



# Estimation and uncertainty quantification for extreme quantile regions

Boris Beranger<sup>1</sup> · Simone A. Padoan<sup>2</sup>  · Scott A. Sisson<sup>1</sup>

Received: 22 April 2019 / Revised: 6 September 2019 / Accepted: 16 September 2019 /  
Published online: 16 December 2019  
© Springer Science+Business Media, LLC, part of Springer Nature 2019

## Abstract

Estimation of extreme quantile regions, spaces in which future extreme events can occur with a given low probability, even beyond the range of the observed data, is an important task in the analysis of extremes. Existing methods to estimate such regions are available, but do not provide any measures of estimation uncertainty. We develop univariate and bivariate schemes for estimating extreme quantile regions under the Bayesian paradigm that outperforms existing approaches and provides natural measures of quantile region estimate uncertainty. We examine the method's performance in controlled simulation studies. We illustrate the applicability of the proposed method by analysing high bivariate quantiles for pairs of pollutants, conditionally on different temperature gradations, recorded in Milan, Italy.

**Keywords** Air pollution · Bayesian nonparametrics · Bernstein polynomials · Extremal dependence · Extreme quantile regions · Max-stable distributions

---

**Electronic supplementary material** The online version of this article (<https://doi.org/10.1007/s10687-019-00364-0>) contains supplementary material, which is available to authorized users.

---

✉ Boris Beranger  
B.Beranger@unsw.edu.au

Simone A. Padoan  
Simone.Padoan@unibocconi.it

Scott A. Sisson  
Scott.Sisson@unsw.edu.au

<sup>1</sup> School of Mathematics and Statistics, University of New South Wales, Sydney, Australia

<sup>2</sup> Department of Decisions Sciences, Bocconi University, Milan, Italy

## 1 Introduction

Estimating quantiles and how they can change depending on influential predictors (i.e. quantile regression) is a recurring problem in many applied fields including medicine, survival analysis, economics, finance and environmental science (e.g. Yu et al. 2003). Assessing high quantiles, associated with fixed, low occurrence probabilities that lie beyond the range of  $n$  existing data observations, is of crucial importance in risk management. The solution to such a problem is not obvious when none of the observed data points exceed this event. It is likely that the exceedance probability,  $p$ , is smaller than  $1/n$ . This is an extreme value problem.

More precisely, consider a random variable  $X$  with distribution function  $F$  defined on  $\mathbb{R}_+ := (0, \infty)$ . For  $0 < p < 1$ , let  $Q(p) := F^{\leftarrow}(1 - p)$  be the  $(1 - p)$ -th quantile of  $F$ , where  $F^{\leftarrow}$  is the left-continuous inverse function of  $F$ , i.e.  $F^{\leftarrow}(x) := \inf\{y : F(y) \geq x\}$ . Let  $X_1, \dots, X_n$  be independent random variables with distribution  $F$ . Interest is in estimating the quantile  $Q(p)$  when the exceedance probability  $p$  is very small. We refer to  $Q$  as an *extreme quantile*. The extreme-value approach (e.g. Beirlant et al. 2004; de Haan and Ferreira 2006) assumes that as the sample size  $n$  grows to infinity a suitable asymptotic probabilistic model can be used to approximate the desired quantile. In the asymptotic setting the exceedance probability depends on  $n$ , i.e.  $p = p_n$  with  $p \rightarrow 0$  as  $n \rightarrow \infty$ . Hence, for a sufficiently large number of observations, extreme-value theory can be used to compute an approximation,  $Q_n$ , of the extreme quantile.

A more challenging problem is estimating, for some fixed probability, an extreme bivariate region – a subset of two-dimensional Euclidean space – in which a future event would fall when none of the two-dimensional observations fall in such a region, and are likely to lie far from it (e.g. de Haan and Ferreira 2006, Ch. 6). In some applications there may not be any specific shape for the critical region, whereas in others it may be well defined. One illustration of the former is air pollution monitoring, where the critical region is a set of combinations of two pollutants' concentrations where at least one of them is at a high level. Here, the shape of the critical region depends on the type of pollutants and the intensity of their dependence.

Einmahl et al. (2013) proposed a simple and practically useful method for defining critical regions which are generated through the level sets of a probability density function  $f$ , under suitable conditions (see Section 2.2). Precisely, a critical region is the level set of  $f$  given by

$$\mathcal{Q} = \{\mathbf{x} \in \mathbb{R}_+^2 : f(\mathbf{x}) \leq \alpha\}, \quad (1.1)$$

with  $\alpha > 0$ , such that  $\mathbb{P}(\mathcal{Q}) = p$  for some very small  $p \in (0, 1)$ . We refer to  $\mathcal{Q}$  as an *extreme quantile region*. In particular  $\mathcal{Q}^c = \{\mathbf{x} \in \mathbb{R}_+^2 : f(\mathbf{x}) > \alpha\}$  is the set with smallest area such that  $\mathbb{P}(\mathcal{Q}^c) = 1 - p$ . See also Cooley et al. (2017) for an alternative approach. Let  $\mathbf{X}_1, \dots, \mathbf{X}_n$  be independent two-dimensional random vectors with distribution function  $F$ . Similar to the univariate case, the extreme-value approach suggests adopting an appropriate asymptotic probabilistic model as  $n$  approaches infinity. From this, a set  $\mathcal{Q}_n$ , that is not fixed but depends on  $n$ , may be derived such that  $\mathbb{P}(\mathcal{Q}_n) = p_n$  with  $p \equiv p_n$  and  $p \rightarrow 0$  as  $n \rightarrow \infty$ , which provides a sufficiently

close approximation to the extreme quantile region. In particular, in both univariate and bivariate cases we assume the mild condition that  $np \rightarrow c \in [0, \infty)$  as  $n \rightarrow \infty$ .

Beyond a definition and a method of estimation of such quantiles, a critical aspect in practice is to provide some quantification of the uncertainty around the estimates. Several estimators of extreme quantiles already exist in the univariate setting (e.g. Beirlant et al. 2004; de Haan and Ferreira 2006; Wang and Li 2015). However, quantifying the uncertainty of these estimators can be difficult as their asymptotic variance depends on the parameters of second-order conditions which can be problematic to estimate in some settings. Estimators of extreme quantile regions also exist (Cai et al. 2011; Einmahl et al. 2013) but to the best of our knowledge measures of their uncertainty are not available. A readily applicable uncertainty quantification for extreme quantiles is the methodological contribution of this article.

In this paper we develop a Bayesian approach for inferring both univariate extreme quantiles as well as extreme quantile regions. In the univariate case, we define a parametric Bayesian method for the extreme quantiles by exploiting the well known and widely used censored-likelihood approach (e.g. Prescott and Walden 1983; Smith 1994; Ledford and Tawn 1996; Huser et al. 2016; Bienvenüe and Robert 2017), based on the likelihood function of the univariate Generalized Extreme Value (GEV) family of distributions (e.g. Sisson and Coles 2003). Inference is performed using an adaptive random-walk Metropolis-Hastings algorithm (Garthwaite et al. 2016). In the bivariate case we combine the univariate approach for the estimation of the GEV marginal parameters with the non-parametric Bayesian approach for the estimation of the extremal dependence proposed by Marcon et al. (2016). As a result we obtain a semi-parametric Bayesian inferential method based on the censored-likelihood corresponding to a suitable two-dimensional extreme-value distribution. Through such a censored-likelihood it is possible to simultaneously estimate the dependence structure of the extreme-value distribution together with the parameters of its margins, which are in turn members of the GEV class. The components of such a pseudo-posterior distribution are combined to produce a pseudo-posterior distribution for the extreme quantile regions. Accordingly this approach allows for the direct estimation of extreme quantile intervals and quantile regions, as well as clear measures of their uncertainty.

There are many settings in environmental science where the analysis of univariate extreme quantiles is of particular concern. These include evaluating tropical cyclone wind speed conditional on certain climate variables (Jagger and Elsner 2009), precipitation conditional on global climate model projections (Friederichs and Thorarinsdottir 2012), both precipitation and wind power given outputs from numerical weather prediction models (Bremnes, 2004a, b), and ozone and particulate matter concentrations conditional on meteorological variables (Porter et al. 2015). However it is clear that some environmental variables depend on others: e.g. wind speed is dependent on wind gust or air pressure (e.g. Marcon et al. 2017), and ozone is dependent on particulate matter (e.g. Falk et al. 2019). Similarly, when designing offshore structures, engineers need estimates of extreme quantiles of oceanographic variables. Due the strong dependence between the involved variables e.g. wave heights and wind speeds, bivariate extreme value tools (in the simplest case) are sometimes necessary (Naess and Karpa 2015). As a consequence, the joint study of two (or

more) dependent variables will produce a more comprehensive and accurate analysis because it is based on more available information.

The goal of this article is to provide an inferential framework (i.e. a method and software to implement it) that can provide estimates of extreme bivariate quantile regions, which can then be used to inform part of a larger study. For example, reliability engineers designing offshore platforms can use this framework as both a visual and quantitative tool to produce extreme quantile curves (such as those in Fig. 5) that can be informative for reliability design. Here, we particularly focus on the problem of air pollution.

