the study of such data is that flooding is likely under the combined conditions of extreme surges and wave heights, though the precise combination may be rather complex. Thus, the immediate issue of quantifying the dependence between constituent variables is supplemented by the need to understand the effect of such dependence on the behavior of specified variable combinations. Further details and analyses of these particular data are given by Coles and Tawn (1994) and Bortot et al. (2000).

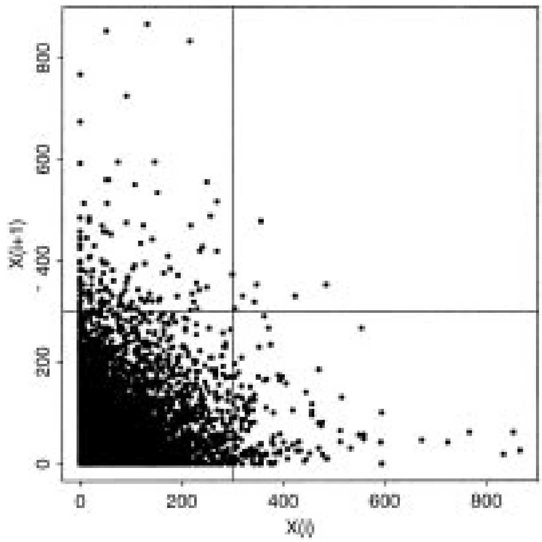

Figure 3. Scatter plot of successive values, i.e. (X_i, X_{i+1}) for $i = 1, \ldots, n-1$, in the rainfall series. The horizontal and vertical lines show the marginal threshold level, $u = 300$ suggested by univariate extreme value methods.

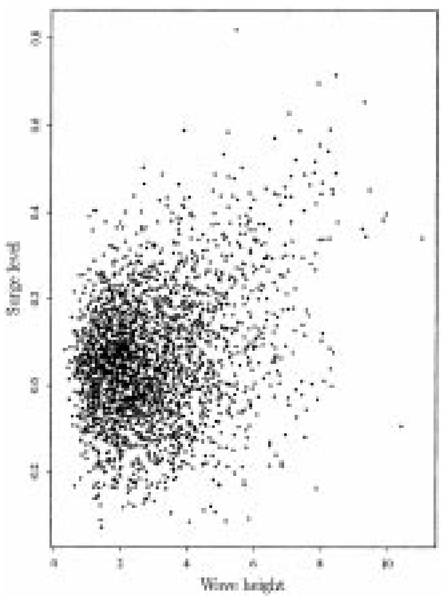

Figure 4. Newlyn: offshore wave heights and surge levels. Both variables are measured in meters.

3. Measures of extremal dependence

3.1. The copula function

For any random vector, (X,Y), the distribution function $F(x, y) = \Pr(X \leq x, Y \leq y)$ comprises a complete description of dependence between X and Y. The influence of marginal aspects can also be removed by observing that, subject to continuity conditions, there is a unique function $C(\cdot, \cdot)$ with domain $\mathscr{A} = [0, 1] \times [0, 1]$ such that

$$F(x, y) = C\{F_X(x), F_Y(y)\},$$

where F_X and F_Y are the marginal distribution functions given by

$$F_X(x) = F(x, \infty) \text{ and } F_Y(y) = F(\infty, y).$$

The function C is the copula; it contains complete information about the joint distribution of X and Y apart from the marginal distributions. In this sense, C describes association between X and Y in a form that is invariant to marginal transformation. Put differently, C is the joint distribution function of X and Y after transformation to variables U and V, with Uniform [0,1] margins, via $(U, V) = \{F_X(X), F_Y(Y)\}$. For more details see Nelsen (1998) and Joe (1997).

For standard distributions the copula is easily evaluated as in the following examples:

Independence. In this case $F(x, y) = F_X(x)F_Y(y)$, so $C(u, v) = uv$ on $\mathscr{A}$.

Perfect dependence. In this case $Y = F_Y^{-1}\{F_X(X)\}$ with probability 1, so $F(x, y) = \min\{F_X(x), F_Y(y)\}$ and $C(u, v) = \min(u, v)$ on $\mathscr{A}$.

Bivariate logistic extreme value distribution. As we will see in Section 4.3, this family is one member of the class of bivariate extreme value distributions, which arise as the class of non-degenerate limit distributions for componentwise maxima. The family has generalized extreme value distributions for F_X and F_Y, and copula

$$C(u, v) = \exp[-\{(-\log u)^{1/\alpha} + (-\log v)^{1/\alpha}\}^{\alpha}] \tag{3.1}$$

on $\mathscr{A}$. The parameter $\alpha \in [0, 1]$ determines the strength of dependence: $\alpha = 1$ gives independence; decreasing α leads to increased dependence with perfect dependence arising in the limit as $\alpha \to 0$.

Gaussian dependence model. If (X,Y) have a bivariate Normal distribution with correlation coefficient ρ, then F_X and F_Y are the distribution functions of Normal random variables, and

$$C(u,v) = \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi(1-\rho^2)^{1/2}} \exp\left\{-\frac{1}{2(1-\rho^2)}(s^2 - 2\rho st + t^2)\right\} ds\, dt$$

on $\mathscr{A}$, where $\Phi(\cdot)$ is the univariate standard Normal distribution function.

These four families of distributions are used as examples throughout the paper. The copulas for many other families of distributions are listed by Joe (1997) and Currie (1999).

3.2. Exploratory data analysis

For statistical applications it is often helpful to summarize dependence, both informally and formally, from observed data. With independent observations $(x_i, y_i), i = 1, \ldots, n$ from an unknown distribution F, it is natural to transform to Uniform marginals, leading to realizations from the associated copula C. Since the marginals of F are unknown, estimates $\hat{F}_X$ and $\hat{F}_Y$ must be used; for example, the marginal empirical distribution functions. Then the pairs $(u_i, v_i), i = 1, \ldots, n$, where

$$u_i = \hat{F}_X(x_i) \text{ and } v_i = \hat{F}_Y(y_i)$$

are independent realizations with approximate distribution C. An informal picture of extremal dependence may then be obtained by examining the large values of u_i and v_i.

Plots of (u_i, v_i) for each of the three data-sets discussed in Section 2 are shown in Figure 5. On these scales the conclusions drawn previously are broadly re-enforced: strong extremal dependence for the temperature data, little evidence for extremal dependence in the rainfall series, and weak but evident extremal dependence for the oceanographic data. The apparent non-uniformity in the marginal behavior of the rainfall data is a consequence of the large number of ties caused by observations of zero rainfall; the plot also suffers from severe discretization of the data. For the wave-surge data the increased density of points close to (1,1) now suggests much more clearly a tendency for the most extreme levels to be associated. Furthermore, the reasonably uniform scatter elsewhere in $\mathscr{A}$ suggests the variables are nearly independent for non-extreme values.

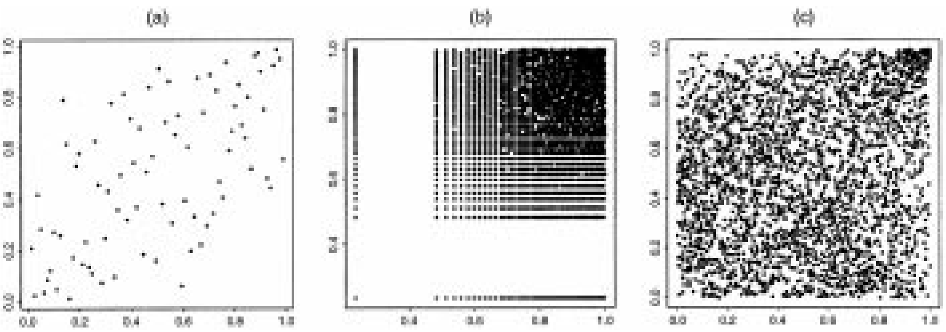

Figure 5. Data sets transformed to have uniform marginal distributions: (a) temperatures, (b) rainfalls, (c) waves and surge levels.

3.3. Summarizing dependence

For both interpretation and inference it is often useful to reduce the information contained in the copula, C, either to a one-dimensional function, or even a single parameter. We examine two elementary measures, χ and $\bar{\chi}$, which provide measures of different aspects of extremal dependence. We argue that both measures are needed in order to obtain a summary that is informative for variables which may be either asymptotically independent or asymptotically dependent.

