

Bayesian Uncertainty Management in Temporal Dependence of Extremes

T. Lugrin, A. C. Davison and J. A. Tawn*

March 4, 2016

Abstract

Both marginal and dependence features must be described when modelling the extremes of a stationary time series. There are standard approaches to marginal modelling, but long- and short-range dependence of extremes may both appear. In applications, an assumption of long-range independence often seems reasonable, but short-range dependence, i.e., the clustering of extremes, needs attention. The extremal index $0 < \theta \leq 1$ is a natural limiting measure of clustering, but for wide classes of dependent processes, including all stationary Gaussian processes, it cannot distinguish dependent processes from independent processes with $\theta = 1$. Eastoe and Tawn (2012) exploit methods from multivariate extremes to treat the subasymptotic extremal dependence structure of stationary time series, covering both $0 < \theta < 1$ and $\theta = 1$, through the introduction of a threshold-based extremal index. Inference for their dependence models uses an inefficient stepwise procedure that has various weaknesses and has no reliable assessment of uncertainty. We overcome these issues using a Bayesian semiparametric approach. Simulations and the analysis of a UK daily river flow time series show that the new approach provides improved efficiency for estimating properties of functionals of clusters.

Keywords: Asymptotic independence; Bayesian semiparametrics; Conditional extremes; Dirichlet process; Extreme value theory; Extremogram; Risk analysis; Threshold-based extremal index

*Thomas Lugrin is a PhD candidate and Anthony Davison is Professor of Statistics, EPFL-FSB-MATHAA-STAT, Station 8, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland, Thomas.Lugrin@epfl.ch, Anthony.Davison@epfl.ch. Jonathan Tawn is Professor of Statistics, Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, UK, j.tawn@lancs.ac.uk. This research was partially supported by the Swiss National Science Foundation.

1 Introduction

Extreme value theory provides an asymptotically justified framework for the statistical modelling of rare events. In the univariate case with independent variables there is a broadly-used framework involving modelling exceedances of a high threshold by a generalised Pareto distribution (Coles, 2001). For extremes of stationary univariate time series, standard procedures use marginal extreme value modelling but consideration of the dependence structure between the variables is essential when assessing risk due to clusters of extremes. Leadbetter *et al.* (1983) and Hsing *et al.* (1988) described the roles of long- and short-range dependence on extremes of stationary time series. Typically an assumption of independence at long range is reasonable, with Ledford and Tawn (2003) giving diagnostic methods for testing this. In practice independent clusters are often identified using the runs method (Smith and Weissman, 1994), which deems successive exceedances to be in separate clusters if they are separated by at least m consecutive non-exceedances of the threshold; Ferro and Segers (2003) provide automatic methods for the selection of m .

Short-range dependence has the most important practical implications, since it leads to local temporal clustering of extreme values. Leadbetter (1983) and O'Brien (1987) provide different asymptotic characterisations of the clustering through the extremal index $0 < \theta \leq 1$, the former with θ^{-1} being the limiting mean cluster size. The case $\theta = 1$ corresponds to there being no more clustering than if the series were independent, and decreasing θ corresponds to increased clustering.

Although the extremal index is a natural limiting measure for such clustering there is a broad class of dependent processes with $\theta = 1$, including all stationary Gaussian processes. Thus the extremal index cannot distinguish the clustering properties of this class of dependent processes from those of white noise. Furthermore, many other functionals of clusters of extremes may be of interest (Segers, 2003), so for practical modelling of the clusters of extreme values above a high threshold, the restriction to $\theta < 1$ is a major weakness.

Smith *et al.* (1997b), Ledford and Tawn (2003), Eastoe and Tawn (2012) and Winter and Tawn (2016) draw on methodology for multivariate extremes to provide models for univariate clustering, and thereby enable the properties of a wide range of cluster functionals to be estimated. The focus of this paper is the improvement of inference techniques for the most general model for these cases, namely the semiparametric conditional extremes model of Heffernan and Tawn (2004). Our ideas can be applied to any cluster functional, but we focus here primarily on the threshold-based extremal index introduced by Ledford and Tawn (2003).

Given the strong connections between multivariate extreme value and clustering modelling, here and in Section 3 we present the developments of the model in parallel for the two situations. Examples of applications for multivariate cases include assessing the risk of joint occurrence of extreme river flows or sea-levels at different locations (Keef

et al., 2009; Asadi *et al.*, 2015), the concurrent high levels of different pollutants at the same location (Heffernan and Tawn, 2004), and simultaneous stock market crashes (Poon *et al.*, 2003). For the time series case, applications include assessing heatwave risks (Reich *et al.*, 2014; Winter and Tawn, 2016), modelling of extreme rainfall events (Süveges and Davison, 2012) and wind gusts (Fawcett and Walshaw, 2006).

For the stationary time series (X_t) , Ledford and Tawn (2003) define the threshold-based extremal index

$$\theta(x, m) = \Pr(X_1 \leq x, \dots, X_m \leq x \mid X_0 > x), \quad (1.1)$$

where x is large, which is the key measure of short-range clustering of extreme values, with $1/\theta(x, m)$ being the mean cluster size when clusters of exceedances of the threshold x are defined via the runs method of Smith and Weissman (1994) with run length m . Furthermore $\theta(x, m)$ converges to the extremal index θ as $x \rightarrow x_F$ and $m \rightarrow \infty$ appropriately (O'Brien, 1987; Kratz and Rootzén, 1997). Many studies have focused on estimating the limit θ (Ferro and Segers, 2003; Süveges, 2007; Robert, 2013), but in applications, all these consider a finite level u as an approximation to the limit x_F . This is equivalent to assuming $\theta(x, m)$ to be constant above u , which is generally not the case in applications (see Figure 1). Additionally Eastoe and Tawn (2012) find that $\theta(x, m)$ is fundamental to modelling the distributions of both cluster maxima, i.e., peaks over threshold, and block maxima, e.g., annual maxima. See Section 4 for more on the relevance of $\theta(x, m)$ for time series extremes.

When considering asymptotically motivated models for the joint distribution of $\mathbf{X}_{-0} = (X_1, \dots, X_m)$ given that $X_0 > x$ for the estimation of $\theta(x, m)$, it is helpful to have a simple characterisation of extremal dependence. The standard pairwise measure of extremal dependence for (X_0, X_j) is

$$\chi_j = \lim_{x \rightarrow x_F} \Pr(X_j > x \mid X_0 > x), \quad j = 1, \dots, m, \quad (1.2)$$

with the cases $\chi_j > 0$ and $\chi_j = 0$ respectively termed asymptotic dependence and asymptotic independence at lag j . A plot of χ_j against j has been termed the extremogram (Davis and Mikosh, 2009), by analogy with the correlogram of time series analysis. When $\chi_j = 0$ for all $j \geq 1$, the extremogram fails to distinguish between different levels of asymptotic independence, but the rate of convergence to zero of $\Pr(X_j > x \mid X_0 > x)$ determines the key characteristics of the tail of the joint distribution (Ledford and Tawn, 1996). Ledford and Tawn (2003) propose using such a measure at each time lag j when the variables are asymptotically independent. An alternative is to combine both approaches by studying a threshold-based version of χ_j ,

$$\chi_j(x) = \Pr(X_j > x \mid X_0 > x), \quad j = 1, \dots, m, \quad (1.3)$$

for a range of large values of x .

In Section 2, we review classical multivariate extreme models, which all entail $\chi_j(x) = \chi_j$, $x > u$ for some high threshold u , and often even $\chi_j > 0$. Instead we consider the conditional formulation of Heffernan and Tawn (2004) that has been subsequently studied more theoretically by Heffernan and Resnick (2007), Das and Resnick (2011), and Mitra and Resnick (2013). This class of models covers $\chi_j \geq 0$ and $\chi_j(x)$ changing with large x ($j = 1, \dots, m$) through modelling dependence within the asymptotic independence class. This model gives estimates of $\theta(x, m)$ that can be constant or vary with x , $x > u$. This additional flexibility comes at a price: inference is required for up to $2m$ parameters, and for an arbitrary m -dimensional distribution G .

The asymptotic arguments for the Heffernan–Tawn model are given in Section 3, for an $(m + 1)$ -dimensional variable with Laplace marginal distributions. Suppose that the monotone increasing transformation $T = (T_0, \dots, T_m)$ transforms $\mathbf{X} = (X_0, \dots, X_m)$ to $\mathbf{Y} = (Y_0, \dots, Y_m)$, with $Y_i = T_i(X_i)$ ($i = 0, \dots, m$), so that \mathbf{Y} has Laplace marginal distributions. In applications the Heffernan–Tawn model corresponds to a multivariate regression with

$$\mathbf{Y}_{-0} \mid \{Y_0 = y\} = \boldsymbol{\alpha}y + y^\beta \mathbf{Z} = \boldsymbol{\alpha}y + \boldsymbol{\mu}y^\beta + \boldsymbol{\psi}y^\beta \mathbf{Z}^* \quad (1.4)$$

where here, and subsequently, the arithmetic is to be understood componentwise, with $\boldsymbol{\alpha} \in [-1, 1]^m$, $\boldsymbol{\beta} \in [-\infty, 1]^m$, $\boldsymbol{\mu} \in \mathbb{R}^m$, $\boldsymbol{\psi} \in \mathbb{R}_+^m$, and \mathbf{Z} an m -dimensional random variable with $\mathbf{Z} \sim G$; \mathbf{Z}^* has zero mean and unit variance for all marginal variables, with $\mathbf{Z}^* = (\mathbf{Z} - \boldsymbol{\mu})/\boldsymbol{\psi}$. We require that (1.4) holds for all $y > u$, where u is a high threshold on the Laplace marginal scale. The parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\psi})$ determine the conditional mean and variance through

$$\mathbb{E}(\mathbf{Y}_{-0} \mid Y_0 = y) = \boldsymbol{\alpha}y + \boldsymbol{\mu}y^\beta, \quad \text{var}(\mathbf{Y}_{-0} \mid Y_0 = y) = \boldsymbol{\psi}^2 y^{2\beta}.$$

Thus $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\psi})$ can be estimated by multivariate regression. The complication for inference is that the error distribution G is in general unknown and arbitrary, apart from the first two moment properties mentioned above. One exception to this is when $\boldsymbol{\alpha} = \mathbf{1}$ and $\boldsymbol{\beta} = \mathbf{0}$, in which case the Heffernan–Tawn model reduces to known asymptotically-dependent models with G directly related to H , as detailed in Section 3.