According to the World Health Organisations (WHO) and the Global Ambient Air Quality Database (2018 update, <https://www.who.int/airpollution/en/>), air pollution kills an estimated 7 million people worldwide each year, 4.2 million of which as a result of exposure to ambient (outdoor) air pollution. Furthermore, 25% of all heart diseases deaths are attributable to air pollution. Based on WHO guidelines (World Health Organization 2006) the European policy on air-quality standards (Guerreiro et al. 2016) regulates emissions of the pollutants particulate matter (PM<sub>10</sub>), ozone (O<sub>3</sub>), nitrogen oxide (NO), nitrogen dioxide (NO<sub>2</sub>) and sulphur dioxide (SO<sub>2</sub>), aiming to reduce the negative impact that they have on human health, the environment and climate. For example, short-term pollutant concentrations (obtained through the empirical quantiles of pollutants time series) that should not be exceeded in order to protect against air pollution peaks are: a daily average concentration of 50  $\mu\text{g}/\text{m}^3$  (90.4 percentile) for PM<sub>10</sub>, a daily maximum concentration of 120  $\mu\text{g}/\text{m}^3$  (93.2 percentile) for O<sub>3</sub>, an hourly average concentration of 200  $\mu\text{g}/\text{m}^3$  (99.8 percentile) for NO<sub>2</sub> and a daily average concentration of 125  $\mu\text{g}/\text{m}^3$  (99.2 percentile) for SO<sub>2</sub>. We refer to such concentrations as the pollutant limit thresholds. In contrast, long-term pollutant thresholds are set based on yearly mean concentrations, rather than using quantiles. See Guerreiro et al. (2016) for details. The above thresholds are determined for each individual pollutant, although it is well known that some pollutants are dependent on each other (e.g. Dahlhaus, 2000; Clapp and Jenkin 2001; Heffernan and Tawn 2004; World Health Organization 2006).

With our methodology (and software) we are able to estimate extreme quantiles while taking into account the dependence between pairs of pollutants, and thereby be informed about potentially dangerous pollutant concentration combinations. This will be of particular interest to air pollution emission regulators and public health analysts. Estimating extreme pollutant concentrations conditional on e.g. meteorological variables can also provide information on whether some predictors are important for understanding the evolution of air pollution at high concentration levels. Here, we investigate this explicitly through an analysis of the behaviour air pollution data recorded over the last  $\sim 20$  years in Milan, Italy.

The remainder of this article is organised as follows. In Section 2 we briefly review the extreme-value approach for approximating the extreme quantiles in both the univariate (Section 2.1) and bivariate (Section 2.2) cases. In Section 3 we introduce our Bayesian semi-parametric approach for inferring extreme quantiles and extreme quantile regions. Section 4 provides an extensive simulation study examining the performance of the proposed method. The article concludes with an analysis of the extremes of air pollution recorded in Milan (Section 5) followed by a Discussion.

## 2 Estimating extreme quantiles

### 2.1 The univariate case

Let  $F$  be a distribution on  $\mathbb{R}_+$ , and assume that  $F$  is in the domain of attraction of the GEV family of distributions,  $F \in \mathcal{D}(G)$ . Then, for  $n = 1, 2, \dots$  there are norming constants  $a_n > 0$  and  $b_n$  such that

$$F^n(a_n x + b_n) \xrightarrow{n \rightarrow \infty} \exp\left(-(1 + \gamma x)_+^{-1/\gamma}\right) =: G(x; \gamma), \quad (2.1)$$

where  $\gamma > 0$  is the extreme-value index that describes the heaviness of the tail of the distribution  $G$  and  $(a)_+ = \max(a, 0)$  (see e.g. de Haan and Ferreira 2006, Ch. 1).

There are several extreme-value based methods for modelling extreme quantiles; see e.g. de Haan and Ferreira (2006, Ch. 1, 4) and Wang and Li (2015) for a compendium. Here, we briefly review one which is useful for our purposes. By de Haan and Ferreira (2006, Theorem 1.1.6) the result in Eq. 2.1 is also equivalent to the result

$$t\{1 - F(a(t)x + U(t))\} \xrightarrow{t \rightarrow \infty} (1 + \gamma x)^{-1/\gamma}, \quad (2.2)$$

where  $U(t) = F^{\leftarrow}(1 - 1/t)$  for  $t > 1$  and  $a(\cdot)$  is a suitable function (de Haan and Ferreira 2006, Theorem 1.1.6). Let  $k \equiv k_n$  and assume that  $k \rightarrow \infty$  and  $k/n \rightarrow 0$  as  $n \rightarrow \infty$ . Set  $t = n/k$  and  $y = a(n/k)x + U(n/k)$ . Then, from Eq. 2.2 we obtain

$$F(y) \approx 1 - \frac{k}{n} \left(1 + \gamma \frac{y - \mu}{\sigma}\right)_+^{-1/\gamma}, \quad n \rightarrow \infty, \quad (2.3)$$

where  $\mu, \sigma > 0$  are location and scale parameters. By few manipulations and de Haan and Ferreira (2006, Theorem 1.1.8) we have that  $a(t) \approx \sigma$  and  $U(t) \approx \mu$  as  $t \rightarrow \infty$ . Result Eq. 2.1 is also equivalent to

$$\frac{U(tx) - U(t)}{a(t)} \xrightarrow{t \rightarrow \infty} \frac{x^\gamma - 1}{\gamma}, \quad (2.4)$$

see de Haan and Ferreira (Theorem 1.1.6, 2006, for details). Since  $Q(p) \equiv U(1/p)$ , from Eq. 2.4 we obtain, using arguments in de Haan and Ferreira (2006, Ch. 3.1),

$$Q(p) \approx \mu + \sigma \frac{\left(\frac{k}{np}\right)^\gamma - 1}{\gamma} \quad \text{as } n \rightarrow \infty. \quad (2.5)$$

A suitable adjustment of the GEV distribution allows us to also derive the approximate quantile function in Eq. 2.5. Specifically, for some threshold  $s > 0$  and for  $x \geq s$ , let  $F_s(x) := \mathbb{P}(X \leq x | X > s)$ . Using de Haan and Ferreira (2006, Theorem 1.1.6), for  $a_n = a(n)$  and  $b_n = U(n)$ , for all  $x \geq s$  we can write

$$F_{a_n s + b_n}(a_n x + b_n) \xrightarrow{n \rightarrow \infty} 1 - \left(1 + \gamma \frac{x - s}{\tilde{\sigma}}\right)_+^{-1/\gamma} =: H((x - s)/\tilde{\sigma}; \gamma), \quad (2.6)$$

where  $\tilde{\sigma} = 1 + \gamma s$  and  $H(\cdot; \gamma)$  is the Generalized Pareto (GP) family of distributions. Set  $y = a_n x + b_n$  and  $t := a_n s + b_n$ , then for large  $t$  and any  $y \geq t$  we obtain

$$\begin{aligned} F(y) &\approx F(t) + \{1 - F(t)\} H((y - t)/\tilde{\sigma}; \gamma) \\ &\approx \exp[-\{1 - F(t)\} \{1 - H((y - t)/\tilde{\sigma}; \gamma)\}] \\ &\approx \exp\left[-t^{-1} \{1 - H((t - \mu)/\sigma; \gamma)\} \{1 - H((y - t)/\tilde{\sigma}; \gamma)\}\right] \\ &= \exp\left(-t^{-1} (1 + \gamma(y - \mu)/\sigma)_+^{-1/\gamma}\right) \equiv G^{1/t}((y - \mu)/\sigma; \gamma), \quad (2.7) \end{aligned}$$

where  $\tilde{\sigma} = \sigma + \gamma(t - \mu)$ . Writing again  $F(y) = 1 - p$  and  $t = n/k$  in Eq. 2.7, and by noting that  $-\log(1 - p) \approx p$  for  $p \rightarrow 0$ , then the quantile function in Eq. 2.5 can be obtained by inverting the expression in Eq. 2.7 with respect to  $p$ .

## 2.2 The bivariate case

Let  $X = (X_1, X_2)$  be a random vector with joint distribution function  $F$  on  $\mathbb{R}_+^2$  with margins  $F_j$ ,  $j = 1, 2$ , and probability density function  $f$ . Assuming that  $F \in \mathcal{D}(G)$ , then there are sequences of norming constants  $a_n > 0$  and  $b_n$  such that

$$F^n(a_n \mathbf{x} + b_n) \xrightarrow{n \rightarrow \infty} G(\mathbf{x}), \quad (2.8)$$

where  $G$  is the so-called bivariate max-stable distribution (de Haan and Ferreira 2006, Ch 6), whose margins  $G(x_j; \gamma_j)$  are members of the GEV family Eq. 2.1 with tail indices  $\gamma_j > 0$ ,  $j = 1, 2$ . In the following we denote  $G_*(\mathbf{x}) := G(x_1^{\gamma_1}, x_2^{\gamma_2})$ ,  $\mathbf{x} \in \mathbb{R}_+^2$ , as an extreme-value distribution with unit-Fréchet margins:  $G_*(x_1, \infty) = \exp(-1/x_1)$  and  $G_*(\infty, x_2) = \exp(-1/x_2)$  for every  $x_1, x_2 > 0$ .

The convergence result in Eq. 2.8 implies convergence at both marginal and dependence levels. For marginal convergence we have that Eqs. 2.1, 2.2 and 2.3 hold for  $F_j$ , with  $a_{n,j}$ ,  $b_{n,j}$ ,  $a_j(t)$ ,  $U_j(t)$  and  $y_j = a_j(n/k)x_j + U_j(n/k)$  with  $j = 1, 2$ . Furthermore,  $a_j(t) \approx \sigma_j > 0$  and  $U_j(t) \approx \mu_j > 0$  as  $t \rightarrow \infty$ . For convergence of the dependence structure, for every  $(x, y) \in (0, \infty]^2 \setminus \{(\infty, \infty)\}$ , from de Haan and Ferreira (2006, Ch 6.1.2) we have that

$$t(1 - F(U_1(tx_1), U_2(tx_2))) \xrightarrow{t \rightarrow \infty} -\log G_*(\mathbf{x}). \quad (2.9)$$

The result in Eq. 2.9 implies the existence of a measure  $\nu$ , named the exponent measure (see de Haan and Ferreira 2006, Ch 6 for details) such that for every  $\mathbf{x} \in \mathbb{R}_+^2$ ,

$$-\log G_*(\mathbf{x}) = \nu(\{\mathbf{v} \in \mathbb{R}_+^2 : v_1 > x_1 \text{ or } v_2 > x_2\}) = 2 \int_0^1 \left( \frac{w}{x_1} \vee \frac{1-w}{x_2} \right) H(dw),$$

where  $H$  is a probability measure on  $[0, 1]$  satisfying the mean constraint  $\int_0^1 w H(dw) = \int_0^1 (1-w) H(dw) = 1/2$ . In the following we denote both the probability measure and its distribution function by  $H$ , with the difference being determined by the context.