3.3.1. Dependence measure χ. One natural summary of extremal dependence is the coefficient χ in equation (1.1). This is generalized to the case of non-identically distributed pairs (X,Y) by transformation to Uniform margins (U,V) and setting

$$\chi = \lim_{u\to 1} \Pr(V > u | U > u).$$

It is more convenient, however, to obtain χ as the limit of an alternative, asymptotically equivalent, function. Observe that:

$$\begin{aligned}\Pr(V > u | U > u) &= \frac{\Pr(U > u, V > u)}{\Pr(U > u)} \\ &= \frac{1 - 2u + C(u,u)}{1 - u} \\ &= 2 - \frac{1 - C(u,u)}{1 - u} \\ &\sim 2 - \frac{\log C(u,u)}{\log u}\end{aligned}$$

as $u \to 1$. Hence, defining

$$\chi(u) = 2 - \frac{\log \Pr(U < u, V < u)}{\log \Pr(U < u)} \quad \text{for } 0 \le u \le 1, \tag{3.2}$$

it follows that

$$\chi = \lim_{u\to 1} \chi(u). \tag{3.3}$$

More than just providing the limit χ, the function $\chi(u)$ can also be interpreted as a quantile-dependent measure of dependence. In particular, the sign of $\chi(u)$ determines whether the variables are positively or negatively associated at the quantile level u, with bounds

$$2 - \log(2u - 1)\big/\log(u) \le \chi(u) \le 1,$$

where the lower bound is interpreted as $-\infty$ for $u \le 1/2$, and 0 for $u = 1$.

It is instructive to determine $\chi(u)$ for each of the earlier models. For independent variables $\chi(u) = 0$; for perfect dependence $\chi(u) = 1$; and for the bivariate logistic extreme value distribution $\chi(u) = 2 - 2^{\alpha}$. In each of these three cases $\chi(u)$ is constant in u, which turns out to be the case for any distribution falling in the class of bivariate extreme value distributions (see Section 4.3). This explains the advantage of the $\chi(u)$ formulation: assessment of the constancy of empirical estimates of $\chi(u)$ provides a diagnostic check for membership of the bivariate extreme value class.

More generally, $\chi(u)$ is a non-trivial function of u. For example, in the case of the Gaussian dependence model, $\chi(u)$ is an integral expression that depends on u and which requires numerical evaluation. It is plotted in Figure 6 for a range of values of the correlation coefficient, ρ. For all values of u, stronger dependence, as measured by $\chi(u)$, is obtained by increasing ρ. For intermediate values of u, $\chi(u)$ is reasonably linear with distinctly different values for all ρ. However, as $u \to 1$, the effect of dependence is diminished, with $\chi(u) \to 0$ for all $\rho < 1$. For $\rho > 0$ the convergence is very slow, and ultimately abrupt; hence, $\chi(u)$ is considerably greater than zero for u close to 1. This has practical implications since estimates of this measure will be derived from empirical observations for which $u < 1$. Thus, estimates of $\chi(u)$ may appear constant and positive, even for asymptotically independent variables.

3.3.2. Dependence measure $\bar{\chi}$. Recent research has highlighted the importance of the class of asymptotically independent distributions in multivariate extreme value modeling. By definition, $\chi = 0$ within this class, so the measure χ is unable to provide information on relative strength of dependence for such models. To overcome these limitations we define a second dependence measure. Denoting the joint survivor

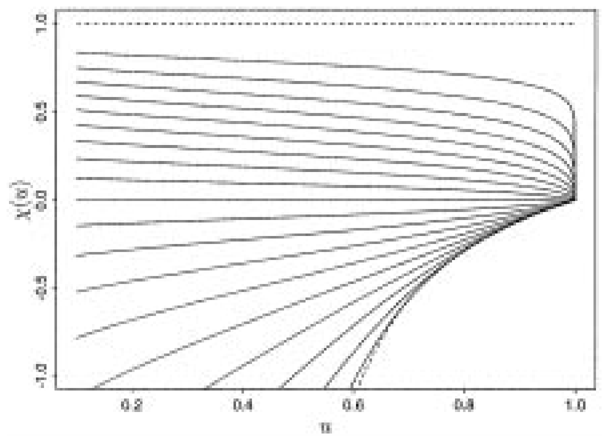

Figure 6. The dependence measure $\chi(u)$ for the Gaussian dependence model: the curves shown (bottom to top) correspond to $\rho = -0.9, -0.8, \ldots, 0.9$. The upper and lower bounds on $\chi(u)$ are shown as dashed lines.

function $\Pr(X > x, Y > y)$ by $\bar{F}(x, y)$, we may write

$$\begin{aligned}\bar{F}(x,y) &= 1 - F_X(x) - F_Y(y) + F(x,y)\\ &= \bar{C}\{F_X(x), F_Y(y)\},\end{aligned} \tag{3.4}$$

where

$$\bar{C}(u,v) = 1 - u - v + C(u,v).$$

Now, by analogy with (3.2) and (3.3), we define

$$\bar{\chi}(u) = \frac{2\log\Pr(U > u)}{\log\Pr(U > u, V > u)} - 1 = \frac{2\log(1-u)}{\log\bar{C}(u,u)} - 1 \text{ for } 0 \leq u \leq 1,$$

where $-1 < \bar{\chi}(u) \leq 1$ for all $0 \leq u \leq 1$; the precise definition is chosen for scaling convenience. To focus on extremal characteristics, analogous to (3.3), we also define

$$\bar{\chi} = \lim_{u \to 1} \bar{\chi}(u),$$

for which $-1 < \bar{\chi} \leq 1$.

For asymptotically dependent variables $\bar{\chi} = 1$: the examples of perfect dependence and the bivariate logistic extreme value distribution are easily seen to satisfy this. The more useful application is to asymptotically independent distributions, for which $\bar{\chi}$ provides a measure that increases with dependence strength. For example, in the case of independent variables, $\bar{C}(u,v) = (1-u)(1-v)$ on $\mathscr{A}$, so $\bar{\chi}(u) = 0$ identically for $u \in [0,1]$ and $\bar{\chi} = 0$. For the Gaussian dependence model,

$$\bar{C}(u,v) = \int_{\Phi^{-1}(u)}^{\infty}\int_{\Phi^{-1}(v)}^{\infty} \frac{1}{2\pi(1-\rho^2)^{1/2}} \exp\left\{-\frac{1}{2(1-\rho^2)}(s^2 - 2\rho st + t^2)\right\} ds\, dt,$$

and it can be verified that

$$\bar{C}(u,u) \sim c_\rho\{-\log(1-u)\}^{-\rho/(1+\rho)}(1-u)^{2/(1+\rho)} \text{ as } u \to 1,$$

where $c_\rho = (1+\rho)^{3/2}(1-\rho)^{-1/2}(4\pi)^{-\rho/(1+\rho)}$ (Ledford and Tawn 1996; 2000). Thus, $\bar{\chi} = \rho$, which provides a useful benchmark for interpreting the magnitude of $\bar{\chi}$ in general models.