Heffernan and Tawn (2004) and Eastoe and Tawn (2012) used a stepwise inference procedure, estimating $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\psi})$ under a working assumption that \mathbf{Z}^* are independent normal variables. After obtaining parameter estimates, they estimated G nonparametrically using the empirical joint distribution of the observed standardised multivariate residuals, i.e., values of \mathbf{Z} for $Y_0 > u$. There are weaknesses in this approach, which loses efficiency in the estimation of $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}, \boldsymbol{\psi})$ and G and in subsequent inferences due to the generally incorrect working assumption of normality. Moreover, as noted by Peng and Qi (2004), the empirical estimation of G leads to poor estimation of the upper tail of the conditional distribution of $\mathbf{Y}_{-0} \mid \{Y_0 = y\}$, so it would be preferable to have a better, yet general, estimator of G . Furthermore, the uncertainty of the parameter estimation is unaccounted-for in the estimation of G and of cluster functionals such as $\theta(x, m)$.

Cheng *et al.* (2014) proposed a Bayesian approach to estimating the Heffernan–Tawn model in a single stage, but their estimation procedure involves changing the structure of the model and adding a noise term in (1.4), thereby allowing the likelihood term to be split appropriately. They also need strong prior information extracted from the stepwise inference method in order to get valid estimates for the model parameters, so this procedure does not really tackle the loss of efficiency of the stepwise estimation procedure.

We propose to overcome these weaknesses by using Bayesian semiparametric inference to estimate the model parameters and the distribution G , simultaneously performing the entire fitting procedure for the dependence model. This gives a new model for G , namely, a mixture of Gaussian distributions, which provides estimates of the conditional distribution of $\mathbf{Y}_{-0} \mid Y_0 = y$ beyond the range of the current estimator and which in theory provides an arbitrary good approximation for G (Marron and Wand, 1992). The Bayesian approach also provides a coherent framework for fitting a parsimonious parametric model; joint estimation of the model parameters enables the imposition of structure between them. For example, in multivariate problems the context may suggest that different components of $\boldsymbol{\alpha}$ may be identical. In the context of time series extremes, for first-order Markov models, it can be shown that $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ involve at most two unknown parameters (Papastathopoulos *et al.*, 2016). Furthermore, when the \mathbf{X}_{-0} are known to be asymptotically dependent on X_0 , this method provides a new approach to modelling.

We show the practical importance of the new approach by applying it to the daily mean flow time series of the River Ray at Grendon Underwood in north-west London, with observations from October 1962 to December 2008. We consider only flows from October to March, as this period typically contains the largest flows and forms an approximately stationary series. An empirical estimate of $\theta(x, m)$, with $m = 4$, is shown in Figure 1 with bootstrap-derived 95% confidence intervals. A major weakness with this estimate is that it cannot be evaluated beyond the range of the data, so a model is needed to evaluate $\theta(x, m)$ for larger x . We select our modelling threshold u to be the empirical 98% marginal quantile of the data. Using the methods in Eastoe and Tawn (2012) we have an estimate of $\theta(x, 4)$ for all $x > u$ using the stepwise estimation method. As seen in Figure 1, this estimate converges to 1 as $x \rightarrow x_F$ but we have no reliable method for deriving confidence intervals. Figure 1 also shows posterior median estimates and 95% credibility intervals obtained using our Bayesian semiparametric method. These show broad agreement with both of the other estimates within the range of the data, but with tighter uncertainty intervals and statistically significant differences in extrapolation of $\theta(x, 4)$, indicating that the new method has the potential to offer marked improvement for estimating $\theta(x, m)$ and other cluster functionals.

The paper is structured as follows. We first briefly present the standard approaches to multivariate extremes in Section 2. We introduce the conditional model of interest in a multivariate framework in Section 3, followed by a section about modelling of dependent time series. Section 5 explains the Bayesian semiparametric inference procedure, which




Figure 1 – Comparison of the empirical, stepwise and Bayesian semiparametric estimates of $\theta(x, 4)$ described in Sections 3 and 5 respectively. Empirical estimate (solid black) and stepwise (long dashed) and Bayesian (solid grey) semiparametric estimates with their respective 95% confidence regions (dashed and shaded), estimated on winter flows (m^3s^{-1}) of the River Ray from 1962 to 2008.

is used in Section 6 to illustrate the efficiency gains of this new inference method on simulated data. In Section 7 we fit our model to the River Ray flow data and show its ability to estimate functionals of time series clusters other than the threshold-based index.

2 Multivariate setup and classical models

Both multivariate and time series extremes involve estimating the probability of events that may never yet have been observed. Suppose that $\mathbf{X} = (X_0, \dots, X_m)$ is an $(m + 1)$ -dimensional variable with joint distribution function $F_{\mathbf{X}}$ and marginal distribution functions F_0, \dots, F_m . We need to estimate the probability $\Pr(\mathbf{X} \in A)$, where $A \subset \mathbb{R}^{m+1}$ is an extreme set, i.e., a set such that for all $\mathbf{x} \in A$, at least one component of $\mathbf{x} = (x_0, \dots, x_m)$ is extreme. To do this we must model $F_{\mathbf{X}}(\mathbf{x})$ for all $\mathbf{x} \in B$, where B is an extreme set that contains A . Let A_i be the subset of A for which component i is largest on a quantile scale, i.e.,

$$A_i = A \cap \{\mathbf{x} \in \mathbb{R}^{m+1} : F_i(x_i) > F_j(x_j), j \in \{0, \dots, m\} \setminus \{i\}\}, \quad i = 0, \dots, m,$$

and let $v_i = \inf\{x_i : \mathbf{x} \in A_i\}$, so that we can write

$$\Pr(\mathbf{X} \in A) = \sum_{i=0}^m \Pr(\mathbf{X} \in A_i \mid X_i > v_i) \Pr(X_i > v_i). \quad (2.1)$$

Thus estimates of marginal and dependence features are required for estimating the conditional probabilities in the sum and marginal distributions determine the second terms of the products in the sum.

Although our approach applies to any form of set A , to focus our arguments we restrict ourselves to identical marginal distributions F with upper endpoint x_F and we set

$$A = A_0 = \{X_0 > x, X_1 \leq x, \dots, X_m \leq x\},$$

so $v_0 = x$ and we estimate only the conditional probability term in (2.1), i.e., $\theta(x, m)$, as defined in (1.1).

Early approaches to modelling the conditional distribution appearing in (1.1) assumed that \mathbf{X} lies in the domain of attraction of a multivariate extreme value distribution (Coles and Tawn, 1994; de Haan and de Ronde, 1998) and applied these asymptotic models above a high threshold. Unlike in the univariate case, there is no finite parametrisation of the dependence structure; it can only be restricted to functions of a distribution H on the m unit simplex S_m with $\int_{S_m} w_i dH(\mathbf{w}) = (m+1)^{-1}$ ($i = 0, \dots, m$), where $\mathbf{w} = (w_0, \dots, w_m)$. Both parametric and non-parametric inference for this class of models has been proposed. Numerous parametric models are available (Kotz and Nadarajah, 2000, Ch. 3; Cooley *et al.*, 2010; Ballani and Schlather, 2011). Nonparametric estimation is also widely studied, mostly based on empirical estimators (de Haan and de Ronde, 1998; Hall and Tadjvidi, 2000; Einmahl *et al.*, 2001; Einmahl and Segers, 2009).

A major weakness of these early methods is that for these models either $\chi_j > 0$ or (X_0, X_j) is independent for all $j = 1, \dots, m$, whereas there are distributions, such as the multivariate Gaussian, with $\chi_j = 0$ but (X_0, X_j) dependent. If $\chi_j > 0$ for any $j = 1, \dots, m$, these models give estimates of $\theta(x, m) \rightarrow c_m$ as $x \rightarrow x_F$, where $c_m < 1$. This class of models is not flexible enough to cover distributions that are dependent at finite levels, but asymptotically independent for all pairs of variables.

3 Threshold-based model for conditional probabilities

3.1 Heffernan–Tawn model

In order to provide a model characterising conditional probabilities, such as those which describe the clustering behaviour, we need a model for multivariate extreme values, and in particular we require a model for the joint distribution of X_1, \dots, X_m given that $X_0 > x$.

In this section we suppose that the marginal distributions of \mathbf{X} are not identical, and that $X_0 \sim F_0$ is in the domain of attraction of a generalised Pareto distribution, i.e., there exists a function $\sigma_u > 0$ such that as $u \rightarrow x_{F_0}$, $(X_0 - u)/\sigma_u$, conditional on $X_0 > u$, converges to a generalised Pareto variable with unit scale parameter and shape parameter ξ .