The extreme quantile regions method introduced in Einmahl et al. (2013) requires also some further assumptions at the density level. The density  $f$  is assumed to be decreasing in each coordinate, outside of  $(0, M]^2$  for some  $M > 0$ , and bounded

away from zero on  $(0, M]^2$ . There exist a nonnegative Lebesgue integrable function  $g$  such that for every  $\mathbf{x} \in \mathbb{R}_+^2$ ,

$$\nu(\{\mathbf{v} \in \mathbb{R}_+^2 : v_1 > x_1 \text{ or } v_2 > x_2\}) = \iint_{\{v_1 > x_1 \text{ or } v_2 > x_2\}} g(\mathbf{v}) d\mathbf{v}.$$

and on  $\mathbb{R}_+^2$ ,

$$tU_1(t)U_2(t)f(U_1(tx_1), U_2(tx_2)) \xrightarrow{t \rightarrow \infty} (\gamma_1\gamma_2)^{-1}x_1^{1-\gamma_1}x_2^{1-\gamma_2}g(\mathbf{x}) =: q(\mathbf{x}). \quad (2.10)$$

We refer to  $g$  and  $q$  as the density of the exponent measure and the *basic density* function, respectively. By the above assumption on  $g$  it follows that  $H$  has a continuous density  $h := \partial H / \partial w$  on  $(0, 1)$  with no atoms at 0 and 1, i.e.  $H(\{0\}) = H(\{1\}) = 0$ . By the homogeneity of  $g$  we have that  $h(w) = 2^{-1}g(w, 1-w)$ , where  $r = x_1 + x_2$  and  $w = x_1/r$ . We refer to  $H$  and  $h$  as the angular measure and density, respectively.

The extreme-value approach for modelling extreme quantile regions  $\mathcal{Q}$  in Eq. 1.1, works with the set  $\mathcal{Q}_n = \{\mathbf{x} \in \mathbb{R}^2 : f(\mathbf{x}) \leq \alpha\}$ , where  $\alpha = \alpha_n$  is not fixed but depends on the sample size  $n$  such that  $\mathbb{P}(\mathcal{Q}_n) = p$  with  $p = p_n \rightarrow 0$  as  $n \rightarrow \infty$ . In particular, Einmahl et al. (2013) suggest focusing on the fixed set  $\mathcal{S} = \{\mathbf{x} \in \mathbb{R}_+^2 : q(\mathbf{x}) \leq 1\} = \{\mathbf{x} \in \mathbb{R}_+^2 : r \geq q_*^{-1}(w), w \in [0, 1]\}$  that we call the *basic set*, where

$$q_*(w) = q(w, 1-w)^{-\frac{1}{1+\gamma_1+\gamma_2}}, \quad q(w, 1-w) = \frac{2w^{1-\gamma_1}(1-w)^{1-\gamma_2}h(w)}{\gamma_1\gamma_2} \quad (2.11)$$

and where  $q(w, 1-w)$  is obtained from the relation  $q(rw, r(1-w)) = q(w, 1-w)/r^{(1+\gamma_1+\gamma_2)}$ . Abusing terminology refer to  $q_*$  as the *angular basic density* function. According to the exponent measure, the size of  $\mathcal{S}$  is

$$\nu(\mathcal{S}) = 2 \int_{[0,1]} q_*^{-1}(w) h(w) dw. \quad (2.12)$$

The idea is to then inflate the basic set  $\mathcal{S}$  into an extreme set,  $\tilde{\mathcal{Q}}_n$ , depending on  $n$ , so that  $\mathbb{P}(\tilde{\mathcal{Q}}_n) \approx p$  for  $n \rightarrow \infty$  and such that  $\tilde{\mathcal{Q}}_n$  is a good approximation of  $\mathcal{Q}_n$ , i.e. so that  $\mathbb{P}(\mathcal{Q}_n \Delta \tilde{\mathcal{Q}}_n)/p \rightarrow 0$  as  $n \rightarrow \infty$ , where  $B \Delta D = B \setminus D \cup D \setminus B$ , for two nonempty sets  $B$  and  $D$ . Here we consider a slightly different definition of the set  $\tilde{\mathcal{Q}}_n$  to that given in Einmahl et al. (2013). Specifically, by exploiting the fact that  $U_j(n/k) \approx \mu_j$  and  $a_j(n/k) \approx \sigma_j$ ,  $j = 1, 2$  as  $n \rightarrow \infty$ , then for large  $n$ , we have that  $\tilde{\mathcal{Q}}_n$  can be approximated as

$$\tilde{\mathcal{Q}}_n \approx \left\{ \left( \mu_1 + \sigma_1 \frac{\left( \frac{kv(\mathcal{S})x_1}{np} \right)^{\gamma_1} - 1}{\gamma_1}, \mu_2 + \sigma_2 \frac{\left( \frac{kv(\mathcal{S})x_2}{np} \right)^{\gamma_2} - 1}{\gamma_2} \right) : (x_1, x_2) \in \mathcal{S} \right\}. \quad (2.13)$$

The approximation in Eq. 2.13 is consistent with the formula in Eq. 2.5 used to approximate the univariate extreme quantiles. In Section 3.2 we show how the bivariate max-stable distribution in Eq. 2.8 can be used to estimate the extreme quantile region in Eq. 2.13.

### 3 Inference

#### 3.1 The univariate case

We describe an approximate Bayesian framework for estimating the extreme quantile in Eq. 2.5 for small  $p$  (where the meaning of “small  $p$ ” is given in Section 1). In particular, we explore the Bayesian paradigm using a censored likelihood (Sisson and Coles 2003), based on Eq. 2.7. Specifically, let  $X_1, \dots, X_n$  be independent and identically distributed random variables with distribution function  $F$  on  $\mathbb{R}_+$ , where  $F \in \mathcal{D}(G)$ . First we define a high threshold  $t$ . Let  $X_{1,n} \leq X_{2,n} \leq \dots \leq X_{n,n}$  be the  $n$  order statistics and  $F_n$  be the empirical distribution function. The threshold may then be defined as  $T = X_{k,n}$  for large  $k$  such that  $1 - F_n(X_{k,n})$  is close to one, for instance  $F_n(X_{k,n}) = k/n = 0.10, 0.05, 0.01$ . Next, let  $\mathbf{y}_{1:n} = (y_1, \dots, y_n)$  be a realisation of  $(X_1, \dots, X_n)$  and  $t$  be the corresponding threshold. Then on the basis of the approximation in Eq. 2.7 we define the censored likelihood function

$$\mathcal{L}(\mathbf{y}_{1:n}; \boldsymbol{\theta}) = \prod_{i=1}^n \mathcal{L}(y_i; \boldsymbol{\theta}),$$

where each contribution to the likelihood depends on the domain where an observation  $y_i$  falls. That is,

$$\mathcal{L}(y_i; \boldsymbol{\theta}) \propto \begin{cases} G^{k/n}(t; \boldsymbol{\theta}), & \text{if } y_i \leq t, \\ \frac{\partial}{\partial y} G^{k/n}(y; \boldsymbol{\theta})|_{y=y_i}, & \text{if } y_i > t, \end{cases} \quad (3.1)$$

where  $G^{k/n}(y; \boldsymbol{\theta}) \equiv G^{k/n}((y - \mu)/\sigma; \gamma)$  with  $\boldsymbol{\theta} = (\mu, \sigma, \gamma)^\top$ ,  $G^{k/n}((y - \mu)/\sigma; \gamma)$  is given in Eq. 2.7 and

$$\frac{\partial}{\partial t} G^{k/n}(y; \boldsymbol{\theta})|_{y=y_i} = G^{k/n}((y_i - \mu)/\sigma; \gamma) \left( 1 + \frac{\gamma(y_i - \mu)}{\sigma} \right)^{-1/\gamma-1} \frac{1}{\sigma} \frac{k}{n}.$$