In Figure 6 we found that the convergence of $\chi(u) \to \chi = 0$ was very slow for the Gaussian dependence structure. Hence, in Figure 7, we plot $\bar{\chi}(u)$ for the Gaussian dependence model to see if the limiting behavior is more apparent at sub-asymptotic levels. Although there is a rapid change in behavior as $u \to 1$, it is clear that $\bar{\chi}(u)$ is approximately linear for $u > 0.5$ and bounded from 1. The plot also shows that estimation

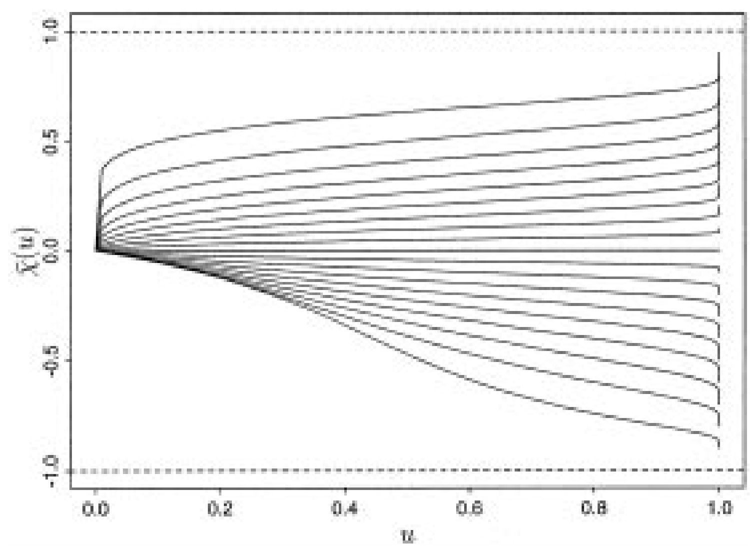

Figure 7. The dependence measure $\bar{\chi}(u)$ for the Gaussian dependence model: the curves shown (bottom to top) correspond to $\rho = -0.9, -0.8, \ldots, 0.9$. The upper and lower bounds on $\bar{\chi}(u)$ are shown as dashed lines.

of $\bar{\chi}$ for data from this distribution is likely to be biased; this is consistent with the findings of Ledford (1999). Furthermore, by comparison with Figure 6, it is much easier to identify from $\bar{\chi}(u)$ whether the variables are asymptotically independent than from $\chi(u)$.

In summary, χ is on the scale $[0, 1]$, with the set $(0, 1]$ corresponding to asymptotic dependence, and the measure $\bar{\chi}$ falls within the range $[-1, 1]$, with the set $[-1, 1)$ corresponding to asymptotic independence. Thus the complete pair $(\chi, \bar{\chi})$ is required as a summary of extremal dependence: $(\chi > 0, \bar{\chi} = 1)$ signifies asymptotic dependence, in which case the value of χ determines a measure of strength of dependence within the class; alternatively, $(\chi = 0, \bar{\chi} < 1)$ signifies asymptotic independence, in which case the value of $\bar{\chi}$ determines the strength of dependence within this class.

3.4. Data examples

Simple empirical estimates of the functions $\chi(u)$ and $\bar{\chi}(u)$ can be constructed on the basis of observed data by using the empirical estimate of $C(u, u)$. Analyzing the behavior of these as $u \to 1$ leads to an informal picture of extremal dependence. The confidence intervals for these estimates of $\chi(u)$ and $\bar{\chi}(u)$ are constructed assuming independence of the observations, that each marginal distribution is estimated exactly by its empirical distribution function, and that the sampling distribution of a proportion is well-approximated by its asymptotic distribution. The construction of the confidence interval uses the delta method. For inference purposes assuming the data are independent appears to be a reasonable assumption for these data, even for the rainfall pairs since the temporal dependence is weak. Ignoring dependence leads to under-estimation of the confidence interval width. As the uncertainty in the marginal distribution is ignored, the intervals will be too narrow, although due to the approximate orthogonality of marginal and dependence features this under-estimation also should be slight. Furthermore, since the estimate of $C(u, u)$ is not normally distributed for u near 0 or 1, care is taken not to draw too strong conclusions from these intervals. For the three data-sets described previously, corresponding plots of estimates and confidence intervals are shown in Figures 8–10.

For the temperature data it appears that $\chi(u) \approx 0.5$ for all u, although the pointwise confidence intervals cover the full range of possible values for χ, so the conclusions are only tentative. The value of $\bar{\chi} = 1$ also seems plausible as a limit of $\bar{\chi}(u)$. As the data are componentwise annual maxima, it is expected that a bivariate extreme value distribution, such as the logistic model, may be appropriate. In fact, the patterns of $\chi(u)$ and $\bar{\chi}(u)$ are consistent with the bivariate logistic extreme value distribution with $\alpha \approx 0.6$.

For the successive values of the rainfall series there is some evidence of dependence since $\chi(u) > 0$ for $u < 1$. The behavior of the graphs of Figure 9 for $u < 0.5$ is explained by the large number of zeroes in the data. The series appears to be asymptotically independent as $\chi = 0$ and $\bar{\chi} < 1$. However, as $\bar{\chi}$ seems slightly greater than zero, the evidence supports some dependence in the series at extreme levels.

The exploratory diagnostics employed here do not offer such clear conclusions about the wave-surge data. For low values of u, $\chi(u) < 0$, indicating slight negative dependence, but for large u it appears that $\chi(u) \approx 0.3$, a feature consistent with an asymptotically dependent distribution that is not itself a bivariate extreme value distribution. However, it seems that $0 < \bar{\chi} < 1$, which contradicts the conclusion of asymptotic dependence. Thus, from such an informal analysis, it is difficult to decide between asymptotic dependence and asymptotic independence for these data.

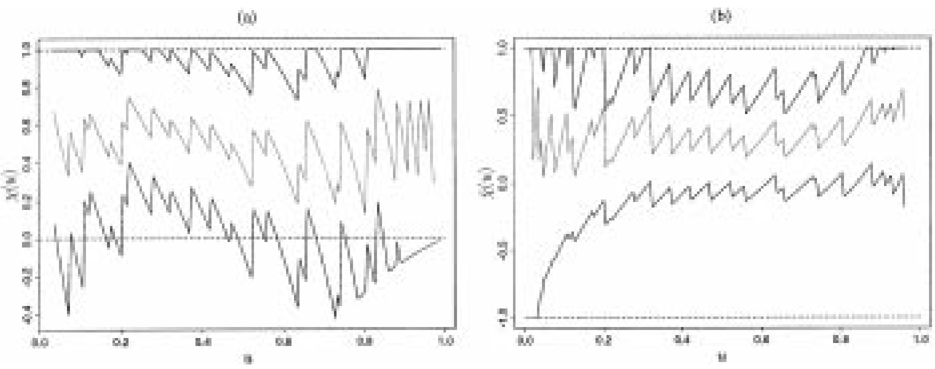

Figure 8. The dependence measures (a) $\chi(u)$ and (b) $\bar{\chi}(u)$ for the temperature data. The dotted and solid lines show the estimate and 95% pointwise confidence intervals respectively. The values $\chi(u) = 0$ and 1 and $\bar{\chi}(u) = -1$ and 1 are shown as dashed lines for reference.

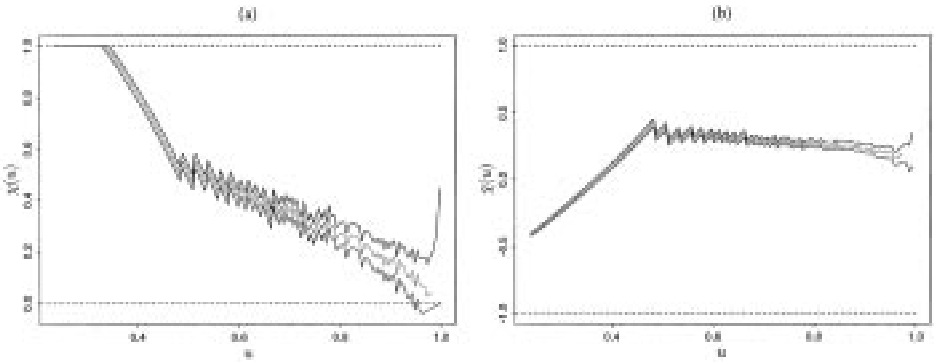

Figure 9. The dependence measures (a) $\chi(u)$ and (b) $\bar{\chi}(u)$ for the rainfall data. The dotted and solid lines show the estimate and 95% pointwise confidence intervals respectively. The values $\chi(u) = 0$ and 1 and $\bar{\chi}(u) = -1$ and 1 are shown as dashed lines for reference.