The joint distribution is modelled via the marginal distributions F_j and a copula for the dependence structure. To study the conditional behaviour of extremes, the copula formulation is most transparent when expressed with Laplace marginals. Let T_j denote the transformation of the marginal distribution of X_j to the Laplace scale, i.e.,

$$T_j(X_j) = \begin{cases} \log \{2F_j(X_j)\}, & X_j < F_j^{-1}(1/2), \\ -\log [2\{1 - F_j(X_j)\}], & X_j > F_j^{-1}(1/2); \end{cases} \quad j = 0, \dots, m,$$

the specification of the F_j is discussed in Section 3.2. Assume there exist m -dimensional functions $\mathbf{a}(x) = \{a_1(x), \dots, a_m(x)\}$ and $\mathbf{b}(x) = \{b_1(x), \dots, b_m(x)\} > \mathbf{0}$ for which

$$\Pr \left[\frac{T_j(X_j) - a_j \{T_0(X_0)\}}{b_j \{T_0(X_0)\}} \leq z_j, j = 1, \dots, m \mid X_0 > u \right] \rightarrow G(\mathbf{z}), \quad u \rightarrow x_{F_0}, \quad (3.1)$$

where all marginal distributions of G are non-degenerate and $\mathbf{z} = (z_1, \dots, z_m) \in \mathbb{R}^m$. Hereafter we write the standardised X_j , or residual, as

$$Z_j = \frac{T_j(X_j) - a_j \{T_0(X_0)\}}{b_j \{T_0(X_0)\}}, \quad X_0 > u, \quad j = 1, \dots, m.$$

Under assumption (3.1), the rescaled conditioning variable $(X_0 - u)/\sigma_u$ is asymptotically conditionally independent of the residual $\mathbf{Z} = (Z_1, \dots, Z_m)$ given $X_0 > u$, as $u \rightarrow x_{F_0}$. That is,

$$\begin{aligned} \Pr \{ \mathbf{Z} \leq \mathbf{z}, (X_0 - u)/\sigma_u > x \mid X_0 > u \} \\ &= \Pr \{ \mathbf{Z} \leq \mathbf{z} \mid (X_0 - u)/\sigma_u > x \} \Pr \{ (X_0 - u)/\sigma_u > x \mid X_0 > u \} \\ &\rightarrow G(\mathbf{z}) \bar{K}(x), \quad u \rightarrow x_{F_0}, \end{aligned} \quad (3.2)$$

where \bar{K} is the generalised Pareto distribution survivor function (4.1) with scale and shape parameters $(1, \xi)$ and G is the limit distribution of the residuals.

Equation (3.1) can be illustrated through the particular case when \mathbf{X} is a centred multivariate Gaussian distribution with correlation matrix elements ρ_{ij} , $i, j = 0, \dots, m$, $i \neq j$. In this case we can derive $a_j(x) = \text{sign}(\rho_{0j})\rho_{0j}^2 x$ and $b_j(x) = x^{1/2}$, and $G(\mathbf{z})$ is a centred multivariate Gaussian distribution with variances $\rho_{0j}^2(1 - \rho_{0j}^2)$ and correlation matrix elements

$$\rho'_{ij} = \frac{\rho_{ij} - \rho_{0i}\rho_{0j}}{\sqrt{(1 - \rho_{0i}^2)(1 - \rho_{0j}^2)}}, \quad i \neq j.$$

Heffernan and Tawn (2004) and Keef *et al.* (2013) showed that under broad conditions, the component functions of $\mathbf{a}(x)$ and $\mathbf{b}(x)$ can be modelled by

$$a_j(x) = \alpha_j x, \quad b_j(x) = x^{\beta_j}, \quad -1 \leq \alpha_j \leq 1, \quad -\infty < \beta_j \leq 1, \quad j = 1, \dots, m.$$

In terms of the dependence structure, α_j and β_j reflect the flexibility of the model. It turns out that (X_0, X_j) are asymptotically dependent only if $\alpha_j = 1$, $\beta_j = 0$, and then

$$\chi_j = \lim_{x \rightarrow x_{F_0}} \Pr \{T_j(X_j) > x \mid T_0(X_0) > x\} = \int_0^\infty \bar{G}_j(-z) e^{-z} dz;$$

with \bar{G}_j the j th marginal survivor function of G ; if $-1 < \alpha_j < 1$, then (X_0, X_j) are asymptotically independent, with positive extremal dependence if $\alpha_j > 0$, negative extremal dependence if $\alpha_j < 0$, and with extremal near-independence if $\alpha_j = 0$ and $\beta_j = 0$.

We set $\beta_j \geq 0$, as when $\beta_j < 0$ all the conditional quantiles for X_j converge to the same value as X_0 increases, which is unlikely in most environmental contexts. If the conditioning threshold u is high enough that the conditional probability on the left of (3.1) is close to its limit, then the Heffernan–Tawn model can be stated as

$$T_j(X_j) = \alpha_j T_0(x) + \{T_0(x)\}^{\beta_j} Z_j, \quad X_0 = x > u, \quad j = 1, \dots, m, \quad (3.3)$$

where $(Z_1, \dots, Z_m) \sim G$ is independent of X_0 , and G can be any distribution with non-degenerate margins.

3.2 Existing inference procedure

We now outline the approach to inference suggested by Heffernan and Tawn (2004). Consider a vector (X_0, \dots, X_m) whose marginal distributions F_0, \dots, F_m each lie in the domain of attraction of a generalised Pareto distribution. We estimate them using the semiparametric estimator of Coles and Tawn (1994),

$$\hat{F}_j(x) = \begin{cases} \tilde{F}_j(x), & x < u, \\ 1 - \left\{1 - \tilde{F}_j(u)\right\} \left(1 + \hat{\xi}_j \frac{x - u}{\hat{\sigma}_{u,j}}\right)_+^{-1/\hat{\xi}_j}, & x \geq u, \end{cases}$$

where \tilde{F}_j is the empirical marginal distribution function of X_j . Here $\hat{\sigma}_{u,j}$ and $\hat{\xi}_j$ are maximum likelihood estimates based on all exceedances of u , ignoring any dependence; their variances can be evaluated by a sandwich procedure (Fawcett and Walshaw, 2007) or by a block bootstrap. The margins F_j are transformed to the Laplace scale through the transformation $\hat{T}_j(X_j)$.

Estimation of the probability of any extreme set of interest involves inference for model (3.3) with the estimators of parameters of the dependence model assumed independent of the parameter estimators of the marginal distribution. This assumption has

been found not to be restrictive in other copula inference contexts (Genest *et al.*, 1995). Heffernan and Tawn (2004) proposed a stepwise inference procedure for estimating the extremal dependence structure, based on the working assumption that the residual variables Z_1, \dots, Z_m are independent and Gaussian with means μ_1, \dots, μ_m and variances $\psi_1^2, \dots, \psi_m^2$. This assumption allows likelihood inference based on the assumed marginal densities,

$$T_j(X_j) \mid \{T_0(X_0) = x\} \sim \mathcal{N}(\alpha_j x + x^{\beta_j} \mu_j, x^{2\beta_j} \psi_j^2), \quad x > u, \quad j = 1, \dots, m.$$

The first step of their procedure consists of a likelihood maximisation performed separately for each j , giving estimates of α_j , β_j and the nuisance parameters μ_j and ψ_j . Additional constraints, arising from results of Keef *et al.* (2013), lead to the likelihood function being zero for certain combinations of parameters (α_j, β_j) . Thus the maximisation is over a subset of $[-1, 1] \times [0, 1]$ for these two parameters. These constraints ensure that the conditional quantiles of $T_j(X_j) \mid T_0(x)$ are ordered in a decreasing sequence for all large x under fitted models corresponding to positive asymptotic dependence, asymptotic independence and negative asymptotic dependence respectively. For details of these constraints see Keef *et al.* (2013), who show that imposing these additional constraints improves inference of the conditional extremes model. Given model (3.3) and the estimates $\hat{\alpha}_j$ and $\hat{\beta}_j$, the second step of the estimation procedure involves multivariate residuals \mathbf{Z} for each data point, using the relation

$$\hat{\mathbf{Z}}_j = \frac{\hat{T}_j(X_j) - \hat{\alpha}_j \hat{T}_0(X_0)}{\{\hat{T}_0(X_0)\}^{\hat{\beta}_j}}, \quad j = 1, \dots, m, \quad X_0 > u,$$

and hence constructing the joint empirical distribution function $\hat{G}(\mathbf{z})$.

An estimator for $\Pr(A \mid X_0 > x)$ for any extreme set A is obtained as follows: sample R independent replicates $X_0^{(1)}, \dots, X_0^{(R)}$ of X_0 conditional on $X_0 > x$ from a generalised Pareto distribution with threshold x ; independently sample $\mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(R)}$ from the joint empirical distribution function \hat{G} ; compute

$$\mathbf{X}_{-0}^{(r)} = \hat{\mathbf{T}}^{-1} \left[\hat{\alpha} \hat{T}_0(X_0^{(r)}) + \{\hat{T}_0(X_0^{(r)})\}^{\hat{\beta}} \mathbf{Z}^{(r)} \right], \quad r = 1, \dots, R,$$

where vector arithmetic is to be understood componentwise and $\hat{\mathbf{T}}^{-1} = (T_1^{-1}, \dots, T_m^{-1})$ is a componentwise back-transformation to the original scale; then the estimator for $\Pr(A \mid X_0 > x)$ is

$$\frac{1}{R} \sum_{r=1}^R I \left\{ (X_0^{(r)}, \mathbf{X}_{-0}^{(r)}) \in A \right\},$$

where I is the indicator function.