Assuming a prior  $\Pi(\boldsymbol{\theta})$  for  $\boldsymbol{\theta} \in \boldsymbol{\Theta} \subseteq \mathbb{R}_+^3$ , we draw samples from the resulting posterior distribution using an adaptive Metropolis-Hastings algorithm. Specifically, we directly apply the adaptive (Gaussian) random-walk Metropolis-Hastings (RWMH) algorithm discussed in Garthwaite et al. (2016). The current state of the chain  $\boldsymbol{\theta}^{(j)}$  at time  $j$  is updated by proposing a draw  $\boldsymbol{\theta}' \sim h(\boldsymbol{\theta}|\boldsymbol{\theta}^{(j)}) = \phi_3(\boldsymbol{\theta}^{(j)}, \tau^{(j)} \Sigma^{(j)})$  where  $\phi_d(\mathbf{a}, A)$  denotes a  $d$ -dimensional Gaussian density function with mean  $\mathbf{a}$  and covariance matrix  $A$ . Because  $h(\boldsymbol{\theta}'|\boldsymbol{\theta}) = h(\boldsymbol{\theta}|\boldsymbol{\theta}')$  is symmetric the acceptance probability of setting  $\boldsymbol{\theta}^{(j+1)} = \boldsymbol{\theta}'$  reduces to

$$\pi^{(j)} = \min \left( \frac{\mathcal{L}(\mathbf{y}_{1:n}; \boldsymbol{\theta}') \Pi(\boldsymbol{\theta}')}{\mathcal{L}(\mathbf{y}_{1:n}; \boldsymbol{\theta}^{(j)}) \Pi(\boldsymbol{\theta}^{(j)})}, 1 \right),$$



otherwise  $\boldsymbol{\theta}^{(j+1)} = \boldsymbol{\theta}^{(j)}$ . Following Haario et al. (2001), the proposal covariance matrix  $\Sigma^{(j)}$  is specified as

$$\Sigma^{(j+1)} = \begin{cases} (1 + [\tau^{(j)}]^2/j)\mathbf{I}_3, & j \leq 100 \\ \frac{1}{j-1} \sum_{k=1}^j (\boldsymbol{\theta}^{(k)} - \bar{\boldsymbol{\theta}}^{(j)})(\boldsymbol{\theta}^{(k)} - \bar{\boldsymbol{\theta}}^{(j)})^\top + ([\tau^{(j)}]^2/j)\mathbf{I}_3, & j > 100, \end{cases} \quad (3.2)$$

where  $\mathbf{I}_d$  is the  $d$ -dimensional identity matrix,  $\bar{\boldsymbol{\theta}}^{(j)} = j^{-1}(\boldsymbol{\theta}^{(1)} + \dots + \boldsymbol{\theta}^{(j)})$ , and  $\tau^{(j)} > 0$  is a scaling parameter that affects the acceptance rate of proposal parameter values. Following Garthwaite et al. (2016) we adaptively update  $\tau$  using a Robbins-Monro process so that

$$\log \tau^{(j+1)} = \log \tau^{(j)} + c(\pi^{(j)} - \pi^*), \quad (3.3)$$

where  $c = (2\pi)^{1/2} \exp(\zeta_0^2/2)/(2\zeta_0)$  is a steplength constant,  $\zeta_0 = -1/\Phi(\pi^*/2)$ , and where  $\Phi$  is the univariate standard Gaussian distribution function. The parameter  $\pi^*$  is the desired overall sampler acceptance probability, here specified as  $\pi^* = 0.234$  following Roberts et al. (1997). This algorithm is summarised in Step 1 of Algorithm 1. Proposed Gaussian updates for the GEV scale parameter are performed on  $\log \sigma$ . See Garthwaite et al. (2016) for further details.

### 3.2 The bivariate case

Inference for the extreme quantile region  $\tilde{\mathcal{Q}}_n$  in Eq. 2.13 requires estimation of the two sets of marginal parameters  $\boldsymbol{\theta}_i = (\mu_i, \sigma_i, \gamma_i)$ ,  $i = 1, 2$ , together with the basic set  $\mathcal{S}$  and its measure  $\nu(\mathcal{S})$ . In particular, the estimation of the basic set  $\mathcal{S}$  and its measure  $\nu(\mathcal{S})$  in Eq. 2.12 requires estimation of the angular density  $h$ . We extend the Bayesian procedure of the univariate case (Section 3.1) to the bivariate setting to simultaneously estimate the marginal parameters and the angular density. This framework utilises the censored likelihood based on the bivariate max-stable distribution in Eq. 2.8.

Specifically, on the basis of the marginal domain of attraction (see Eq. 2.1) and the approximation in Eq. 2.3 for each marginal distribution, we define the transformations

$$z_i := z_i(\cdot) \equiv z_i(\cdot; \boldsymbol{\theta}) = \frac{k}{n} \left( 1 + \gamma_i \frac{\cdot - \mu_i}{\sigma_i} \right)_+^{-1/\gamma_i}, \quad i = 1, 2, \quad (3.4)$$

where  $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ . From the bivariate domain of attraction (2.8) and the max-stability property (de Haan and Ferreira 2006, Ch 6), for large  $\mathbf{y}$ , where  $\mathbf{y} = \mathbf{a}_n \mathbf{x} + \mathbf{b}_n$ , we have

$$F(\mathbf{y}) \approx G(\mathbf{y}; \boldsymbol{\vartheta}^*),$$

where  $\boldsymbol{\vartheta}^* = (\boldsymbol{\theta}^*, h)$ ,  $\boldsymbol{\theta}^* = (\boldsymbol{\theta}_1^*, \boldsymbol{\theta}_2^*)$ ,  $\boldsymbol{\theta}_i^* = (\mu_i^*, \sigma_i^*, \gamma_i)$ ,  $i = 1, 2$ , and where  $h$  is the angular density. Replicating the arguments involving the Pareto distribution for deriving the approximations in Eqs. 2.6 and 2.7 for each margin, then for  $\mathbf{y} \geq \mathbf{t}$ , where  $\mathbf{t} = (t_1, t_2)$  is a large threshold, we obtain the further approximation

$$F(\mathbf{y}) \approx \exp(-L(\mathbf{z}; \boldsymbol{\vartheta})) =: \tilde{G}(\mathbf{z}; \boldsymbol{\vartheta}), \quad \mathbf{y} \geq \mathbf{t}, \quad (3.5)$$

where  $\mathbf{z} = (z_1, z_2)$  with  $z_i = z_i(y_i; \boldsymbol{\theta})$  given by Eq. 3.4 and where  $\boldsymbol{\vartheta} = (\boldsymbol{\theta}, h)$ .

Following Marcon et al. (2016) we model the angular density using Bernstein polynomials. Writing  $L(\mathbf{z}) = (z_1 + z_2)A(v)$  with  $v = z_2/(z_1 + z_2)$ , we model Pickands dependence function  $A(v)$  through a Bernstein polynomial of degree  $\kappa = 0, 1, \dots$  as

$$A_\kappa(v; \boldsymbol{\beta}_\kappa) = \frac{1}{\kappa + 1} \sum_{j=0}^{\kappa} \beta_j \text{Be}(v; j + 1, \kappa - j + 1), \quad (3.6)$$

where  $\boldsymbol{\beta}_\kappa = (\beta_1, \dots, \beta_\kappa)^\top$  is a parameter vector satisfying suitable conditions so that  $A_\kappa$  in Eq. 3.6 defines a valid Pickands dependence function (Marcon et al. 2016, Section 3.1), and  $\text{Be}(\cdot; a, b)$  is the beta density function with parameters  $a > 0$  and  $b > 0$ . Modelling Pickands dependence function with a polynomial in Bernstein form is equivalent to modelling the angular distribution with a Bernstein polynomial. The corresponding angular density in Bernstein form is then

$$h_{\kappa-1}(w; \boldsymbol{\eta}_\kappa) = \sum_{j=0}^{\kappa-2} (\eta_{j+1} - \eta_j) \text{Be}(w; j + 1, \kappa - j - 1), \quad (3.7)$$

where  $w \in [0, 1]$  and the elements of  $\boldsymbol{\eta}_\kappa = (\eta_0, \dots, \eta_{\kappa-2})^\top$  must satisfy suitable conditions so that  $h_{\kappa-1}$  in Eq. 3.7 is a valid angular density (Marcon et al. 2016, Section 3.1). The vectors of coefficients  $\boldsymbol{\beta}_\kappa$  and  $\boldsymbol{\eta}_\kappa$  are related via a one-to-one relationship (Marcon et al. 2016, Proposition 3.2).

Let  $(\mathbf{y}_1, \dots, \mathbf{y}_n)$  be a sample of independent bivariate observations from a distribution  $F$  on  $\mathbb{R}_+^2$ , for which  $F \in \mathcal{D}(G)$ . Then, exploiting the approximation  $\tilde{G}$  in Eq. 3.5 we construct the censored likelihood function

$$\mathcal{L}(\boldsymbol{\vartheta}) = \prod_{i=1}^n \mathcal{L}(\mathbf{y}_i; \boldsymbol{\vartheta}),$$

where each likelihood contribution depends on the domain where  $\mathbf{y}_i$  falls. Specifically,

$$\mathcal{L}(\mathbf{y}; \boldsymbol{\vartheta}) \propto \begin{cases} \tilde{G}(\mathbf{z}; \boldsymbol{\vartheta})|_{\mathbf{z}=\mathbf{z}(t)}, & \text{if } \mathbf{y} \leq t \\ -\frac{\partial z_1(y_1)}{\partial y_1} \tilde{G}(z_1, z_2; \boldsymbol{\vartheta}) L^{(z_1)}(z_1, z_2; \boldsymbol{\vartheta})|_{z_1=z_1(y_1), z_2=z_2(t_2)}, & \text{if } y_1 > t_1, y_2 \leq t_2 \\ -\frac{\partial z_2(y_2)}{\partial y_2} \tilde{G}(z_1, z_2; \boldsymbol{\vartheta}) L^{(z_2)}(z_1, z_2; \boldsymbol{\vartheta})|_{z_1=z_1(t_1), z_2=z_2(y_2)}, & \text{if } y_1 \leq t_1, y_2 > t_2 \\ \mathcal{J} \cdot \tilde{G}(\mathbf{y}; \boldsymbol{\vartheta}) (L^{(z_1)}(\mathbf{z}; \boldsymbol{\vartheta}) L^{(z_2)}(\mathbf{z}; \boldsymbol{\vartheta}) - L^{(z_1, z_2)}(\mathbf{z}; \boldsymbol{\vartheta})|_{\mathbf{z}=\mathbf{z}(\mathbf{y})}), & \text{if } \mathbf{y} > t, \end{cases} \quad (3.8)$$