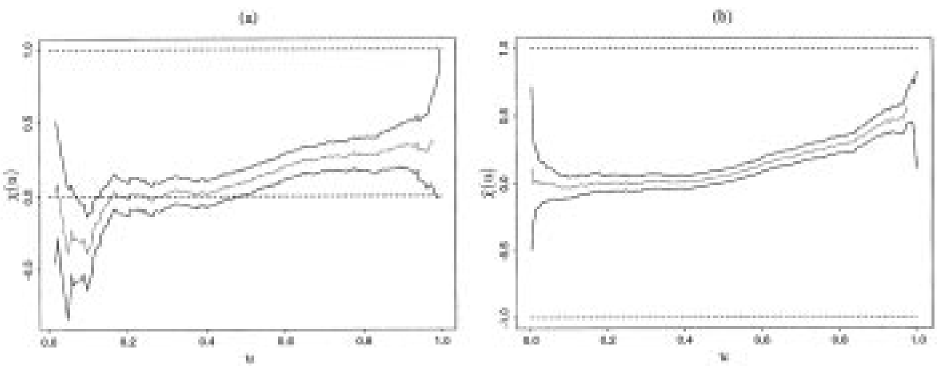

Figure 10. The dependence measures (a) $\chi(u)$ and (b) $\bar{\chi}(u)$ for the wave and surge level data. The dotted and solid lines show the estimate and 95% pointwise confidence intervals respectively. The values $\chi(u) = 0$ and 1 and $\bar{\chi}(u) = -1$ and 1 are shown as dashed lines for reference.

4. Extremal dependence: theory and models

Having obtained natural measures of extremal dependence that are limiting values of simple dependence functions, we now establish the connection between these measures and standard models for bivariate extremes. In particular, we consider the roles of the parameters χ and $\bar{\chi}$ in the context of exploratory and formal inference for extreme value models. Essentially, there are two classes of models: models for componentwise block maxima, in which the raw data are themselves extreme values; and threshold methods, in which a complete series of data is available, and extreme value models are used to characterize the series above high marginal thresholds only. We consider each class separately in Sections 4.3 and 4.4 respectively.

As in earlier sections it is convenient to dis-entangle dependence and marginal aspects. In practice data may again be transformed marginally to have standard distributions, using empirical, parametric or semi-parametric methods, prior to a study of dependence. There may be some gain in estimation efficiency by combining the marginal and dependence estimation into one inferential step, but for presentation it is convenient to assume that data have already been transformed to a standard form. Furthermore, since our dependence measures are marginally invariant, there is no loss of generality in imposing a fixed marginal scale. In contrast to the earlier sections in which uniform margins were used, it is now convenient to assume a standard Fréchet marginal scale, so that X and Y each have distribution function $F(z) = \exp(-1/z)$ for $z > 0$, leading to $\Pr(X > z) = \Pr(Y > z) \sim z^{-1}$ as $z \to \infty$.

4.1. *Links with χ and $\bar{\chi}$*

With (X,Y) having Fréchet marginal distributions it follows from Section 3.3.1 that

$$\chi = \lim_{z \to \infty} \Pr\{Y > z \mid X > z\}.$$

There are also a variety of alternative summaries of asymptotic dependence; see de Haan (1985), Tawn (1988) and Weintraub (1991). Since each is implicitly related to χ however, it follows that each has the same limitation in being identical across the class of asymptotically independent distributions.

It is less well-known that $\bar{\chi}$ also arises in connection with asymptotic models for (X,Y). In particular Ledford and Tawn (1996, 1997, 1998), de Haan and de Ronde (1998) and Peng (1999) each demonstrate that, under broad conditions, the joint survivor function of an arbitrary random pair (X,Y), with unit Fréchet marginal distributions, satisfies the asymptotic condition

$$\Pr\{X > z, Y > z\} \sim \mathscr{L}(z)\{\Pr(X > z)\}^{1/\eta} \text{ for large } z, \tag{4.1}$$

where $\mathscr{L}(z)$ is a slowly varying function as $z \to \infty$ and η, the coefficient of tail dependence, lies in the range $(0,1]$. It follows from (4.1) and (3.4) that

$$\bar{C}(u,u) \sim \mathscr{L}\{(1-u)^{-1}\}(1-u)^{1/\eta} \text{ as } u \to 1.$$

Hence,

$$\bar{\chi}(u) \sim \frac{2\log(1-u)}{\{\log \mathscr{L}\{(1-u)^{-1}\} + \frac{1}{\eta}\log(1-u)\}} - 1 \to 2\eta - 1, \qquad \text{as } u \to 1,$$

and so $\bar{\chi} = 2\eta - 1$. Furthermore, if $\eta = 1$ and $\mathscr{L}(z) \to c$ as $z \to \infty$, with $0 < c \leq 1$, then $(\chi = c, \bar{\chi} = 1)$, and the variables are asymptotically dependent of degree c.

The work of Ledford and Tawn identified the parameter η as pivotal in characterizing extremal dependence. By relating η to $\bar{\chi}$ we have shown in this paper that the dependence parameter for asymptotically independent variables can also be motivated by the same elementary considerations that were used to obtain χ.

4.2. *Parametric inference for χ and $\bar{\chi}$*

Section 3 described informal graphical methods for determining whether data are asymptotically independent or asymptotically dependent. We now develop formal assessment procedures for this, based on ideas in Ledford and Tawn (1996).

Let $T = \min(X, Y)$. Then from the joint tail condition (4.1),

$$\Pr(T > z) = \Pr\{X > z, Y > z\} \sim \mathscr{L}(z) z^{-1/\eta} \text{ as } z \to \infty. \tag{4.2}$$

Hence, η is the shape parameter of the univariate variable T, and standard univariate extreme value techniques applied to the variable T lead to inferences on η (or equivalently $\bar{\chi} = 2\eta - 1$) and hence on the asymptotic status of dependence in the pair (X, Y). Possible standard estimators to use for shape parameter estimation are threshold-based likelihood estimators (Davison and Smith, 1990) and the semi-parametric estimators of Dekkers et al. (1989). Since simulation studies indicate similar performance across such estimators, we base our analyses on likelihood methods in the subsequent presentation.

It is also clear from equation (4.2) that the parameter χ may be estimated as the scale parameter of the univariate variable T, subject to fixing $\eta = 1$ and setting $\mathscr{L}(z) = \chi$ for large z. However, with this choice, (4.2) is only correct to first order, so the procedure is inappropriate for asymptotically independent variables, since then $\chi = 0$ and it is the second-order term which is dominant. Hence, both $\bar{\chi}$ and χ should be considered together, with the estimate of χ being valuable only when the estimate of $\bar{\chi}$ is not significantly less than 1. Consequently, an appropriate way to test $\chi = 0$ against $\chi > 0$ is through a test of $\bar{\chi} < 1$ against $\bar{\chi} = 1$. This can easily be performed using a generalized likelihood ratio test based on tail model (4.2).

Inference for χ and $\bar{\chi}$ above, and for the parametric dependence models in the following sections, is based on standard likelihood methods. Consequently, the standard errors quoted are based on the asymptotic normality of maximum likelihood estimators. However, the uncertainty this measures only reflects the assumptions made in constructing the likelihood function. Here we assume that observations are independent and that each marginal distribution is estimated exactly by its empirical distribution function. These assumptions, discussed in Section 3.4, together with the assumption that the parametric specification of the model is correct, are each likely to lead to a slight under-estimation of the true uncertainty of the estimates of dependence features.

The procedure for subsequent analysis now depends on whether the data are componentwise block maxima or a complete series: these are considered separately in Sections 4.3 and 4.4.

4.3. Modeling componentwise block maxima

This methodology is appropriate when, for example, only the annual maxima of a process recorded at, say, two locations are available. The Oxford and Worthing temperature data are of precisely this form. Asymptotic models for such data are derived using arguments analogous to the classical univariate result. Thus, if $(X_1, Y_1), (X_2, Y_2), \ldots, (X_n, Y_n)$ is an independent and identically distributed series of random vectors with standard Fréchet margins, let

$$M_{X,n} = \max\{X_1, \ldots, X_n\} \text{ and } M_{Y,n} = \max\{Y_1, \ldots, Y_n\},$$

and define the vector of componentwise maximum by $\mathbf{M}_n = \{M_{X,n}, M_{Y,n}\}$. Then, as $n \to \infty$, subject to weak regularity conditions, the limiting distribution of the normalized vector $n^{-1}\mathbf{M}_n$ has distribution function within the bivariate extreme value class. That is,

$$\Pr(M_{X,n}/n \le x, M_{Y,n}/n \le y) = \{F(nx, ny)\}^n \to G(x, y) \text{ as } n \to \infty,$$

with

$$G(x, y) = \exp\{-V(x, y)\}, \tag{4.3}$$

where

$$V(x, y) = \int_0^1 \max\left(\frac{w}{x}, \frac{1-w}{y}\right) 2dH(w) \tag{4.4}$$

for some distribution function H on the interval $[0, 1]$ satisfying the moment constraint

$$\int_0^1 w dH(w) = 1/2. \tag{4.5}$$

Thus, either of the functions V or H determine dependence in this limiting representation. If $H(w)$ has a density function $h(w)$ then

$$h\left(\frac{x}{x+y}\right) = -\frac{(x+y)^3}{2} \frac{\partial^2 V(x, y)}{\partial x \partial y},$$

so it is straightforward to move between V and H in the bivariate case. Coles and Tawn

(1991) give the equivalent relationship in the multivariate case and when H has atoms of mass on its boundaries.