In the rest of the paper, we are interested in estimating the conditional probability $\theta(x, m)$, corresponding to $A = \{X_0 > x, X_1 < x, \dots, X_m < x\}$, for which the Heffernan–Tawn model provides a characterisation. Under the assumption that the limit (3.1) approximately holds for some subasymptotic u , Eastoe and Tawn (2012) obtain

$$\theta(x, m) = \int_x^\infty G\{\mathbf{z}(x, y)\}k_x(y) dy, \quad x > u, \quad (3.4)$$

where $k_x(y)$ is the generalised Pareto density for threshold x , with scale parameter $1 + \xi(x - u)$ and shape parameter ξ , and $\mathbf{z}(x, y)$ is an m -dimensional vector with elements

$$z_j(x, y) = \frac{T_j(x) - \alpha_j T_0(y)}{\{T_0(y)\}^{\beta_j}}, \quad j = 1, \dots, m. \quad (3.5)$$

A Monte Carlo approximation to the integral (3.4) gives the estimator

$$\hat{\theta}(x, m) = \frac{1}{R} \sum_{r=1}^R \hat{G}\{\mathbf{z}(x, X_0^{(r)})\},$$

where $\mathbf{z}(x, X_0^{(r)})$ is given by expression (3.5) with $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ replaced by estimates. Monte Carlo variability can be reduced by using the same pseudo-random sequence when generating samples for different values of x . Eastoe and Tawn (2012) use a bootstrap method to get confidence bounds for $\hat{\theta}(x, m)$, but de Carvalho and Ramos (2012) found this to be unreliable.

Four main weaknesses of this inference procedure justify developing a more comprehensive approach. First, the working assumption needed for the likelihood maximisation is that the residuals have independent Gaussian distributions, and it is hard to quantify how this affects inference. Second, ignoring the variability of $\hat{\boldsymbol{\alpha}}$, $\hat{\boldsymbol{\beta}}$ estimated in the first step leads to underestimation of the uncertainty in the estimate for the residual distribution G and hence also for $\theta(x, m)$. Third, the empirical estimation of G restricts estimates of extremal conditional probabilities, as simulated \mathbf{Z} values provide no extrapolation over observed values of \mathbf{Z} . Fourth, the inability to impose natural constraints on $(\alpha_1, \dots, \alpha_m)$ and $(\beta_1, \dots, \beta_m)$ leads to inefficiency.

4 Modelling dependence in time

Consider a stationary time series $\{X_t\}$ satisfying appropriate long-range dependence properties and with marginal distribution F . The threshold-based extremal index $\theta(x, m)$ summarises the key extremal dependence in time series. In the block-maxima context, the distribution of the block maximum $M_n = \max\{X_1, \dots, X_n\}$ at a level x is $\{F(x)\}^{n\theta(x, m)}$ (O'Brien, 1987; Kratz and Rootzén, 1997). The associated independent

series $\{X_t^*\}$, having the same marginal distribution as $\{X_t\}$ but independent observations, has $M_n^* = \max\{X_1^*, \dots, X_n^*\}$ with distribution function $\{F(x)\}^n$. So $\Pr(M_n < x) = \{\Pr(M_n^* < x)\}^{\theta(x,m)}$, with $\theta(x,m)$ accounting for the dependence.

The most popular approach to dealing with short-range dependent in such series is the peaks over threshold (POT) approach formalised by Davison and Smith (1990). This approach consists of selecting a high threshold u , identifying independent clusters of exceedances of u , picking the maximum Y of each cluster, and then fitting to these cluster maxima the generalised Pareto distribution

$$\Pr(Y < x \mid Y > u) = 1 - \left(1 + \xi \frac{x - u}{\sigma_u}\right)_+^{-1/\xi}, \quad x > u. \quad (4.1)$$

The limiting results are used as an approximation for data at subasymptotic levels with limit distribution (4.1) taken as exact above a selected value of u .

Alternatives to the POT approach include modelling the series of all exceedances, for example using a Markov chain (Smith *et al.*, 1997a; Winter and Tawn, 2016), but they depend heavily on the validity of the underlying modelling assumptions and so may be inappropriate.

Eastoe and Tawn (2012) consider the threshold-based extremal index as part of a model for the distribution of cluster maxima. Specifically they show that, for a given high threshold u , the cluster maxima, defined by the runs method with run-length m , have approximate distribution function

$$1 - \frac{\theta(x,m)}{\theta(u,m)} \left(1 + \xi \frac{x - u}{\sigma_u}\right)_+^{-1/\xi}, \quad x > u, \quad (4.2)$$

where the parameters ξ and $\sigma_u > 0$ determine the marginal distribution of the original series, and $x_+ = \max(x, 0)$. Eastoe and Tawn (2012) show how using the information in $\theta(x,m)$ can improve over the POT approach. Distribution (4.2) reduces to the generalised Pareto model asymptotically as $u \rightarrow x_F$, and more generally when $\theta(x,m) = \theta(u,m)$ for all $x > u$. When estimates of $\theta(x,m)$ vary appreciably above u , this equality condition for $\theta(x,m)$ provides a diagnostic for situations where the POT method is inappropriate.

In our approach, when a Markov property can reasonably be assumed for a time series, the α_j and β_j have a structure that we want to exploit. Papastathopoulos *et al.* (2016) and Kulik and Soulier (2015) characterise the form of $a_j(x)$ and $b_j(x)$ under very weak assumptions. If the conditions needed for the Heffernan–Tawn simplification — $a_1(x) = \alpha_1 x$ and $b_1(x) = x^{\beta_1}$ — hold, then for positively associated first order Markov processes, Papastathopoulos *et al.* (2016) show that either $(\alpha_j, \beta_j) = (1, 0)$, or $(0, \beta^j)$ or (α^j, β) for some $\alpha \in [0, 1)$ and $\beta \in [0, 1)$. The first case corresponds to asymptotic dependence at all time lags, and the other two to different forms of decaying dependence under asymptotic independence. If $\{X_t\}$ follows an asymptotically dependent Markov process, then no parameters need be estimated, rather than $2m$. If the process is well-approximated by an asymptotically independent Markov process then either of the last

two cases applies, and the number of parameters in the parametric component of the model reduces from $2m$ to 1 or 2. In the case of a Gaussian AR(1) process (X_t) with standard Gaussian margins

$$X_{t+1} = \rho X_t + \varepsilon_t, \quad \varepsilon_t \stackrel{\text{iid}}{\sim} \mathcal{N}(0, 1 - \rho^2), \quad \rho \in (-1, 1),$$

Heffernan and Tawn (2004) get normalising parameters $\alpha_j = \text{sign}(\rho)\rho^{2j}$ and $\beta_j = 1/2$, $j = 1, \dots, m$; the distribution $G(\mathbf{z})$ is a centred multivariate Gaussian with variances $\rho^{2j}(1 - \rho^{2j})$ and correlation matrix elements

$$\rho'_{ij} = \frac{\text{sign}(\rho^{i+j})\rho^{j-i}\sqrt{1 - \rho^{2i}}}{\sqrt{1 - \rho^{2j}}}, \quad i < j.$$

5 Bayesian Semiparametrics

5.1 Overview

Since the m -dimensional residual distribution G in the Heffernan–Tawn model (3.3) is unknown, the approach described in Section 3.2 uses the joint empirical distribution function, which cannot model the tails of the conditional distribution of X_j in (3.3). Our proposed Bayesian approach instead takes G to be a mixture of a potentially infinite number of multivariate Gaussian distributions through the use of a Dirichlet process. This approach can model any G and capture its tails, and has the major benefit of allowing joint estimation of $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$ and G .

Below we introduce Dirichlet processes and describe an approach to approximate Monte Carlo sampling from them. We then describe Bayesian semiparametric inference and the specification of prior distributions, and discuss implementation issues. Throughout we assume that we have n observations from the distribution of $(X_1, \dots, X_m) \mid X_0 > u$, or equivalently from $\mathbf{Z} = (Z_1, \dots, Z_m)$ if $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ were known.

5.2 Dirichlet process mixtures for the residual distribution

Consider a bivariate problem, $m = 1$; with known α and β . If we are to estimate the distribution function of $X_1 \mid X_0 > u$ this is equivalent to estimating the distribution function G of the univariate random variable $Z_1 \sim G$. If G is estimated nonparametrically, its prior must be a distribution over a space of distributions. In this context, a widely used prior is the Dirichlet process (Ferguson, 1973) mixture. A simple model structure takes G to be a mixture of an unknown number of distributions Q_k having parameters $\boldsymbol{\lambda}_k$, $k = 1, 2, \dots$, so that the Dirichlet process boils down to a distribution on the space of mixture distributions P for $\{\boldsymbol{\lambda}_k\}$. If $\boldsymbol{\lambda}_k \mid P \sim P$, then the distribution of P is the Dirichlet process DP (γP_0) , where P_0 is the centre distribution and $\gamma > 0$ the concentration parameter (Hjort *et al.*, 2010).