where

$$\mathcal{J} = \frac{\partial z_1(y_1)}{\partial y_1} \cdot \frac{\partial z_2(y_2)}{\partial y_2},$$

$$L^{(z_1)}(\mathbf{z}; \boldsymbol{\vartheta}) = \frac{\partial}{\partial z_1} L(\mathbf{z}; \boldsymbol{\vartheta}) = A_\kappa(v; \boldsymbol{\beta}_\kappa) - v A'_\kappa(v; \boldsymbol{\beta}_\kappa),$$

$$L^{(z_2)}(\mathbf{z}; \boldsymbol{\vartheta}) = \frac{\partial}{\partial z_2} L(\mathbf{z}; \boldsymbol{\vartheta}) = A_\kappa(v; \boldsymbol{\beta}_\kappa) + (1 - v) A'_\kappa(v; \boldsymbol{\beta}_\kappa),$$

$$L^{(z_1, z_2)}(\mathbf{z}; \boldsymbol{\vartheta}) = \frac{\partial^2}{\partial z_1 \partial z_2} L(\mathbf{z}; \boldsymbol{\vartheta}) = -\frac{1}{z_1 + z_2} v(1 - v) A''_\kappa(v; \boldsymbol{\beta}_\kappa),$$

and where  $v = z_2/(z_1 + z_2)$ ,  $z_i(\cdot; \theta)$ ,  $i = 1, 2$ , is as in Eq. 3.4,  $A'_\kappa$  and  $A''_\kappa$  are the first and second derivatives of  $A_\kappa$  with respect to  $v$ , and where we re-express the full vector of unknown parameters as  $\vartheta = (\theta_1, \theta_2, \kappa, \beta_\kappa)$ . We infer the marginal GEV parameters and the extremal dependence structure through a semiparametric Bayesian approach. In particular, we combine the inferential scheme for each marginal parameter set  $\theta_1, \theta_2$  described above with the trans-dimensional MCMC scheme for inferring the dependence structure over the unknown number of elements in  $\beta_\kappa$  suggested by Marcon et al. (2016) (see also Antoniano-Villalobos and Walker 2013; Fan and Sisson 2011). The latter takes into account that at each MCMC iteration the dimension of  $\beta_\kappa$  changes with  $\kappa$  and that the size of  $\beta_\kappa$  is potentially infinite (see Marcon et al. 2016, Section 3.2 for details).

While it is more convenient to derive the likelihood function using the representation of extremal dependence through Pickands dependence function (Marcon et al. 2016), it is simpler to define a prior distribution on the angular distribution. In this case, the prior distribution for  $\beta_\kappa$  can be deduced from the prior distribution on  $\eta_\kappa$  by exploiting the one-to-one relationship between the two (Marcon et al. 2016). Accordingly, the MCMC algorithm may be implemented using the likelihood function parametrised through Pickands dependence function and the prior distribution for this model that is induced from the prior on the angular distribution.

Recall that for all  $\kappa \geq 3$ ,  $\eta_\kappa$  must satisfy some suitable constraints so that  $h_{\kappa-1}$  in Eq. 3.7 is a valid angular density. Specifically, the elements of  $\eta_\kappa$  must be a nondecreasing sequence in  $[0, 1]$  such that their sum is equal to  $\kappa/2$  (Marcon et al. 2016, Proposition 3.1). Then, following Marcon et al. (2016), for any fixed  $\kappa \geq 3$ , the prior on  $(\eta_\kappa, \kappa)$ , with  $\eta_\kappa = (\eta_0, \dots, \eta_{\kappa-1})$ , is defined as

$$\Pi(\eta_\kappa, \kappa) = \Pi(\eta_1, \dots, \eta_{\kappa-2} | p_0, p_1, \kappa) \Pi(p_1 | \kappa, p_0) \Pi(p_0) \Pi(\kappa), \quad (3.9)$$

where  $\eta_0 = p_0$  and  $\eta_{\kappa-1} = 1 - p_1$  represent the point masses at the endpoints of the simplex  $(H(\{0\}) \text{ and } H(\{1\}))$ ,  $\Pi(\kappa) = \text{NegBin}(\kappa - 3 | m_{NB}, \sigma_{NB})$  with mean  $m_{NB} > 0$  and variance  $\sigma_{NB} > 0$ ,  $\Pi(p_0) = \text{Unif}(0, 1/2)$  and  $\Pi(p_1 | \kappa, p_0) = \text{Unif}(a, b)$  with  $a = \max\{0, (\kappa - 1)p_0 - \kappa/2 + 1\}$  and  $b = (p_0 - \kappa/2 - 1)/(\kappa - 1)$ . The conditional prior distribution on the remaining parameters is set as

$$\Pi(\eta_1, \dots, \eta_{\kappa-2} | p_0, p_1, \kappa) = \text{Unif}(\mathcal{D}_\eta),$$

where the domain  $\mathcal{D}_\eta$  is a suitable set such that the elements of  $\eta_\kappa$  satisfy the required conditions above specified (see Marcon et al. 2016, Section 3.2 for details). Notice that, although the theory in Section 2.2 assumes models for which the condition  $H(\{0\}) = H(\{1\}) = 0$  holds, our proposed inferential method is capable of estimating point masses at the vertices 0 and 1, if this exists.

Within the trans-dimensional MCMC update, the pair  $(\eta_{\kappa^{(j)}}^{(j)}, \kappa^{(j)})$  at time  $j$  is updated through the proposal distribution

$$q(\eta_\kappa, \kappa | \eta_{\kappa^{(j)}}^{(j)}, \kappa^{(j)}) = q_\eta(\eta_\kappa | \kappa) q_\kappa(\kappa | \kappa^{(j)})$$

where  $q_\eta(\eta_\kappa | \kappa) = \Pi(\eta_\kappa | \kappa)$  is the conditional prior implied by Eq. 3.9, and  $q_\kappa(\kappa | \kappa^{(j)})$  is defined such that if  $\kappa^{(j)} = 3$ , it places mass on  $\kappa = 4$  with probability 1

and if  $\kappa^{(j)} > 3$  it places mass on  $\kappa^{(j)} - 1$  and  $\kappa^{(j)} + 1$  with equal probability. Using  $q_\eta(\eta_\kappa|\kappa) = \Pi(\eta_\kappa|\kappa)$  means that these terms cancel in the between-model acceptance probability, whether implemented under the  $\eta_\kappa$  or  $\beta_\kappa$  parameterisation.

The full MCMC sampler is summarised in Algorithm 1, where  $\mathcal{L}(\theta_1, \theta_2, \kappa, \beta_\kappa)$  indicates the bivariate censored likelihood (3.8). Separate Robbins-Monro RWMH updates are implemented for each set of marginal parameters  $\theta_i, i = 1, 2$ , and the above scheme for the dependence parameter updates. See Marcon et al. (2016, Section 3.2) for further details.

---

**Algorithm 1** Trans-dimensional MCMC scheme.

---

```

1 Initialize: Set  $\vartheta^{(0)} = (\theta_1^{(0)}, \theta_2^{(0)}, \kappa^{(0)}, \beta_{\kappa^{(0)}}^{(0)})$ ,  $\eta_{\kappa^{(0)}}^{(0)}$ ,  $\tau_i^{(0)}$  and  $\Sigma_i^{(0)}$  for  $i = 1, 2$ .
2 for  $j = 0$  to  $M$  do
3   Step 1: Marginal component 1:
4     Draw proposal  $\theta_1' \sim MVN(\theta_1^{(j)}, \tau_1^{(j)} \Sigma_1^{(j)})$ .
5     Compute acceptance probability
6     
$$\pi_1 = \min \left( \frac{\mathcal{L}(\theta_1', \theta_2^{(j)}, \kappa^{(j)}, \beta_{\kappa^{(j)}}^{(j)}) \Pi(\theta_1')}{\mathcal{L}(\theta_1^{(j)}, \theta_2^{(j)}, \kappa^{(j)}, \beta_{\kappa^{(j)}}^{(j)}) \Pi(\theta_1^{(j)})}, 1 \right).$$

7     Draw  $U_1 \sim \mathcal{U}(0, 1)$ . If  $\pi_1 < U_1$  then set  $\theta_1^{(j+1)} = \theta_1'$  else set
8      $\theta_1^{(j+1)} = \theta_1^{(j)}$ .
9     Update  $\Sigma_1^{(j)}$  according to (3.2).
10    Update  $\tau_1^{(j)}$  according to (3.3).
11  Step 2: Marginal component 2:
12    Draw proposal  $\theta_2' \sim MVN(\theta_2^{(j)}, \tau_2^{(j)} \Sigma_2^{(j)})$ .
13    Compute acceptance probability
14    
$$\pi_2 = \min \left( \frac{\mathcal{L}(\theta_1^{(j+1)}, \theta_2', \kappa^{(j)}, \beta_{\kappa^{(j)}}^{(j)}) \Pi(\theta_2')}{\mathcal{L}(\theta_1^{(j+1)}, \theta_2^{(j)}, \kappa^{(j)}, \beta_{\kappa^{(j)}}^{(j)}) \Pi(\theta_2^{(j)})}, 1 \right).$$

15    Draw  $U_2 \sim \mathcal{U}(0, 1)$ . If  $\pi_2 < U_2$  then set  $\theta_2^{(j+1)} = \theta_2'$  else set
16     $\theta_2^{(j+1)} = \theta_2^{(j)}$ .
17    Update  $\Sigma_2^{(j)}$  according to (3.2).
18    Update  $\tau_2^{(j)}$  according to (3.3).
19  Step 3: Dependence structure:
20    Draw proposal  $\kappa' \sim q_\kappa(\kappa|\kappa^{(j)})$  and  $\eta_{\kappa'}' \sim q_\eta(\eta_\kappa|\kappa')$ , and compute  $\beta_{\kappa'}'$ .
21    Set  $c = 1/2$  if  $\kappa^{(j)} = 3$  or  $c = 1$  if  $\kappa^{(j)} > 3$ .
22    Compute acceptance probability
23    
$$\pi_3 = \min \left( c \frac{\Pi(\kappa')}{\Pi(\kappa^{(j)})} \frac{\mathcal{L}(\theta_1^{(j+1)}, \theta_2^{(j+1)}, \kappa', \beta_{\kappa'}')}{\mathcal{L}(\theta_1^{(j+1)}, \theta_2^{(j+1)}, \kappa^{(j)}, \beta_{\kappa^{(j)}}^{(j)})}, 1 \right).$$

24    Draw  $U_3 \sim \mathcal{U}(0, 1)$ . If  $\pi_3 < U_3$  then set  $\kappa^{(j+1)} = \kappa'$ ,  $\eta_{\kappa^{(j+1)}}^{(j+1)} = \eta_{\kappa'}'$ ,
25     $\beta_{\kappa^{(j+1)}}^{(j+1)} = \beta_{\kappa'}'$  else set  $\kappa^{(j+1)} = \kappa^{(j)}$ ,  $\eta_{\kappa^{(j+1)}}^{(j+1)} = \eta_{\kappa^{(j)}}^{(j)}$ ,  $\beta_{\kappa^{(j+1)}}^{(j+1)} = \beta_{\kappa^{(j)}}^{(j)}$ .