Using the property that V must be homogeneous of order -1, it is easy to establish that, except for the special case of independence, all bivariate extreme value distributions are asymptotically dependent, and that therefore $\bar{\chi} = 1$ for the entire family. Furthermore,

$$\chi = 2 - V(1,1),$$

or equivalently,

$$\chi = 2\int_0^1 \min(w, 1-w)dH(w),$$

so that χ is a single parameter summary of either V or H. For example, the results stated previously for the bivariate logistic extreme value distribution are a special case of this result, with

$$V_\alpha(x,y) = (x^{-1/\alpha} + y^{-1/\alpha})^\alpha, \tag{4.6}$$

$$H_\alpha(w) = \frac{1}{2}\left[\{w^{(1-\alpha)/\alpha} - (1-w)^{(1-\alpha)/\alpha}\}\{w^{1/\alpha} + (1-w)^{1/\alpha}\}^{\alpha-1} + 1\right], \tag{4.7}$$

where $0<\alpha<1$. Thus, $V_\alpha(1,1) = 2^\alpha$ and $\chi = 2 - 2^\alpha$. Similarly, when $\alpha = 1$, corresponding to independence, $V_1(x,y) = x^{-1} + y^{-1}$ and H_1 consists of half-unit mass atoms at $\{0\}$ and $\{1\}$, i.e. $H_1(\{0\}) = H_1(\{1\}) = 1/2$ and $\chi = 0$.

The usual application of limit result (4.3) is to assume that it is the exact distribution function of marginally transformed componentwise maxima over blocks of large, but finite, length n. Assuming this asymptotic argument is reasonable leads either to independence, for which $\chi = \bar{\chi} = 0$, or asymptotic dependence, for which $\bar{\chi} = 1$ and $\chi > 0$, with larger values of χ indicating stronger dependence. Thus, within this family, and for modeling this format of data, it is χ rather than $\bar{\chi}$ which serves as a meaningful dependence measure. Since $\chi(u)$ is constant for any member of this family, evidence of non-constancy in $\chi(u)$ is indicative of a lack of model fit. Moreover, in cases where the estimate of $\bar{\chi}$ is significantly less than 1, this also suggests that a bivariate extreme value distribution is not a good model. This situation may seem unlikely, but arises when (X,Y) are asymptotically independent and n, the block size, is not sufficiently large enough to apply the asymptotic arguments leading to (4.3).

Both parametric and nonparametric procedures are available for model estimation. An added complication of nonparametric procedures is that they do not necessarily provide estimators that are consistent with the asymptotic characterization of extreme value models (4.3); see Pickands (1981) and Capéraà et al. (1997). Parametric modeling is more straightforward. This consists of selecting a parametric family for V (or, equivalently, H), and likelihood-based inference on the model parameters. The assumption of a finite parameter space means that only a sub-family of the entire class of bivariate extreme value distributions is achievable with a parametric model, and there is some art in obtaining models that are tractable and flexible, while satisfying the functional constraints. One commonly used family is the logistic model; this has $V(x, y)$ given by expression (4.6) and copula (3.1). As already observed, this family spans the space of χ, implying some degree of flexibility, but it does have the limitation of being exchangeable for all values of α. Other parametric models for V or H can be found in Coles and Tawn (1991, 1994) or Joe (1990, 1994).

As an example of this methodology we consider the Oxford and Worthing annual maximum temperature series. A 75% threshold for T leads to the maximum likelihood estimate $\hat{\eta} = 0.980$ with a standard error of 0.448. Consequently, there is no evidence to suggest $\bar{\chi} \neq 1$, though the power of this test is extremely low. This is generally the case with componentwise block maxima modeling: the data are too sparse to obtain strong evidence about model validity. In this example we tentatively proceed with a bivariate extreme value model, assuming, in particular, that the logistic model is appropriate. Supporting evidence for this assumption is that the points in Figure 5 are reasonably symmetric around the line $u = v$, consistent with the exchangeability property of the logistic model. The maximum likelihood estimate of α is $\hat{\alpha} = 0.585$ with a standard error of 0.053. This conclusion is consistent with the earlier empirical analysis: from Section 3.4, $\chi = 2 - 2^{\alpha}$, and replacing α with its estimate leads to $\hat{\chi} = 0.500$, a value consistent with the plot in Figure 8. Adopting the informal interpretation of χ, this corresponds to a reasonably strong degree of dependence, even within the asymptotically-dependent class.

4.4. Threshold methods

If an entire series of vector measurements is available, rather than just block maxima, then improvements in efficiency and flexibility can be obtained by using more general point process characterizations of extremal behavior. As above, let $(X_1, Y_1), (X_2, Y_2), \ldots$ be an independent series of realizations of the random vector (X, Y) on $\mathbb{R}_+^2$ with standard Fréchet marginal distributions. The joint tail behavior can be characterized by specifying a sequence of point processes on $\mathbb{R}_+^2$:

$$P_n = \left\{ \left(\frac{X_i}{n}, \frac{Y_i}{n} \right) \quad : \quad i = 1, \ldots, n \right\}.$$

As $n \to \infty$, $P_n \to P$ on $\mathbb{R}_+^2 / \{\mathbf{0}\}$, where P is a Poisson process. Furthermore, the intensity function of this limiting process has a particular structure, most easily stated on coordinate transformation to ''radial'' and ''angular'' components:

$$R = (X + Y)/n \text{ and } W = X/(X + Y).$$

In this coordinate system the intensity function of P is given by

$$\nu(dr \times dw) = \frac{dr}{r^2} \times 2\, dH(w), \tag{4.8}$$

where H is the dependence measure of the associated componentwise block maxima vector (cf. equation (4.4)), and must therefore satisfy the functional constraint (4.5). Thus, the intensity function ν factorises into two components: r^{-2} corresponding to the choice of standard Fréchet margins; and $dH(w)$, determining dependence at asymptotic levels of the underlying variable (X,Y). This result was derived by de Haan (1985) and used as the basis of inference by Coles and Tawn (1991, 1994) and Joe et al. (1992).

The most important characteristic for model extrapolation is the decay rate of joint tail probabilities. It follows immediately from (4.8) that if $A \subset B_1$, where $B_1^c = \{(x,y) \ : \ x<x_0, y<y_0\}$ for large x_0 and y_0, then

$$\Pr\{(X,Y) \in tA\} = \Pr\{(X/t, Y/t) \in A\} \approx \frac{1}{t}\Pr\{(X,Y) \in A\} \tag{4.9}$$

for all $t \geq 1$. Consequently, from this limiting point process characterization, joint tail probabilities decay at the same rate regardless of the form of limiting dependence, which affects only the scale factor $\Pr\{(X,Y) \in A\}$.

Inference for the limiting Poisson process model may proceed in various ways under the basic assumption that the limiting Poisson process is a valid approximation above high enough thresholds. One possibility is to adopt a parametric family for H, such as the logistic model (4.7). Then the associated Poisson process likelihood may be constructed from which the parameters can be inferred. Alternative nonparametric procedures have been proposed by Einmahl et al. (1997) and de Haan and de Ronde (1998). The basis of such methods is to map nonparametric estimates of probabilities within a set A that contains data, via equation (4.9), to obtain probability estimates for the set tA that may contain no observed data.