The definition of a Dirichlet process states that for any $p = 1, 2, \dots$, and any finite measurable partition $\{B_1, \dots, B_p\}$ of the space of the $\boldsymbol{\lambda}_k$,

$$\{P(B_1), \dots, P(B_p)\} \sim \text{Dirichlet}\{\gamma P_0(B_1), \dots, \gamma P_0(B_p)\}.$$

The interpretation of the Dirichlet process parameters stems from the properties

$$\mathbb{E}\{P(B_i)\} = P_0(B_i), \quad \text{var}\{P(B_i)\} = \frac{P_0(B_i)\{1 - P_0(B_i)\}}{\gamma + 1}, \quad i = 1, \dots, p,$$

so the $\text{DP}(\gamma P_0)$ prior is closer to its mean P_0 and less variable for large values of γ . A constructive characterisation of the Dirichlet process is the stick-breaking representation (Sethuraman, 1994)

$$P(\cdot) = \sum_{k=1}^{\infty} w_k \delta_{\boldsymbol{\lambda}_k}(\cdot), \quad (5.1)$$

where $\delta_{\boldsymbol{\lambda}}$ denotes a distribution concentrated on $\boldsymbol{\lambda}$, and $\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_2, \dots$ are independent, P_0 -distributed, and independent of the random weights $w_k \geq 0$, which satisfy $\sum_{k=1}^{\infty} w_k = 1$. The stick-breaking process takes its name from the computation of the weights: define V_1, V_2, \dots as the breaks independently sampled from a $\text{Beta}(1, \gamma)$ distribution. The weights are then

$$w_1 = V_1, \quad w_k = V_k \prod_{i=1}^{k-1} (1 - V_i), \quad k = 2, 3, \dots \quad (5.2)$$

Ishwaran and Zarepour (2000) use formulation (5.1) to express the Dirichlet process in terms of the random variables w_k and $\boldsymbol{\lambda}_k$. They also introduce index variables c_1, \dots, c_n that describe the components of the mixture in which the observations z_1, \dots, z_n lie, giving a stick-breaking representation in terms of the index variables c_i rather than the random variables $\boldsymbol{\lambda}_k$.

A key step in deriving a posterior distribution is the truncation of the sum in the stick-breaking representation, i.e., replacing the infinite sum in (5.1) by a sum up to N . This is achieved by imposing $V_N = 1$ in (5.2). The accuracy of the stick-breaking approximation improves exponentially fast in terms of the \mathcal{L}_1 error (Ishwaran and James, 2001), as

$$\|M_N - M_{\infty}\|_1 \leq 4 \left[1 - \mathbb{E} \left\{ \left(\sum_{k=1}^{N-1} w_k \right)^n \right\} \right] \approx 4n \exp \left(-\frac{N-1}{\gamma} \right),$$

where M_N is the marginal density

$$\int \prod_{i=1}^n \left\{ \sum_{k=1}^N w_k Q_k(dZ_i | \boldsymbol{\lambda}_k) \right\} \text{DP}(dP_N).$$

For example, the error is smaller than 10^{-29} when truncating the stick-breaking sum representation at $N = 150$ as in our real data analysis (Section 7) for which $\gamma \leq 2$ and $n = 154$.

Taking into account the transformations discussed above, a simple model for G involving the Dirichlet process prior is

$$\begin{aligned} Z \mid c, \Lambda &\stackrel{\text{ind}}{\sim} Q_c = Q(\cdot; \boldsymbol{\lambda}_c), \\ c \mid P_N &\stackrel{\text{iid}}{\sim} P_N = \sum_{k=1}^N w_k \delta_k(\cdot), \\ (\Lambda, \mathbf{w}) &\sim \pi_\Lambda(\cdot) \pi_{\mathbf{w}}(\cdot), \end{aligned} \tag{5.3}$$

where Λ is the matrix with rows $\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_N$, and $\mathbf{w} = (w_1, \dots, w_N)$, with $\sum_{k=1}^N w_k = 1$ for some suitably large N . To lighten the notation we write $Z_1 \mid \boldsymbol{\lambda}_c$ instead of $Z_1 \mid c, \Lambda$ in what follows. Taking the Q_k ($k = 1, \dots, N$) to be normal distributions with means $\mu_{1,k}$ and variances $\psi_{1,k}^2$ leads to Λ being a $2 \times N$ matrix, with rows $\boldsymbol{\lambda}_k = (\mu_{1,k}, \psi_{1,k}^2)$. Model (5.3) is made more flexible by adding a hyperprior for the concentration parameter γ .

5.3 Multivariate semiparametric setting

We now specify the features of our algorithm, finally yielding model (5.5). We must add a further element to (5.3): covariates, i.e., the parametric part of the Heffernan–Tawn model (3.3), to recognise that $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are unknown. This is achieved using a covariate-dependent Dirichlet process, and it can be formulated in terms of the truncated stick-breaking representation as

$$P_{|x}(\cdot) = \sum_{k=1}^N w_k \delta_{\boldsymbol{\lambda}_k(x)}(\cdot), \tag{5.4}$$

so that a single output of the stick-breaking procedure gives rise to a whole family of distributions indexed by x . Our data are the n observations from m -dimensional variables \mathbf{X}_{-0} , given X_0 is large. We assume that, conditional on $T(X_0) > u$, $T(\mathbf{X}_{-0})$ has a mixture of multivariate normal distributions, $\sum_{k=1}^{\infty} w_k \mathcal{N}_m$, where the mean vector $M_k(x)$ and the covariance matrix $\Psi_k(x)$ of the k th normal component depend on the value x of $T(X_0)$. For parsimony, the variance matrix $\Psi_k(x)$ is taken to be diagonal with diagonal elements $\{\Psi_{1,k}(x), \dots, \Psi_{m,k}(x)\}$, as the mixture structure is considered flexible enough to capture the dependence between the elements of \mathbf{X}_{-0} .

As we use the truncated version of the stick-breaking representation (5.1), the conditional distribution for the weights w_i is a generalised Dirichlet distribution (Connor and Mosimann, 1969), written as GDirichlet. This gives the final form of our semiparametric model:

$$\begin{aligned}
T(X_j) \mid \{T(X_0) = x, \alpha_j, \beta_j, M_{j,c}(x), \Psi_{j,c}(x)\} &\stackrel{\text{ind}}{\sim} \mathcal{N}\{M_{j,c}(x), \Psi_{j,c}(x)\}, \quad j = 1, \dots, m, \quad x > u, \\
M_{j,c}(x) &= \alpha_j x + \mu_{j,c} x^{\beta_j}, \quad \Psi_{j,c}(x) = x^{2\beta_j} \psi_{j,c}^2, \\
c \mid \mathbf{w} &\sim \sum_{k=1}^N w_k \delta_k(\cdot),
\end{aligned} \tag{5.5}$$

where the prior for $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\mu}_k, \boldsymbol{\psi}_k)$, with $\boldsymbol{\mu}_k = (\mu_{1,k}, \dots, \mu_{m,k})$ and $\boldsymbol{\psi}_k = (\psi_{1,k}, \dots, \psi_{m,k})$, takes the form

$$\begin{aligned}
\mathbf{w} \mid \gamma &\sim \text{GDirichlet}(1, \gamma, \dots, 1, \gamma), \\
\gamma &\sim \text{Gamma}(\eta_1, \eta_2), \\
\alpha_j &\stackrel{\text{iid}}{\sim} \mathcal{U}(0, 1), \quad \beta_j \stackrel{\text{iid}}{\sim} \mathcal{U}(0, 1), \quad j = 1, \dots, m, \\
\mu_{j,k} &\stackrel{\text{iid}}{\sim} \mathcal{N}(0, \psi_{(\mu)}^2), \quad \psi_{j,k}^2 \stackrel{\text{iid}}{\sim} \text{Inv-Gamma}(\nu_{1,j}, \nu_{2,j}), \quad j = 1, \dots, m, \quad k = 1, \dots, N,
\end{aligned}$$

with positive hyper-parameters $(\eta_1, \eta_2, \psi_{(\mu)}^2, \nu_1, \nu_2)$. The Keef *et al.* (2013) conditions mentioned in Section 3.2 are built into the likelihood terms for α_j and β_j , so the Metropolis–Hastings scheme systematically rejects proposals outside the support of the posterior.

5.4 Implementation issues

The semiparametric multivariate Bayesian model (5.5) has an added benefit of allowing us to structure the parametric component of the model. Assuming \mathbf{X} to be a first order Markov process yields different structures discussed in Section 4. For example, the different forms of decaying dependence in the class of asymptotic independence can be modelled by setting the priors as $\alpha \sim \mathcal{U}(0, 1)$ and $\beta \sim \mathcal{U}(0, 1)$, independently. The appropriate structure can be determined using standard diagnostics, and if adopted in the modelling will lead to substantially improved efficiency. Imposing continuous priors on α and β induces a restriction to the class of asymptotically independent series, but both parameters can be arbitrary close to the boundaries of their support, ensuring that the behaviour of $\theta(x, m)$ and the extremal structure of dependence of the series are not affected at the high levels of interest. A reversible jump procedure (Green, 1995) could be added to the current algorithm in order to enable α and β to have prior masses on the support boundaries to ensure positive posterior probability of asymptotic dependence, see Coles and Pauli (2002) for an example of this type of construction.

The shape and scale for the prior variances of the components $\psi_{j,k}^2$ are taken to be $\nu_1 = \nu_2 = 2$ to make the model prefer numerous components with smaller variances to a few dispersed components. The posterior distribution for γ depends on the logarithm of the last weight in the truncation (cf. Appendix A) and can be numerically unstable, so a vague gamma prior truncated at small values is needed to ensure convergence. Conjugacy

of the prior densities allows analytical calculation of the posterior distributions for all parameters in model (5.5) except α and β , for which a Metropolis–Hastings step is needed. We use a regional adaptive scheme in Roberts and Rosenthal (2009) to avoid the choice of specific proposal variances. The posterior densities are mainly derived from Ishwaran and James (2002), and are given in Appendix A.

As noticed by Porteus *et al.* (2006), the Gibbs sampler used for model (5.5) leads to a clustering bias, because the weights do not satisfy the weak ordering $E(w_1) \geq \dots \geq E(w_{N-1})$. Papaspiliopoulos and Roberts (2008) suggested two different label switching moves to improve the mixing of the algorithm. Components cannot be simply swapped, as this would change the joint distribution of the weights. Label switching is not to be understood in its exact sense within this framework: if a switch between two mixture components is proposed and accepted, then only their means and variances are swapped and the index variables c_i of the data points belonging to these components are renumbered accordingly. We use this approach and adapt it to our semiparametric framework.