```

---

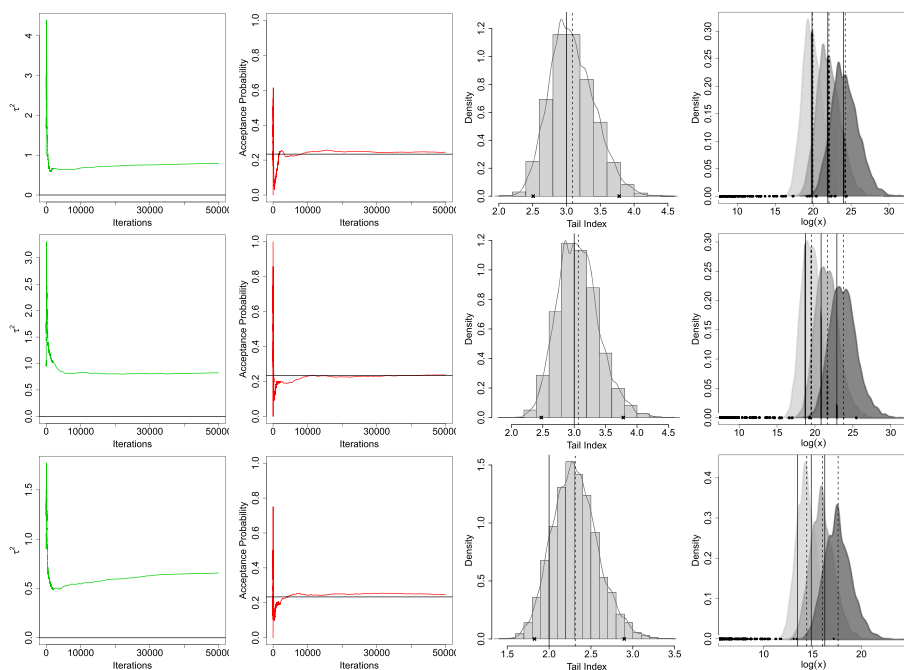
## 4 Simulation experiments

### 4.1 Univariate

We generate  $n = 1500$  observations from each of three distributions: Fréchet with location, scale and shape parameters equals to  $\psi_0 = 3$ ,  $\varsigma_0 = 1$  and  $\xi_0 = 1/3$ , respectively; Half- $t$  with scale and degrees of freedom equal to  $\sigma_0 = 1$  and  $\nu_0 = 1/3$ ; Inverse Gamma with shape parameters equal to  $\eta_0 = 1/2$  and  $\lambda_0 = 1$ . We recall that  $\text{Fr}(y; \psi, \varsigma, \xi) = \exp(-((y - \psi)/\varsigma)^{-\xi})$ , with  $y, \xi > 0$ , is the Fréchet family of distributions with scale  $\varsigma > 0$  and location  $\psi < y$  parameters. These distributions are in the domain of attraction of a GEV distribution with tail indices  $\gamma_0 = 3, 3$  and  $2$ , respectively (Beirlant et al. 2004, Ch. 2). The univariate likelihood function (3.1) is used, censoring observations below the 90-th empirical quantile. We specify a prior distribution for  $\theta$  as a product of uniform prior distributions on the real line for  $\mu$ ,  $\log(\sigma)$  and  $\gamma$ , i.e.  $\Pi(\theta) \equiv \Pi(\mu, \sigma, \gamma) := \Pi(\mu)\Pi(\log(\sigma))\Pi(\gamma)$ . This improper prior distribution, with  $\Pi(\mu, \sigma, \gamma) \propto 1/\sigma$  with  $\sigma > 0$ , leads to a proper posterior distribution (Northrop and Attalides 2016). We run the MCMC sampler for  $M = 50,000$  iterations for each dataset.

Each row in Fig. 1 shows the estimation results for each dataset: Fréchet (top), Half- $t$  (middle) and inverse Gamma (bottom). The columns on the left present trace plots of the scaling parameter  $\tau^2$  (initialised at  $\tau^{(0)} = 1$ ) and the sampler average acceptance probability. Through the Robbins-Monro process both quantities converge rapidly, with the sampler acceptance rate effectively achieving the target (“optimal”) acceptance rate of  $\pi^* = 0.234$  (solid horizontal line) after no more than  $m = 30,000$  iterations, which we remove as sampler burn-in. The centre-right panels illustrate histogram and kernel density estimates of the posterior distribution of the tail index  $\gamma$ , with dashed and solid vertical lines indicating the posterior mean and true value, respectively. The crosses along the horizontal axis show the lower and upper bounds of the estimated 95% credibility interval. In each case the posterior for  $\gamma$  puts most of its mass where the true value lies.

Finally, the rightmost panels display the estimated posterior densities of quantiles corresponding to the exceedance probabilities  $p = 1/750$  (light grey),  $1/1500$  and  $1/3000$  (dark grey), derived from Eq. 2.5. Vertical dashed and solid lines again represent the estimated posterior mean and true quantile values with the latter that are on the log scale: 19.86, 21.94, 24.02 for the Fréchet; 18.73, 20.81, 22.89 for the Half- $t$ ; 13.48, 14.87, 16.25 for the Inverse Gamma. The corresponding estimated central 95% credible intervals are: (17.67, 22.73), (19.41, 25.36), (21.16, 27.95); (17.22, 22.38), (18.93, 25.00), (20.68, 27.60) and (12.58, 16.58), (13.86, 18.58), (15.14, 20.57), respectively. The points on the  $x$ -axis are the upper 5% of the observed dataset. As the sample size is  $n = 1500$ , the largest value is a realisation of an event occurring with probability  $1/1500$ , corresponding to the second investigated quantile. For example, on the log scale, the largest observation from the Fréchet distribution (top row of Fig. 1) is 24.36. This means that the probability of observing an event taking the value 24.36 or greater is less than  $1/750$ , and is more likely to be an event with probability closer to  $1/3000$ .



**Fig. 1** Univariate extreme quantile region results for  $n = 1,500$  Fréchet (top row), Half- $t$  (middle) and inverse-Gamma (bottom) distributed data. Columns illustrate sampler scaling parameter  $\tau^2$  (left) and overall acceptance probability against sampler iteration (centre-left) with target sampler acceptance rate of  $\pi^* = 0.234$  indicated by horizontal line. Centre-right column shows histogram and kernel density estimates of tail index  $\gamma$  after removing  $m = 30,000$  iterations burn-in. Crosses are the lower and upper bounds of the estimated 95% credibility interval. Right column illustrates log-scale posterior densities of quantiles corresponding to the exceedance probabilities  $p = 1/750$  (light grey),  $1/1500$  and  $1/3000$  (dark grey). Posterior mean and true tail index are indicated by dashed and solid lines respectively, observed data indicated by points on the  $x$ -axis

In all cases, the true tail indices and quantiles are always included inside the estimated central 95% credible intervals, confirming the accuracy of our proposed method for estimating such quantities. In Section 2 of the Supplementary Material we investigate the sensitivity of the above results to variations in the experimental setup. In particular we consider a higher censoring threshold (at the 95-th empirical quantile) and alternative prior specifications for  $\mu$ ,  $\sigma$ ,  $\gamma$ . Performance is qualitatively similar to the above in each case.