In theory this procedure works for both asymptotically independent and dependent variables, but there are practical difficulties in the case of asymptotic independence. As we have seen in Section 4.3, the corresponding dependence measure H is degenerate, with atoms of mass on $\{0\}$ and $\{1\}$. This leads to problems of inference, since each P_n will have a behavior that is inadmissible for the limit process P. Moreover, even if the independence limit is correctly identified, probabilities of events are likely to be poorly estimated at sub-asymptotic levels using equation (4.9). Both these difficulties can be avoided by extending the limiting point process representation to take account of the degree of dependence within the class of asymptotically independent distributions. Thus, if η is the coefficient of tail dependence corresponding to (X,Y), following Ledford and Tawn (1997), we define a sequence of point processes on $\mathbb{R}_+^2$ by

$$\tilde{P}_n = \left\{ \left(\frac{X_i}{n^\eta}, \frac{Y_i}{n^\eta} \right) \ : \ i = 1, \ldots, n \right\}.$$

Then, on restriction to the interior of $\mathbb{R}^2_+$, $\tilde{P}_n \to \tilde{P}$ as $n \to \infty$, where $\tilde{P}$ is a Poisson process. Using coordinate system (4.4) the intensity function of this limiting process is

$$\tilde{\nu}(dr \times dw) = \frac{dr}{r^{(1+\eta)/\eta}} \times d\tilde{H}(w). \tag{4.10}$$

Thus, we obtain a similar representation as in the case of asymptotic dependence, but with a different normalization of the process and a consequent modification to the rate at which ν decays with r. Furthermore, the associated angular measure $\tilde{H}$ has different constraints from those of H in the limit (4.8); see Ledford and Tawn (1997) for details.

An analogous property to (4.9) can also be derived for this model. In this case, for $A \subset B_2$, where

$$B_2 = \{(x, y) \ : \ x + y > r_0, w_0 \le x/(x+y) \le 1 - w_0\}, \tag{4.11}$$

for large r_0 and small $w_0 > 0$, it follows from intensity (4.10) that

$$\Pr\{(X, Y) \in tA\} \approx \frac{1}{t^{1/\eta}} \Pr\{(X, Y) \in A\} \tag{4.12}$$

for $t \ge 1$. Consequently, the decay rate of tail probabilities is determined by the coefficient of tail dependence, or equivalently by $\bar{\chi}$. For fixed $\bar{\chi}$, the impact of other aspects to the dependence structure is to determine the scale factor $\Pr\{(X, Y) \in A\}$. The extra generality in the decay of joint tail probabilities in representation (4.12), relative to (4.9) is, however, slightly offset by the additional restriction that the region A is bounded away from the axes $x = 0$ and $y = 0$.

Apart from the additional issue of the estimation of η, considerations for inference remain the same as for the original Poisson process model. For parametric estimation, representation (4.10) leads to a likelihood over regions for which the Poisson limit is sustainable, with the advantage that η can be estimated simultaneously with $\tilde{H}$, leading to improved estimation efficiency for η as inference is based on the complete data in the joint tail region. Alternatively, using the procedure of Section 4.2 to estimate η, for example, the approximation (4.12) may be used to map nonparametric estimates of probabilities of large events to extreme events, as described previously.

Each of the rainfall and wave-surge data-sets are of a format for which the point process model is plausible. Strictly, inference for the rainfall data is complicated by the fact that the data represent lagged pairs in the original series, which induces dependence across the points; we will ignore this issue here. After conversion to Fréchet margins, estimates of $\bar{\chi}$ based on the likelihood method of Section 4.2, and standard errors in parentheses, are 0.00 (0.044) and 0.77 (0.138) for the rainfall and wave-surge data respectively, each being consistent with the empirical estimates obtained earlier. Both estimates were based on a 75% threshold, but these estimates were found to be stable over a wide range of threshold choices. Consequently, it seems appropriate to model the wave-surge data using the classical point process representation, whilst the rainfall data would be better handled using the modified point process representation that incorporates asymptotic independence.

Further model assessment requires an examination of the (R,W) pairs after transformation to Fréchet marginals, focusing, in particular, on the distribution of the W values conditional on large values of R. Figure 11 shows kernel density estimates of $W\,|\,R>r$, for a range of values of r, for the rainfall and wave-surge data-sets. All density estimates are based on the same smoothing parameter and have been adjusted for end-effects. For the rainfall data the conditional density is clearly degenerating to spikes of mass at 0 and 1 as r increases. In contrast, for the wave-surge data, the distribution of W values is relatively stable as r increases. This seems to confirm our earlier assessment of the asymptotic forms of dependence for these data-sets.

The detailed structure of the dependence can be further examined by estimating the intensities of the appropriate limiting point process representation, i.e. intensities (4.8) and (4.10) respectively. Figure 12 shows that for the wave-surge data the estimated parametric logistic model, with $\hat{\alpha} = 0.66\ (0.013)$, provides a reasonable description of the variation of W for values exceeding the 0.9 quantile of the variable R by comparison with the kernel density estimate for these data. For such a high threshold the common choice of smoothing parameter used in Figure 11 is probably under-smoothing the kernel estimate, hence the apparent over-estimation by the logistic model for $W \in (0.2, 0.8)$. Owing to the asymptotic independence of the rainfall data, the detailed behavior of the W values can only be studied for large R subject to the exclusion of W values close to the boundaries of $[0, 1]$. Figure 13 shows estimates of the density of $W\,|\,(R>r, 0.1 \leq W \leq 0.9)$ for a range of threshold choices of r. If the limiting Poisson process were a good approximation then the density would be stable with respect to r. Although less stable than the density for the wave-surge data, the distribution is more stable than the unrestricted conditional density. This kernel density provides an estimate of $d\tilde{H}(w)\big/dw$.

Finally, consider the issue of extrapolation. The methods described here can be used

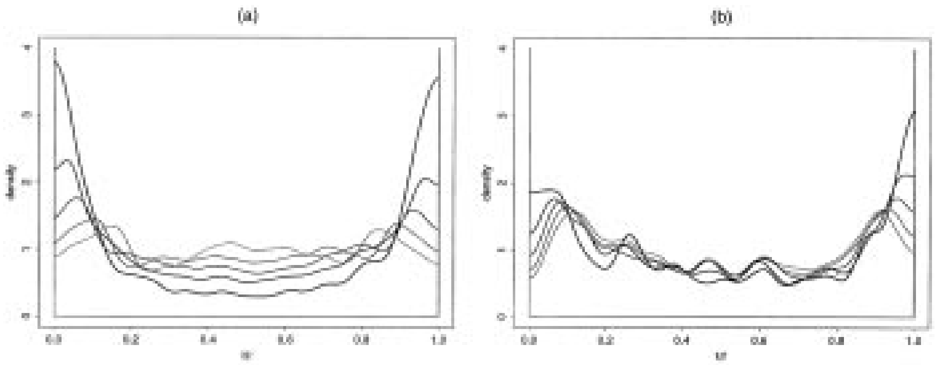

Figure 11. The estimated conditional density of $W\,|\,R>r$: (a) for the rainfall data, (b) for the wave and surge level data. The different estimates correspond to different values of r; the thin line through to the thick line correspond to r being the $0.5, 0.6, \ldots, 0.9$ quantile of the variable R respectively.

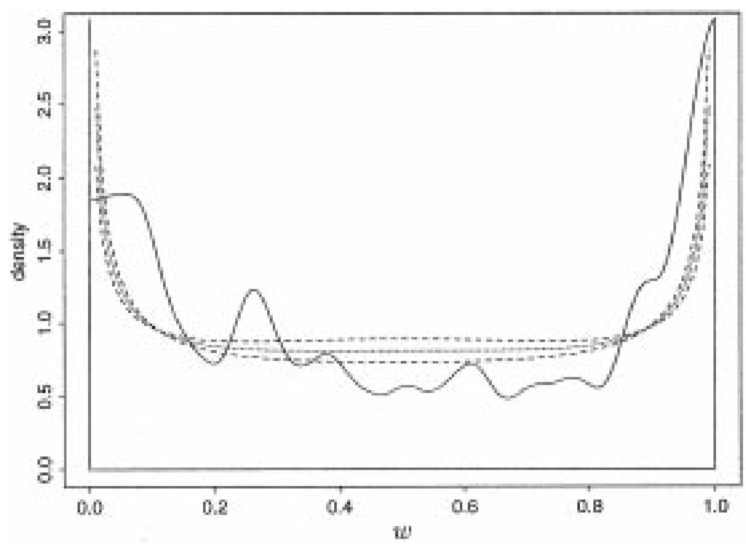

Figure 12. The estimated limiting density $dH(w)/dw$ for the wave and surge level data. The estimated logistic density and pointwise 95% confidence intervals (dotted and dashed curves respectively), and the solid line is the kernel density estimate of the variable $W \mid R > r$, with r the 0.9 quantile of the variable R.