The results presented in Sections 6 and 7 are promising, but two aspects would benefit from improvement. Bayesian semiparametric inference provides a valuable approach to uncertainty in the Heffernan–Tawn model, but the procedure is not fully Bayesian, since the marginal distribution is fitted using maximum likelihood estimation. With further work we could include the fit for the marginal distribution within the fit for the dependence structure, but we would have to account for the temporal dependence between observations in order not to introduce bias. The second possibility for improvement pertains to the sampling of α_j and β_j in (3.3): the special cases corresponding to the boundaries of their support should correspond to Dirac masses, so reversible jump Markov chain Monte Carlo sampling (Green, 1995) could be used.

6 Simulation study

6.1 Bivariate data

We start by showing how the Bayesian semiparametric approach to inference can improve over the stepwise approach in a bivariate setting. The working assumption of Gaussianity for the residual variable Z is key to the stepwise process, and if it fails badly then the stepwise approach may perform poorly relative to the Bayesian semiparametric approach. To illustrate this we take Z to have a bimodal density, either a mixture of Gaussian densities, or a mixture of Laplace densities. As the former is a special case of the structure of the mixture components in the dependent Dirichlet process, we may expect a clear improvement in that case, but it is less clear what to expect in the latter case.

We generated data (X, Y) directly from the Heffernan–Tawn model with parameters (α, β) subject to $X > u$, for large $u > 0$, as follows:

1. Simulate X as $u + \text{Exp}(1)$;

2. Independently simulate Z from the required mixture model;
3. Let $Y = \alpha X + X^\beta Z$.

We selected the mixture for the bimodal distribution of Z such that the simulated (X, Y) data are split into two clear clusters for large X (see left panel of Figure 2 for an example).

We simulated 1000 data sets each with 400 points, roughly twice the number of exceedances available in our river flow application, and fitted the conditional model using the stepwise and the Bayesian semiparametric approaches. We compare the methods through the relative efficiency, measured as the ratio of the root mean squared error (RMSE) for the Bayesian approach to the RMSE of the stepwise approach. The estimators we consider in order to compute the efficiency are the mode, the mean, and the median of the posterior distribution of α and β for the Bayesian approach and the maximum likelihood estimators of α and β for the stepwise approach.

The benefits of the Bayesian semiparametric approach are clearly found, with similar relative efficiencies whether Z is simulated with a Gaussian or Laplace mixture. The relative efficiency is broadly 0.6 for α and in the range 0.5 – 0.65 for β depending on which of the three summary measures of the posterior distribution is chosen. The posterior number of components in the mixture is concentrated around 2 and 3, so the Bayesian semiparametric approach seems to capture the distribution of Z well. Figure 2 shows the joint sampling distribution of the estimators of (α, β) based on the two inference methods. The contours are similar, but suggest that the Bayesian approach estimates the true (α, β) more precisely.

Of key importance is the practical implication of this improvement, which is more naturally measured in terms of improved performance for estimating the threshold-based extremal index. Specifically we estimate $\theta(x, 1) = \Pr(Y < x \mid X > x)$, which requires accurate estimation of the distribution of Z as well as of α and β . The relative efficiency for $\theta(x, 1)$ is computed, where the true value for $\theta(x, 1)$ is obtained from a huge simulation from the true model. The relative efficiency varies over x , with values of 0.95 and 0.90 for the 99% and 99.9% quantiles, suggesting slight improvements within the range of the data. The relative efficiency reduces to 0.69 at the 99.99% quantile, suggesting that the real benefits in the Bayesian semiparametric approach arise when we extrapolate.

6.2 AR(1) process

We now compare the performances of the empirical, the stepwise and the Bayesian semiparametric inference procedures in estimating the threshold-based extremal index of a stationary time series. The data are generated from a first-order Markov process with Gaussian copula and exponential margins. This is equivalent to having a standard Gaussian AR(1) process and using the probability integral transform to obtain exponential marginal distributions. In Gaussian margins this process has lag τ autocorrelation $\rho_\tau = \rho^\tau$, where we consider the set of $\{-0.75, -0.5, \dots, 0.5, 0.75\}$ for the true value of ρ . For each of




Figure 2 – Results for simulation study based on the Heffernan–Tawn model with a bimodal residual distribution with two Laplace components. Left panel: simulated data, with the two components shown as squares and circles. Right panel: kernel density estimate based on 1000 estimates of (α, β) from the stepwise method (grey) and from the Bayesian semiparametric method (black) using the posterior medians as the summary statistic. The same contour levels are used for both density estimates. The true value is shown as a cross.

these values of ρ , the process is asymptotically independent, with extremal index $\theta = 1$, but it exhibits dependence at any subasymptotic threshold when $\rho \neq 0$. The true value for $\theta(x, m)$ is evaluated by computing the ratio of multivariate normal integrals

$$\theta(x, m) = \Pr(X_0 > x, X_1 < x, \dots, X_m < x) / \Pr(X_0 > x). \quad (6.1)$$

using the methods of Genz and Bretz (2009) and Genz *et al.* (2014). The use of exponential margins ensures that the GPD marginal model is exact for all thresholds, so any bias in the estimation of $\theta(x, m)$ can be attributed to inference for the dependence structure. A similar approach was taken by Eastoe and Tawn (2012).

The three methods are applied to 1000 data sets of length 8000, approximately the length of the winter flow data studied in Section 7. This procedure is repeated for each value of ρ in the range $\{-0.75, \dots, 0.75\}$. The empirical method simply estimates each of the probabilities in expression (6.1) empirically. Often called the runs estimate (Smith and Weissman, 1994), this is not defined beyond the largest value of the sample, whereas the other two methods do not suffer this weakness. In each case the marginal threshold u for the modelling and inference is fixed at the 95% empirical quantile of each series. Unlike the stepwise procedure, the Bayesian semiparametric approach enables us to constrain $\{(\alpha_j, \beta_j) : j = 1, \dots, m\}$, and this allows us to exploit our knowledge of the Markovian structure of the process to impose the constraints on the α_j and β_j discussed in Section 4, thus reducing the number of parameters from $2m$ to 1 or 2.

Level	m=1		m=4	
	Empirical	Stepwise	Empirical	Stepwise
98%	88	101	88	100
99%	68	92	71	98
99.99%	–	59	–	57

Table 1 – Ratios (%) of RMSEs computed with estimates of $\theta(x, m)$; the numerator of these efficiencies is always the RMSE estimate derived from the posterior median in the Bayesian semiparametric approach, and the denominator is either the RMSE corresponding to the runs estimate (Empirical) or to the stepwise estimate (Stepwise). Empty cells correspond to high levels of x for which estimates of $\theta(x, m)$ cannot be evaluated.

We estimate $\theta(x, m)$ for a range of high quantiles x and for declustering run-lengths $m = 1$ and 4. Table 1 shows the ratios of RMSEs of the posterior median of $\theta(x, m)$ from the Bayesian semiparametric approach and the empirical and the stepwise estimators in the particular case when $\rho = 0.5$. The Bayesian semiparametric estimator is always superior to the empirical estimator, with the advantage improving as x increases. For the stepwise approach the results are similar to those in Section 6.1: the two estimators are similar at low levels but the Bayesian semiparametric estimator performs better at higher levels. Figure 3 summarises the results for all values of ρ , showing a major improvement of our method over the stepwise approach for negative autocorrelation and short run-length, with increased gain at higher levels. In order to assess the effectiveness of imposing the Markovian structure in the Bayesian semiparametric approach, we also fitted the 1000 simulated time series with unconstrained α and β in the case $\rho = 0.5$. The efficiency of the unconstrained approach only declines relative to the constrained approach at high quantiles. For example the 68% in the bottom right of Table 1 increases to 75%.

We expect the Bayesian approach to gain accuracy in terms of frequentist coverage of $\theta(x, m)$, as it fits the data in one stage and thus provides a better measure of uncertainty. To assess this we considered bootstrap confidence intervals for the stepwise method and credible intervals for the Bayesian method, both of the type $[L_\alpha, \infty)$. Here L_α is the α th quantile of the distribution of the estimator, considering bootstrap estimates for the former and posterior samples for the latter. Using the same 1000 simulated data sets as earlier in this section, we computed the proportion of times that the true value of $\theta(x, m)$ would fall in these confidence or credible intervals, for a range of α -confidence levels from 5% to 95%, different run-lengths m , and several levels x for $\theta(x, m)$. The coverage performance is summarised in Figure 4, which shows the difference between the calculated and the nominal coverage α . Zero coverage error means perfect uncertainty assessment; positive and negative errors mean one-sided over- and under-coverage respectively. The stepwise approach over-estimates coverage for both levels of x and all α . At relatively low




Figure 3 – Ratios (%) of RMSEs computed with estimates of $\theta(x, m)$, the numerator being the RMSE derived from the posterior median in our approach and the denominator being the RMSE corresponding to the stepwise estimate. The cases $m = 1$ (circles) and $m = 4$ (diamonds) are illustrated for x at the 99% (empty symbols) and 99.99% (filled symbols) levels.




Figure 4 – Coverage error for $\theta(x, m)$ computed for confidence levels 0.05, 0.1, ..., 0.95, for $x = 98\%$ (empty symbols) and $x = 99.999\%$ (filled symbols), with $m = 1$, for the stepwise approach (circles) and the Bayesian semiparametric approach (diamonds).

x -levels of $\theta(x, m)$, the gain in coverage accuracy by the Bayesian approach is remarkable in particular at mid-coverage levels, but it shows no marked improvement for larger x .