## 4.2 Bivariate

Estimating and quantifying bivariate extreme quantile regions is more challenging than the univariate setting. Here we examine data simulated from two distributions defined on  $\mathbb{R}_+^2$ : the positive bivariate Cauchy and a positive truncated bivariate- $t$  density. The former has been previously considered in the literature (e.g. Einmahl et al. 2013; Cai et al. 2011), and we consider the second one as an alternative flexible model. Two further examples (the so-called Asymmetric and Clover densities) are

investigated in the Supplementary Material (Section 3). As for the univariate setting, we simulate  $n = 1500$  observations from each density and marginally censor all observations that fall below the corresponding 90-th marginal empirical quantile. We are again interested in estimating quantile regions for events with probability  $p = 1/750$ ,  $1/1500$  and  $1/3000$ , corresponding to regions that contain very little or no observed data. Details about the selected distributions are:

- **Bivariate Cauchy distribution on  $\mathbb{R}_+^2$ :** The Cauchy probability density function is

$$f(\mathbf{x}) = \frac{2}{\pi (1 + x_1^2 + x_2^2)^{3/2}}, \quad \mathbf{x} \in \mathbb{R}_+^2.$$

The Cauchy distribution is very heavy-tailed with tail indices  $\gamma_1 = \gamma_2 = 1$ , and angular density  $h(w) = 2^{-1} (w^2 + (1 - w)^2)^{-3/2}$ . The angular basic density function is  $q_*(w) = (w^2 + (1 - w)^2)^{1/2}$  and the associated basic set is

$$\mathcal{S} = \left\{ \mathbf{x} \in \mathbb{R}_+^2 : r > (w^2 + (1 - w)^2)^{-1/2}, w \in [0, 1] \right\}.$$

- **Truncated bivariate- $t$  distribution on  $\mathbb{R}_+^2$ :** We consider a truncated two-dimensional Student- $t$  distribution with unit-variances, correlation  $0 \leq \rho \leq 1$  and  $\nu > 0$  degrees of freedom. The bivariate- $t$  probability density function is

$$f(\mathbf{x}) = \frac{t_{2,\nu}(\mathbf{x}; \rho)}{T_{2,\nu}(\mathbf{0}; \rho)}, \quad \mathbf{x} \in \mathbb{R}_+^2,$$

where  $t_{d,\nu}$  denotes the  $d$ -dimensional centred Student- $t$  probability density function with correlation  $0 \leq \rho \leq 1$  and  $\nu > 0$  degrees of freedom,  $T_{2,\nu}$  is the  $d$ -dimensional centred Student- $t$  distribution function, and  $\mathbf{0} = (0, 0)^\top$ . The tail indices are  $\gamma_1 = \gamma_2 = 1/\nu$  and vary with the degree of freedom and the angular density is

$$h(w) = \frac{1}{2\nu w^3} \sqrt{\frac{\nu+1}{1-\rho^2}} \left( \frac{1-w}{w} \right)^{\frac{1-\nu}{\nu}} \frac{t_{1,\nu+1} \left( \sqrt{\frac{\nu+1}{1-\rho^2}} \left( \left( \frac{1-w}{w} \right)^{\frac{1}{\nu}} - \rho \right) \right)}{1 - T_{1,\nu+1} \left( -\rho \sqrt{\frac{\nu+1}{1-\rho^2}} \right)}.$$

The angular basic density is equal to

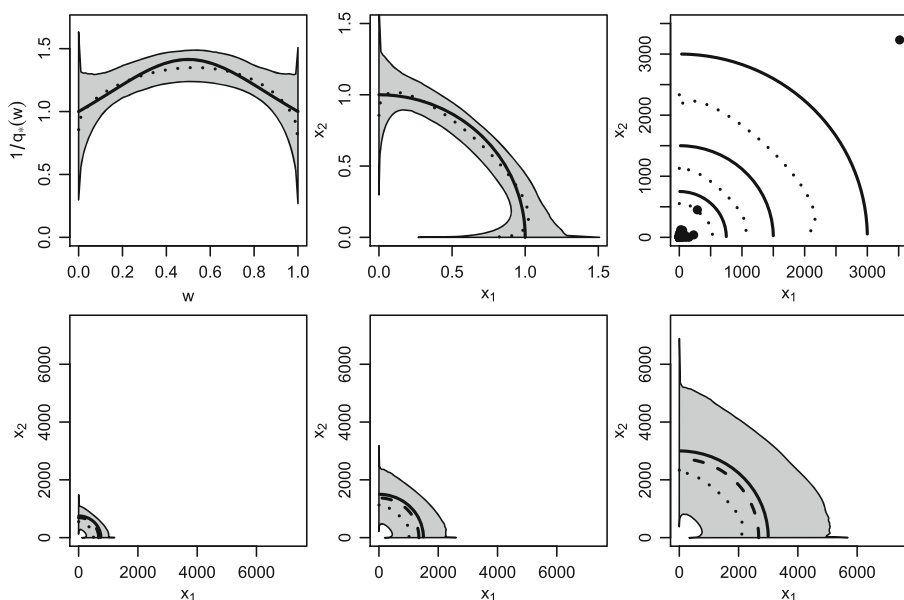
$$q_*(w) = \left( \nu w^{-(1+2/\nu)} \sqrt{\frac{\nu+1}{1-\rho^2}} \frac{t_{1,\nu+1} \left( \sqrt{\frac{\nu+1}{1-\rho^2}} \left( \left( \frac{1-w}{w} \right)^{\frac{1}{\nu}} - \rho \right) \right)}{1 - T_{1,\nu+1} \left( -\rho \sqrt{\frac{\nu+1}{1-\rho^2}} \right)} \right)^{-\frac{\nu}{\nu+2}},$$

and the associated basic set is

$$\mathcal{S} = \left\{ \mathbf{x} \in \mathbb{R}_+^2 : r > \left( \frac{\nu}{w^{1+\frac{2}{\nu}}} \sqrt{\frac{\nu+1}{1-\rho^2}} \frac{t_{1,\nu+1} \left( \sqrt{\frac{\nu+1}{1-\rho^2}} \left( \left( \frac{1-w}{w} \right)^{\frac{1}{\nu}} - \rho \right) \right)}{1 - T_{1,\nu+1} \left( -\rho \sqrt{\frac{\nu+1}{1-\rho^2}} \right)} \right)^{\frac{\nu}{\nu+2}}, w \in [0, 1] \right\}.$$

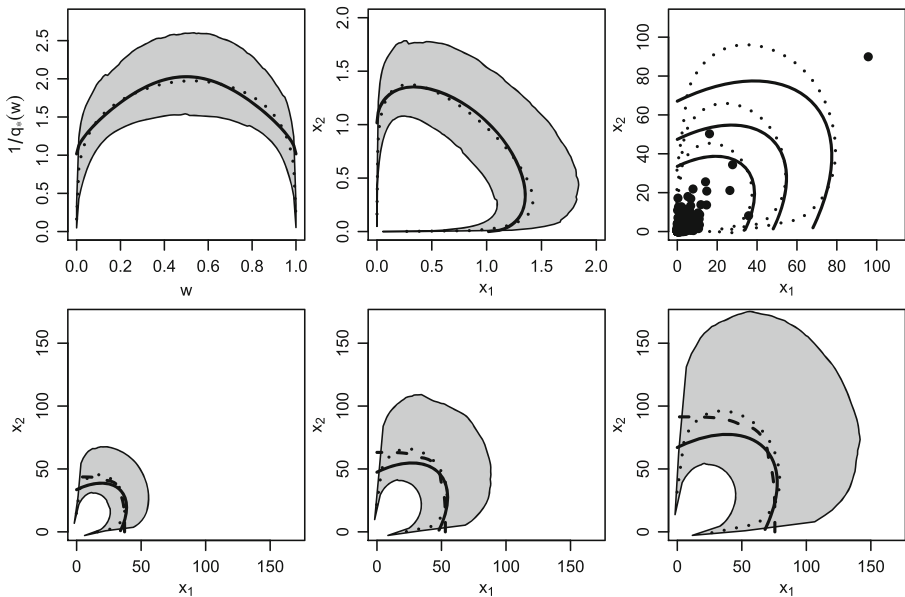
Details concerning the derivation of the angular density are available in [Appendix](#). It is well known that a bivariate- $t$  distribution defined on  $\mathbb{R}^2$  is in the domain of attraction of the Extremal- $t$  model (see e.g. Beranger and Padoan [2015](#)). The angular distribution corresponding to this model places mass on all the subsets of the unit simplex, which in the bivariate case corresponds to the interval  $(0, 1)$  and the vertices  $\{0\}$  and  $\{1\}$ . However, when considering a truncated bivariate- $t$  distribution defined on  $\mathbb{R}_+^2$ , it can be shown that the corresponding angular distribution does not put mass at the endpoints, i.e.  $H(\{0\}) = H(\{1\}) = 0$  (see [Appendix](#)). Note that when  $\nu = 1$  and  $\rho = 0$  the truncated bivariate- $t$  distribution reduces to the bivariate Cauchy on  $\mathbb{R}_+^2$ .

As before, we specify a uniform prior distribution  $\Pi(\mu_i, \sigma_i, \gamma_i)$  for each margin  $i = 1, 2$ . As the dependence structures for the above models are known to be smooth (see solid lines in the top left panels of Figs. 2–3), it is expected they can be well modelled through relatively low degree polynomials. Hence we set the prior distribution for the polynomial degree as  $\Pi(\kappa) = \text{NegBin}(m_{NB} = 3.2, \sigma_{NB} = 4.48)$ . Even though the selected models do not permit mass in the corners of the simplex (i.e.  $H(\{0\}) = H(\{1\}) = 0$ ), in the analysis we still allow for the possibility of non-zero point masses at the endpoints of the simplex by specifying  $\Pi(p_0) = \Pi(p_1) = \text{Unif}(0, 0.1)$ . This prior is slightly different to those in Marcon et al. ([2016](#)) to better represent our prior belief that  $p_0$  and  $p_1$  are likely to be small for these data. See Marcon et al. ([2016](#)) for an analysis of alternative prior specifications.