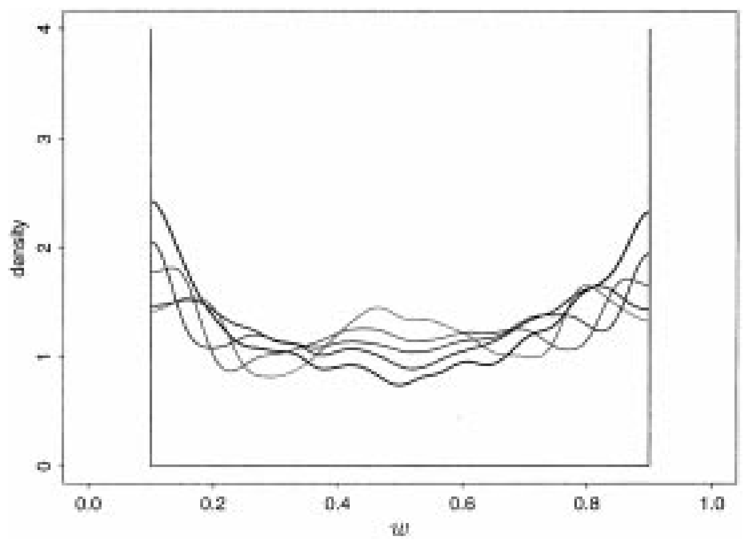

Figure 13. The estimated conditional density of $W \mid (R > r, 0.1 < W < 0.9)$ for the rainfall data. The different estimates correspond to different values of r; the thin line through to the thick line correspond to r being the $0.5, 0.6, \ldots, 0.9$ quantile of the variable $R \mid (0.1 < W < 0.9)$ respectively.

only for estimation of joint tail probabilities where the failure region A is contained in a set of the form B_2, given by equation (4.11). We illustrate extrapolation only for the wave-surge data as the rainfall data are effectively independent at high levels. In practice some large combinations of waves and surges produce worse coastal flooding than others, in a way that depends on the characteristics of the flood defence design. The probability of such events depends heavily on the marginal distributions of the process, which makes the general issue of extrapolation too complex for a simple illustrative example; see Coles and Tawn (1994) and de Haan and de Ronde (1998) for worked examples. Instead, we presume that the wave-surge data have already been transformed to variables (X,Y) with unit Fréchet marginal distributions, and that the region of interest is given by

$$E_v = \{(x,y) \in \mathbb{R}_+^2 \ : \ x > v \text{ and } y > v\}.$$

Figure 14 illustrates such a region in the case $v = 1000$. Here, $E_v \not\subset B_2$, where B_2 is given by (4.11), however, for all v, E_v can be approximated by a set $E_v^* \subset B_2$ with $\Pr(E_v^*) \approx \Pr(E_v)$. Then, from approximation (4.12), we have that $\Pr(E_{tu}) \approx t^{-1/\eta} \Pr(E_u)$ for suitably large u and $t \geq 1$. Defining $p = \Pr(E_{v_p})$ leads to

$$\log v_p = -\eta \log p + \log u + \eta \log \Pr(E_u). \tag{4.13}$$

For two choices of u, each shown in Figure 14, we estimate $\Pr(E_u)$ empirically. Replacing $\Pr(E_u)$ and η in expression (4.13) by their estimates gives the estimates for $\log v_p$ that are shown in Figure 15. Pointwise 95% confidence intervals show that the effect of threshold choice on the extrapolation is very slight. The greater impact of η is illustrated by showing estimates based on both $\eta = 0.5$ and $\eta = 1$. The fact that the confidence intervals contain the $\eta = 1$ values is a consequence of these data being consistent with asymptotic dependence.

5. Discussion

Our aim in this paper has been to give an overview of the issues concerning the use of multivariate extreme value techniques and to synthesize informal exploratory, and formal modeling, procedures. We have developed the argument in the simplest case of bivariate, independent and identically distributed variables, but the issues extend to more complex modeling situations, for which the techniques presented here, appropriately generalized, still apply. The distinction between asymptotic dependence and asymptotic independence has been found to be crucial for both model development and data application. Thus, the limitation of the classical models for multivariate extremes to the asymptotically dependent case, for which $\bar{\chi} = 2\eta - 1 = 1$, is seen to be a strong restriction. That this issue is of considerable practical importance is supported by the empirical study of Bruun and Tawn (1998).

We have said little about marginal estimation, which we presume to have been

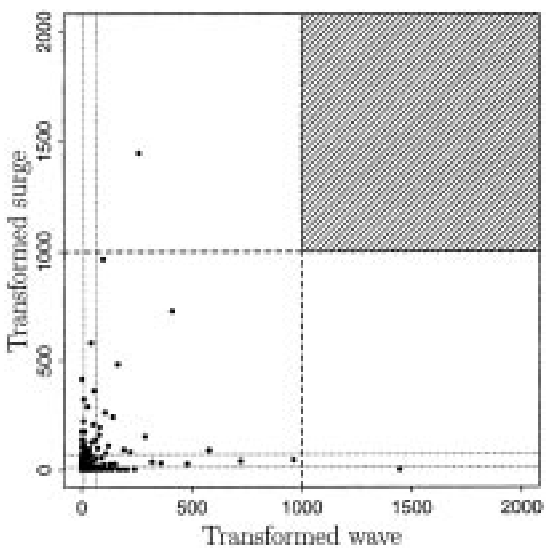

Figure 14. Plot of the wave and surge level data transformed to have unit Fréchet marginal distributions. The plot shows the shaded region E_v with $v = 1000$, and the dotted lines indicate two choices of joint threshold levels: $u = 10$ and $u = 70$.

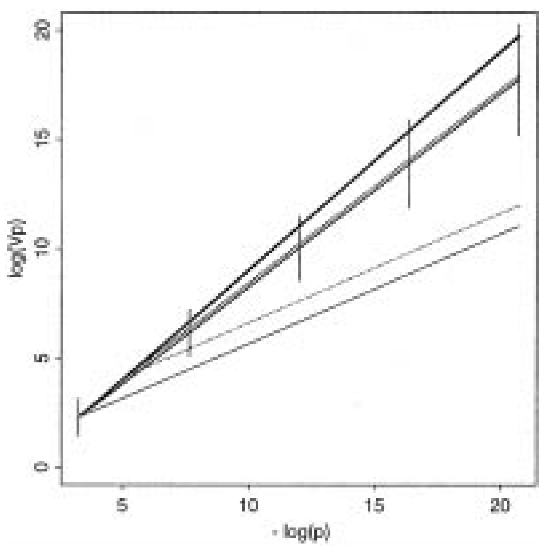

Figure 15. Plot showing the effect of $\bar{\chi}$ on extrapolations of the joint tail. Here $\log v_p$ is plotted against $-\log p$ where $\bar{F}(v_p, v_p) = p$ and $\bar{F}$ is the joint survivor function for wave and surge level variables transformed to have Fréchet marginals. The solid and dotted lines correspond to the lower and higher threshold levels (shown in Figure 14) respectively. The bottom, middle and top lines for each line type correspond in order to $\eta = 0.5, \hat{\eta}$ and 1. The vertical lines indicate pointwise 95% confidence intervals for estimates of $\log v_p$ for the lower joint threshold choice only.

undertaken prior to the assessment of dependence. In practice, if marginal and dependence components are estimated simultaneously, there may be a transfer of information between model components. The effect on the estimation of dependence parameters is often slight due to the near-orthogonality of marginal and dependence parameters. However, there may be substantial improvements in marginal parameter estimation as information also transfers between margins; see Shi et al. (1992), Einmahl et al. (1997) and Barão and Tawn (1999).