7 Data analysis

River flooding can badly damage properties and have huge insurance costs. Large-scale floods in the UK since the year 2000 have caused insurance losses of £5 billion, and more than £400 million is spent each year on flood defences. Modelling the dependence of extreme water-levels is key to accurate prediction of flood risk.

Our application uses daily flows of the River Ray at Grendon Underwood, north-west of London, for the 47 winters from 1962 to 2008. We assume stationarity of the series over the winter months. River flows in this catchment are typically short-range dependent: after heavy rainfall the flow can reach high values before decreasing gradually as the river returns to its baseflow regime. We thus expect the flow to be dependent at extreme levels and at small lags, so a small run length m is required. For illustration we take $m = 1, 7$, with the former being the more appropriate.

Standard graphical methods (Coles, 2001) were used to choose the 95% empirical quantile as the marginal threshold. A sensitivity analysis on a range of thresholds gave results similar to those below. The Heffernan–Tawn model was then fitted to the data transformed to Laplace margins, with u as the 98% empirical quantile and $m = 1, 7$. A higher threshold was selected for the dependence modelling to ensure the independence of X_0 and \mathbf{Z} in approximation (3.3).

We first investigate the asymptotic structure of the data at different lags. We use $\chi_j(x) = \Pr(X_j > x \mid X_0 > x)$, $j = 1, \dots, m$, whose limit χ_j (cf. Section 1) either measures the degree of association within the asymptotic dependence class when $\chi_j > 0$ or indicates asymptotic independence when $\chi_j = 0$. A Monte Carlo integration similar to that used for estimating the posterior distribution of $\theta(x, m)$ is applied to get the posterior distribution of $\chi_j(x)$ for selected values of x and $j = 1, \dots, 7$ depicted in Figure 5, where the posterior densities are summarised using highest density regions (Hyndman, 1996). Convergence of $\chi_j(x)$ to 0 at all lags is supported by the model. As expected we observe a monotone decay in extremal dependence over time lag. The flexibility of the conditional model is well illustrated here, as the procedure establishes positive dependence at any finite level but anticipates asymptotic independence of successive daily flows.

From the estimates of $\chi_j(x)$ we expect $\theta = 1$ but $\theta(x, m) < 1$ for $x < x_F$. We computed estimates of $\theta(x, m)$ for several values of x based on the posterior distribution fitted with the Bayesian semiparametric approach and compared them with the stepwise approach and the empirical estimates. We give block bootstrap confidence intervals for the stepwise and empirical estimates, with a block length ensuring that winters are not split between blocks. Figure 1 shows estimates of $\theta(x, 4)$ obtained with the three methods, with $m = 4$ to show an intermediate estimate, as $\theta(x, 7) \leq \theta(x, 4) \leq \theta(x, 1)$. The three methods broadly agree at historical levels, with wider confidence intervals for the runs estimate. For higher x , the stepwise estimate predicts stronger dependence than does the Bayesian estimate.




Figure 5 – Posterior summaries for $\chi_j(x)$, $j = 1, \dots, 7$, with x the 95% (grey), 99% (blue), and 99.99% (red) marginal quantiles of the River Ray data. The different shades of colour indicate the 95%, 80%, and 50% highest density regions of the posterior densities. The black segments indicate the posterior medians.

Table 2 shows that the three methods agree closely for $m = 1$, with the posterior distribution giving slightly tighter credible intervals than the two other methods. For $m = 7$, the Bayesian approach seems to improve a little on the stepwise estimates when compared to the empirical estimates at low levels, partly because of the Markov constraints on the α_j and β_j , which also reduce the uncertainty. In terms of convergence of $\theta(x, m)$ to the extremal index θ , we observe that at very high levels and for both values of m , the estimates of $\theta(x, m)$ tend to the same values, which indicates coherence in the approach. This also illustrates the lesser concern of the choice of the run-length when we are interested in tail probabilities, typically when estimating cluster maxima distributions.

The Bayesian semiparametric approach appears to offer a more coherent basis for the extrapolation to the required levels and uncertainty quantification for design purposes than does the stepwise method. Although we assessed the performance of Bayesian semiparametric estimates of $\theta(x, m)$, but other conditional probabilities could also be estimated using this approach.

Level	$m = 1$			$m = 7$		
	Empirical	Stepwise	Bayesian	Empirical	Stepwise	Bayesian
95%	62 _(58,68)	62 _(61,70)	62 _(57,66)	35 _(33,41)	36 _(33,42)	33 _(27,38)
99%	90 _(83,95)	87 _(83,91)	88 _(83,91)	80 _(70,88)	69 _(66,79)	76 _(71,80)
99.9%	–	96 _(93,99)	97 _(94,99)	–	90 _(86,96)	98 _(96,99)
99.99%	–	99 _(96,100)	100 _(98,100)	–	98 _(93,100)	100 _(99,100)

Table 2 – Estimation of the threshold-based extremal index $\theta(x, m)$ (%) for four different levels of x and $m = 1, 7$ on the Ray River winter flow data, with 95% confidence intervals (CI) given as subscripts. Empirical: runs estimator (block bootstrap CI); Stepwise: Heffernan–Tawn method (block bootstrap CI); Bayesian: posterior median from the Bayesian semiparametric approach (quantiles of the posterior distribution).

A Posterior densities for the semiparametric model

Posterior density for $\boldsymbol{\mu}_k$: The posterior density for $\boldsymbol{\mu}_k = (\mu_{1,k}, \dots, \mu_{m,k})$ is multivariate Gaussian with independent margins, i.e.,

$$\mu_{j,k} \mid \mathbf{X}_j, \mathbf{X}_0, \psi_{j,k}^2, \alpha_j, \beta_j \stackrel{\text{ind}}{\sim} \mathcal{N} \left(M_{(\mu_{j,k})}, S_{(\mu_{j,k})}^2 \right), \quad j = 1, \dots, m, \quad k = 1, \dots, N,$$

with posterior mean and variance

$$M_{(\mu_{j,k})} = S_{(\mu_{j,k})}^2 \left(\frac{1}{\psi_{j,k}^2} \sum_{i \in C_k} \frac{X_{j,i} - \alpha_j X_{0,i}}{X_{0,i}^{\beta_j}} \right), \quad S_{(\mu_{j,k})}^2 = \left(\frac{n_k}{\psi_{j,k}^2} + \frac{1}{\psi_{(\mu),j}^2} \right)^{-1},$$

where $\mathbf{X}_j = (X_{j,1}, \dots, X_{j,n})$ are the observations at the j th lag, $C_k = \{i : c_i = k\}$, and $n_k = |C_k|$ is the number of observations in component k ; the $\psi_{(\mu),j}^2$ are the variance parameters of the prior for the components' means.

Posterior density for $\boldsymbol{\psi}_k^2$: The multivariate posterior density for the components' variances can be split into independent parts,

$$\psi_{j,k}^2 \mid \mathbf{X}_j, \mathbf{X}_0, \mu_{j,k}, \alpha_j, \beta_j \stackrel{\text{ind}}{\sim} \text{Inv-Gamma}(N_{1,j,k}, N_{2,j,k}), \quad j = 1, \dots, m, \quad k = 1, \dots, N,$$

with parameters

$$N_{1,j,k} = \frac{n_k}{2} + \nu_{1,j}, \quad N_{2,j,k} = \frac{1}{2} \sum_{i \in C_k} \frac{\left(X_{j,i} - \alpha_j X_{0,i} - \mu_{j,k} X_{0,i}^{\beta_j} \right)^2}{X_{0,i}^{2\beta_j}} + \nu_{2,j}.$$

Posterior density for \mathbf{c} : The posterior density is such that

$$c_i \mid \mathbf{X}, \boldsymbol{\mu}, \boldsymbol{\psi}^2, \boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{w} \stackrel{\text{ind}}{\sim} \sum_{k=1}^N W_{k,i} \delta_k, \quad i = 1, \dots, n,$$

where here for convenience \mathbf{X} , $\boldsymbol{\mu}$ and $\boldsymbol{\psi}^2$ are the matrices with rows $(\mathbf{X}_0, \dots, \mathbf{X}_m)$, $(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_N)$, and $(\boldsymbol{\psi}_1^2, \dots, \boldsymbol{\psi}_N^2)$ respectively; the stick-breaking weights are defined as

$$W_{k,i} = \frac{w_k}{\bar{W}_i} \prod_{j=1}^m \left[\frac{1}{X_{0,i}^{\beta_j} \psi_{j,k}^2} \exp \left\{ -\frac{1}{2} \frac{\left(X_{j,i} - \alpha_j X_{0,i} - \mu_{j,k} X_{0,i}^{\beta_j} \right)^2}{X_{0,i}^{2\beta_j} \psi_{j,k}^2} \right\} \right],$$

with $\bar{W}_i = \sum_{k=1}^N W_{k,i}$ ($i = 1, \dots, n$) constants that make the weights sum to 1.

Posterior density for \mathbf{w} : The posterior density for \mathbf{w} is generalised Dirichlet,

$$\mathbf{w} \mid \mathbf{c}, \gamma \sim \text{GDirichlet}(a_1, b_1, \dots, a_{N-1}, b_{N-1}),$$

where

$$a_k = 1 + n_k, \quad b_k = \gamma + \sum_{j=k+1}^N n_j, \quad k = 1, \dots, N - 1.$$

Posterior density for γ : The posterior density for the concentration parameter γ is

$$\text{Gamma} \left(N + \eta_1 - 1, \frac{\eta_2}{1 - \eta_2 \log w_N} \right) \mathbf{1}_{[\varepsilon, \infty)},$$

with $\varepsilon > 0$, typically $\varepsilon = 0.5$, and $\mathbf{1}$ is the indicator function.

References

- Asadi, P., Davison, A. C. and Engelke, S. (2015) Extremes on river networks. *Annals of Applied Statistics* **9**, 2023–2050.
- Ballani, F. and Schlather, M. (2011) A construction principle for multivariate extreme value distributions. *Biometrika* **98**, 633–645.
- de Carvalho, M. and Ramos, A. (2012) Bivariate extreme statistics, II. *Revstat* **10**, 83–107.
- Cheng, L., Gilleland, E., Heaton, M. J. and AghaKouchak, A. (2014) Empirical Bayes estimation for the conditional extreme value model. *Stat* **3**, 391–406.
- Coles, S. (2001) *An Introduction to Statistical Modeling of Extreme Values*. London: Springer.
- Coles, S. G. and Pauli, F. (2002) Models and inference for uncertainty in extremal dependence. *Biometrika* **89**, 183–196.
- Coles, S. G. and Tawn, J. A. (1994) Statistical methods for multivariate extremes: An application to structural design (with discussion). *Journal of the Royal Statistical Society Series C* **43**, 1–48.
- Connor, R. J. and Mosimann, J. E. (1969) Concepts of independence for proportions with a generalization of the Dirichlet distribution. *Journal of the American Statistical Association* **64**, 194–206.
- Cooley, D., Davis, R. A. and Naveau, P. (2010) The pairwise beta distribution: a flexible parametric multivariate model for extremes. *Journal of Multivariate Analysis* **101**, 2103–2117.

- Das, B. and Resnick, S. I. (2011) Conditioning on an extreme component: model consistency with regular variation on cones. *Bernoulli* **17**, 226–252.
- Davis, R. A. and Mikosh, T. (2009) The extremogram: a correlogram for extreme events. *Bernoulli* **15**, 977–1009.
- Davison, A. C. and Smith, R. L. (1990) Models for exceedances over high thresholds (with discussion). *Journal of the Royal Statistical Society Series B* **52**, 393–442.
- Eastoe, E. F. and Tawn, J. A. (2012) Modelling the distribution of the cluster maxima of exceedances of subasymptotic thresholds. *Biometrika* **99**, 43–55.
- Einmahl, J. H. J., de Haan, L. and Piterbarg, V. I. (2001) Nonparametric estimation of the spectral measure of an extreme value distribution. *The Annals of Statistics* **29**, 1401–1423.
- Einmahl, J. H. J. and Segers, J. (2009) Maximum empirical likelihood estimation of the spectral measure of an extreme-value distribution. *The Annals of Statistics* **37**, 2953–2989.
- Fawcett, L. and Walshaw, D. (2006) A hierarchical model for extreme wind speeds. *Journal of the Royal Statistical Society Series C* **55**, 631–646.
- Fawcett, L. and Walshaw, D. (2007) Improved estimation for temporally clustered extremes. *Environmetrics* **18**, 173–188.
- Ferguson, T. S. (1973) A Bayesian analysis of some nonparametric problems. *The Annals of Statistics* **1**, 209–230.
- Ferro, C. A. T. and Segers, J. (2003) Inference for clusters of extreme values. *Journal of the Royal Statistical Society Series B* **65**, 545–556.
- Genest, C., Ghoudi, K. and Rivest, L.-P. (1995) A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika* **82**, 543–552.
- Genz, A. and Bretz, F. (2009) *Computation of Multivariate Normal and t Probabilities. Lecture Notes in Statistics*. Heidelberg: Springer.
- Genz, A., Bretz, F., Miwa, T., Mi, X., Leisch, F., Scheipl, F. and Hothorn, T. (2014) *mvtnorm: Multivariate Normal and t Distributions*. R package version 1.0-2.
- Green, P. J. (1995) Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**, 711–732.

- de Haan, L. and de Ronde, J. (1998) Sea and wind: Multivariate extremes at work. *Extremes* **1**, 7–45.
- Hall, P. and Tadjvidi, N. (2000) Distribution and dependence-function estimation for bivariate extreme-value distributions. *Bernoulli* **6**, 835–844.
- Heffernan, J. E. and Resnick, S. I. (2007) Limit laws for random vectors with an extreme component. *The Annals of Applied Probability* **17**, 537–571.
- Heffernan, J. E. and Tawn, J. A. (2004) A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society Series B* **66**, 497–546.
- Hjort, N. L., Holmes, C., Müller, P. and Walker, S. G. (2010) *Bayesian Nonparametrics*. New York: Cambridge University Press.
- Hsing, T., Hüsler, J. and Leadbetter, M. R. (1988) On the exceedance point process for a stationary sequence. *Probability Theory and Related Fields* **78**, 97–112.
- Hyndman, R. J. (1996) Computing and graphing highest density regions. *The American Statistician* **50**, 120–126.
- Ishwaran, H. and James, L. F. (2001) Gibbs sampling methods for stick-breaking priors. *Journal of the American Statistical Association* **96**, 161–172.
- Ishwaran, H. and James, L. F. (2002) Approximate Dirichlet process computing in finite normal mixtures: smoothing and prior information. *Journal of Computational and Graphical Statistics* **11**, 508–532.
- Ishwaran, H. and Zarepour, M. (2000) Markov chain Monte Carlo in approximate Dirichlet and beta two-parameter process hierarchical models. *Biometrika* **87**, 371–390.
- Keef, C., Papastathopoulos, I. and Tawn, J. A. (2013) Estimation of the conditional distribution of a multivariate variable given that one of its components is large: Additional constraints for the Heffernan and Tawn model. *Journal of Multivariate Analysis* **115**, 396–404.
- Keef, C., Tawn, J. and Svensson, C. (2009) Spatial risk assesment for extreme river flows. *Journal of the Royal Statistical Society Series C* **58**, 601–618.
- Kotz, S. and Nadarajah, S. (2000) *Extreme Value Distributions: Theory and Applications*. London: Imperial College Press.
- Kratz, M. F. and Rootzén, H. (1997) On the rate of convergence for extremes of mean square differentiable stationary normal processes. *Journal of Applied Probability* **34**, 908–923.

- Kulik, R. and Soulier, P. (2015) Heavy tailed time series with extremal independence. *Extremes* **18**, 273–299.
- Leadbetter, M. R. (1983) Extremes and local dependence in stationary sequences. *Probability Theory and Related Fields* **65**, 291–306.
- Leadbetter, M. R., Lindgren, G. and Rootzén, H. (1983) *Extremes and Related Properties of Random Sequences and Processes*. New York: Springer.
- Ledford, A. W. and Tawn, J. A. (1996) Statistics for near independence in multivariate extreme values. *Biometrika* **83**, 169–187.
- Ledford, A. W. and Tawn, J. A. (2003) Diagnostics for dependence within time series extremes. *Journal of the Royal Statistical Society Series B* **65**, 521–543.
- Marron, J. S. and Wand, M. P. (1992) Exact mean integrated squared error. *The Annals of Statistics* **20**, 712–736.
- Mitra, A. and Resnick, S. I. (2013) Modeling multiple risks: hidden domain of attraction. *Extremes* **16**, 507–538.
- O’Brien, G. L. (1987) Extreme values for stationary and Markov sequences. *The Annals of Probability* **15**, 281–291.
- Papaspiliopoulos, O. and Roberts, G. O. (2008) Retrospective Markov chain Monte Carlo methods for Dirichlet process hierarchical models. *Biometrika* **95**, 169–186.
- Papastathopoulos, I., Strokorb, K., Tawn, J. A. and Butler, A. (2016) Extreme events of Markov chains. To appear.
- Peng, L. and Qi, Y. (2004) Discussion of “A conditional approach for multivariate extreme values”, by J. E. Heffernan and J. A. Tawn. *Journal of the Royal Statistical Society Series B* **66**, 541–542.
- Poon, S.-H., Rockinger, M. and Tawn, J. A. (2003) Modelling extreme-value dependence in international stock markets. *Statistica Sinica* **13**, 929–953.
- Porteus, I., Ihler, A., Smyth, P. and Welling, M. (2006) Gibbs sampling for (coupled) infinite mixture models in the stick breaking representation. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence*.
- Reich, B. J., Shaby, B. A. and Cooley, D. (2014) A hierarchical model for serially-dependent extremes: a study of heat waves in the Western US. *Journal of Agricultural, Biological, and Environmental Statistics* **19**, 119–135.
- Robert, C. Y. (2013) Automatic declustering of rare events. *Biometrika* **100**, 587–606.

- Roberts, G. O. and Rosenthal, J. S. (2009) Examples of Adaptive MCMC. *Journal of Computational and Graphical Statistics* **18**, 349–367.
- Segers, J. (2003) Functionals of clusters of extremes. *Advances in Applied Probability* **35**, 1028–1045.
- Sethuraman, J. (1994) A constructive definition of Dirichlet priors. *Statistica Sinica* **4**, 639–650.
- Smith, R. L., Tawn, J. A. and Coles, S. G. (1997a) Markov chain models for threshold exceedances. *Biometrika* **84**, 249–268.
- Smith, R. L., Tawn, J. A. and Coles, S. G. (1997b) Markov chain models for threshold exceedances. *Biometrika* **84**, 249–268.
- Smith, R. L. and Weissman, I. (1994) Estimating the extremal index. *Journal of the Royal Statistical Society Series B* **56**, 515–528.
- Süveges, M. (2007) Likelihood estimation of the extremal index. *Extremes* **10**, 41–55.
- Süveges, M. and Davison, A. C. (2012) A case study of a "Dragon-King": the 1999 Venezuelan catastrophe. *The European Physical Journal Special Topics* **205**, 131–146.
- Winter, H. C. and Tawn, J. A. (2016) Modelling heatwaves in central France: a case study in extremal dependence. *Journal of the Royal Statistical Society Series C* **65**, 345–365.