Additional complications also arise in practice. For example, multivariate data, such as the lagged rainfall pairs, are often time-dependent. In such circumstances, and notwithstanding the extra inferential difficulties, the various measures of extremal dependence introduced can play an important role in modeling (Ledford and Tawn, 2000). For example, provided the multivariate data have only weak long-range dependence and that they are asymptotically independent at all lags, then the extremal behavior of the series is asymptotically equivalent to that of an independent multivariate series. Therefore, testing for asymptotic independence in this context is similar to testing for clustering conditions similar to the Leadbetter et al. (1983) $D'(u_n)$ condition.

In situations where no asymptotic clustering is identified, it is the sub-asymptotic temporal behavior of the series that is important. To study this, the measure $\bar{\chi}$ between lagged variables becomes relevant; see Bortot and Tawn (1998) for the study of 1-dimensional asymptotically independent Markov chains. For processes with asymptotic clustering the level of dependence is appropriately measured by a (multivariate) extremal index; see Leadbetter et al. (1983) and Nandagopalan (1994). In this case, for 1-dimensional Markov chains, results in Smith et al. (1997) and Yun (1998) suggest that χ is highly influential in determining the extremal index value.

Hence, in summary, the two diagnostic measures χ and $\bar{\chi}$ provide informative and complementary information about the form of extremal dependence in multivariate series. The parameter χ has a very simple relationship with multivariate extreme value models, but being model-free, provides a robust assessment of extremal dependence. Its limitations are apparent for asymptotically independent data, for which $\bar{\chi}$ provides a more appropriate, also model-free, measure of extremal dependence.

Acknowledgment

We are grateful to EPSRC for funding the work of all three authors and to the editor for encouraging us to develop what was originally a chapter from a short course on extreme values, given at the RSS Conference September 1998 and the Newton Institute in October 1998, into this paper.

References

Barão, M.I. and Tawn, J.A., ''Extremal analysis of short series with outliers: sea-levels and athletics records,'' *Appl. Statist.* 48, 469–487, (1999).

Bortot, P. and Tawn, J.A., ''Models for the extremes of Markov chains,'' *Biometrika* 85, 851–867, (1998).

Bortot, P., Coles, S.G., and Tawn, J.A., ''The multivariate Gaussian tail model: an application to oceanographic data,'' *Appl. Statist.*, 49, 31–49 (2000).

Bruun, J.T. and Tawn, J.A., ''Comparison of approaches for estimating the probability of coastal flooding,'' *Appl. Statist.* 47, 405–423, (1998).

Capéraà, P., Fougères, A.-L., and Genest, C., ''Estimation of bivariate extreme value copulas,'' *Biometrika* 84, 567–577, (1997).

Coles, S.G., ''A temporal study of extreme rainfall,'' in *Statistics for the Environment 2: Water Related Issues*, (V. Barnett and K.F. Turkman, eds.), Wiley, Chichester, 61–78, (1994).

Coles, S.G. and Tawn, J.A., ''Modelling extreme multivariate events,'' *J. R. Statist. Soc. B* 53, 377–392, (1991).

Coles, S.G. and Tawn, J.A., ''Statistical methods for multivariate extremes: an application to structural design (with discussion),'' *Appl. Statist.* 43, 1–48, (1994).

Coles, S.G. and Tawn, J.A., ''A Bayesian analysis of extreme rainfall data,'' *Appl. Statist.* 45, 463–478, (1996).

Currie, J.E., ''Directory of coefficients of tail dependence,'' *Department of Mathematics and Statistics Technical Report*, ST-99-06, Lancaster University, 1999.

Davison, A.C. and Smith, R.L., ''Models for exceedances over high thresholds (with discussion),'' *J. R. Statist. Soc. B* 52, 393–442, (1990).

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

Einmahl, J.H.J., de Haan, L. and Sinha, A.K., ''Estimating the spectral measure of an extreme value distribution,'' *Stoch. Proc. Appl.* 70, 143–171, (1997).

Geffroy, J., ''Contributions à la théorie des valeurs extrème,'' *Publ. Inst. Statist. Univ. Paris* 7/8, 37–185, (1958/59).

de Haan, L., ''Extremes in higher dimensions: the model and some statistics.'' In *Proc. 45th Sess. Int. Statist. Inst.*, paper 26.3. International Statistical Institute, The Hague, 1985.

de Haan, L. and de Ronde, J., ''Sea and wind: multivariate extremes at work,'' *Extremes* 1, 7–45, (1998).

Joe, H., ''Families of min-stable multivariate exponential and multivariate extreme value distributions,'' *Statist. Probab. Lett.* 9, 75–81, (1990).

Joe, H., ''Multivariate extreme-value distributions with applications to environmental data,'' *Canadian. J. Statist.* 22, 47–64, (1994).

Joe, H., *Multivariate Models and Dependence Concepts*, Chapman & Hall, London, 1997.

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

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

Ledford, A.W., ''Extreme values of the bivariate normal distribution,'' Submitted, 1999.

Ledford, A.W. and Tawn, J.A., ''Statistics for near independence in multivariate extreme values,'' *Biometrika* 83, 169–187, (1996).

Ledford, A.W. and Tawn, J.A., ''Modelling dependence within joint tail regions,'' *J. R. Statist. Soc. B* 59, 475–499, (1997).

Ledford, A.W. and Tawn, J.A., ''Concomitant tail behavior for extremes,'' *Adv. Appl. Probab.* 30, 197–215, (1998).

Ledford, A.W. and Tawn, J.A., ''Diagnostics for dependence within time-series extremes,'' Submitted to *J. R. Statist. Soc. B* (2000).

Mardia, K.V., ''Asymptotic independence of bivariate extremes,'' *Calcutta Statist. Assoc. Bull.* 13, 172–178, (1964).

Nandagopalan, S., ''On the multivariate extremal index,'' *Jnl of Research, National Inst. of Standards and Technology* 99, 543–550, (1994).

Nelsen, R.B., *An Introduction to Copulas*, Springer-Verlag, New York, 1998.

Peng, L., ''Estimation of the coefficient of tail dependence in bivariate extremes,'' *Statist. Probab. Lett.* 43, 399–409, (1999).

Pickands, J., ''Multivariate extreme value distributions,'' *Bull. Int. Statist. Inst.* 49, 859–878, (1981).

Resnick, S.I., *Extreme Values, Point Processes and Regular Variation*, Springer-Verlag, New York, 1987.

Sibuya, M., ''Bivariate extreme statistics,'' *Ann. Inst. Statist. Math.* 11, 195–210, (1960).

Shi, D., Smith, R.L., and Coles, S.G., ''Joint versus marginal estimation for bivariate extremes,'' *Department of Statistics Technical Report*, 2074, University of North Carolina at Chapel Hill, 1992.

Smith, R.L., ''Extreme value theory,'' In *Handbook of Applicable Mathematics*, (W. Ledermann, ed), 7, chapter 14, 437–472, Chichester, Wiley, 1990.

Smith, R.L., ''Multivariate threshold methods.'' In *Extreme Value Theory & Applications*, (J. Galambos, J. Lechner and E. Simiu, eds) Kluwer, Dordrecht, 225–248, (1994).

Smith, R.L., Tawn, J.A., and Coles, S.G., ''Markov chain models for thresholds exceedances,'' *Biometrika* 84, 249–268, (1997).

Tawn, J.A., ''Bivariate extreme value theory: models and estimation,'' *Biometrika* 75, 397–415, (1988).

Tiago de Oliveira, J., ''Structure theory of bivariate extremes, extensions,'' *Est. Mat., Estat. e. Econ.* 7, 165–195, (1962/63).

Weintraub, K.S., ''Sample and ergodic properties of some min-stable processes,'' *Annals Probab.* 19, 706–723, (1991).

Yun, S., ''The extremal index of a higher-order stationary Markov chain,'' *Ann. Appl. Probab.* 8, 408–437, (1998).