**Fig. 2** Extreme quantile analysis for the Cauchy distribution. True (solid lines), posterior mean estimate (dotted) and 90% credible regions (grey) for the inverse of the angular basic density (top left), the basic set  $\mathcal{S}$  (top middle), and the quantiles with probability  $p = 1/750, 1/1500$  and  $1/3000$  (bottom left to bottom right). Dashed lines in bottom panels shows the EdHK point estimator. Top right panel illustrates the simulated dataset (points) and true and posterior mean estimated quantiles





**Fig. 3** Extreme quantile analysis for the bivariate- $t$  distribution with  $\nu_0 = 2$  degrees of freedom and correlation  $\rho_0 = 0.5$ . True (solid lines), posterior mean estimate (dotted) and 90% credible regions (grey) for the inverse of the angular basic density (top left), the basic set  $\mathcal{S}$  (top middle), and the quantiles with probability  $p = 1/750, 1/1500$  and  $1/3000$  (bottom left to bottom right). Dashed lines in bottom panels shows the EdHK point estimator. Top right panel illustrates the simulated dataset (points) and true and posterior mean estimated quantiles

We run Algorithm 1 for  $M = 50,000$  iterations and determine the burn-in period  $m$  by visual inspection of trace plots of the marginal scaling parameters  $\tau_1, \tau_2$  and the overall acceptance rates of the marginal ( $\pi_i^* = 0.234, i = 1, 2$ ) and dependence proposals.

For each draw from the posterior distribution, we can construct the angular density  $h_{\kappa-1}(w; \eta)$  using Eq. 3.7, combine this with the marginal shape indices  $\gamma_1, \gamma_2$  into  $q(w)$  Eq. 2.11, and compute the angular basic density  $q_*(w)$  (see Section 2.2 for details). Hence, for each fixed  $w \in (0, 1)$ , we obtain samples from the posterior of the angular basic density  $q_*(w)$ . Its inverse  $q_*^{-1}(w)$  and pointwise central credible intervals are shown in the top-left panels of Figs. 2–3.

The boundary of the basic set  $\mathcal{S}$  corresponds to those points  $(x, y) \in (0, \infty)^2$  such that  $x + y = q_*^{-1}(w)$ , i.e. the points  $(wq_*^{-1}(w), (1-w)q_*^{-1}(w))$ . For fixed  $w \in (0, 1)$ , the posterior distribution of this boundary is available, from which we may calculate the posterior mean and 90% central credible intervals. Computed over all  $w \in (0, 1)$  the (pointwise) posterior mean and credible intervals for the basic set are illustrated in the top-centre panels of Figs. 2–3.

To construct the extreme quantile regions for a given small probability  $p$  (top-right panel and bottom row of Figs. 2–3), consider a point  $w \in (0, 1)$ . As before, we may construct the posterior distribution of the points  $(x, y) = (wq_*^{-1}(w), (1-w)q_*^{-1}(w)) \in \mathbb{R}_+^2$  with angle  $w$  and radial component  $q_*^{-1}(w)$ , as

an estimate of a point on the boundary of  $\mathcal{S}$ . For each  $(x, y)$  point in this posterior, we may compute the quantile  $\tilde{Q}_n$  associated with probability  $p$  via Eq. 2.13, leading to a posterior distribution approximating the  $1 - p$  bivariate quantile level at a particular (estimated) point in  $\mathcal{S}$ . We then use the marginal 0.05 and 0.95 quantiles of this bivariate posterior to define an approximate 90% credible region, and the marginal posterior means to produce a mean estimate. This procedure is repeated for other  $w \in (0, 1)$ .

The top right panel in Figs. 2–3 provides a comparison between the true extreme quantile region (solid line) and the posterior mean estimate (dotted line), for  $p = 1/750, 1/1500$  and  $1/3000$ , with the observed data (points) overlaid. The bottom panels in each figure illustrate each extreme quantile region separately, but with the estimated 90% pointwise credible regions and the point estimates of the quantile regions given in Einmahl et al. (2013) (dotted lines) are included for comparison. We refer to the latter as the EdHK estimator.

Figure 2 illustrates the results for the Bivariate Cauchy distribution. The estimated (pointwise posterior mean) inverse angular basic density  $q_*^{-1}(w)$  (dotted line; top-right panel) describes the behaviour of the true function (solid line) well, and is fully included in the (pointwise) estimated 90% credible regions. The curvature in the dependence structure around  $w = 0.5$  is not as pronounced as the true  $q_*^{-1}(w)$ , which consequently impacts on the estimated basic set and extreme quantile regions. The bottom panels show that the mean extreme quantile levels (dotted) line are close to the true levels (solid line) and provide a similar fit to the EdHK point estimate (dashed) for each probability  $p = 1/750, 1/1500$  and  $1/3000$ . In addition all three curves consistently appear in the centre of the 90% credible regions.

Finally, Fig. 3 presents analysis for the truncated bivariate- $t$  distribution. The dependence structure  $q_*^{-1}(w)$  is very accurately estimated within the interior of the simplex, although towards the endpoints the estimated  $q_*^{-1}(w)$  appears to approach 0 whereas the true value is approximately 1. Producing quantile regions that seem drawn to the origin for low values of either component, may at first appear erroneous. However, inspection of the simulated dataset (top-right panel) reveals that there are no extreme observations that lie along the axes, thereby explaining the behaviour of our estimator. On balance, the estimated posterior means follow the general shape of the true quantile regions and provide similar results than the competitor EdH estimates (dashed lines). The true quantile regions are amply included in the (pointwise) estimated 90% credible regions.

Overall, the proposed methodology is able to accurately estimate both marginal and dependence structure of process extremes. As a point estimator, the posterior mean of the extreme quantile region performs well in recovering the true region, while being responsive to the dependence structure within the data itself, as with the bivariate- $t$  data analysis (Fig. 3). It also seems to perform more consistently than the EdHK estimator (e.g. Figs. 5 and 6 of the Supplementary Material). The credible regions both provide some measure of the uncertainty inherent with low probability events, while also allowing for judgements regarding whether exceptionally high observations can still be considered to belong to events with particular probabilities (e.g. whether the single large outlier in Fig. 3 can be considered a  $1/1500$  or  $1/3000$  probability event).

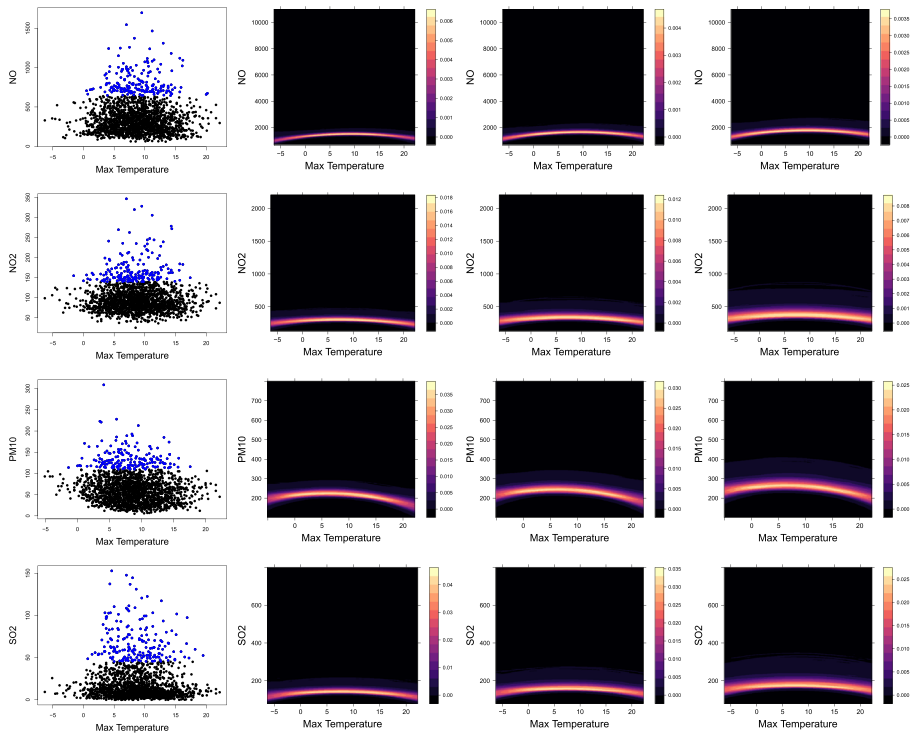
## 5 Analysis of extreme air pollution levels in Milan

Understanding the extremes of air pollutants is of critical importance, especially in the context of climate change (De Sario et al. 2013). In recent work, Martins et al. (2017) used standard univariate extreme value theory to compare the air quality between the two largest urban regions in Brazil, Heffernan and Tawn (2004) and Beranger and Padoan (2015) estimated the extremal dependence between pollutants in Leeds, U.K., Coles and Pan (1996) studied the extreme temporal behaviour of airborne NO<sub>2</sub> particles in Milan between 1984–1994, and Vettori et al. (2018) performed a spatial analysis of air pollution in Los Angeles, U.S.A.

Here we study extreme air pollution levels recorded in Milan, Italy, over the winter period October 31st–February 27/28th, between December 31st 2001 and December 30th 2017. We examine the daily mean level of PM<sub>10</sub> and daily maximum levels of NO, NO<sub>2</sub> and SO<sub>2</sub>. These data were also considered by Falk et al. (2019, Table 3) although there the objective was estimating joint exceedance probabilities (with excesses in all variables).

Here, the aim is to estimate bivariate extreme quantile regions given by a high concentration of the two pollutants and with a small probability of occurrence. This analysis is important for understanding the long term impact of air quality on health. When modelling univariate extremes it is common to express marginal parameters via a regression model which may be functions of space or other covariates (Padoan et al. 2010). For air pollution in particular, evidence is overwhelming that interactions between the pollution and temperatures cannot be disregarded (De Sario et al. 2013; Cheng and Kan 2012; Katsouyanni et al. 1993; Roberts 2004). The leftmost panels of Fig. 4 show scatterplots of each pollutant against the daily maximum temperature. All four indicate a possible quadratic relationship between pollutant and maximum temperature, and so we write the  $i$ -th marginal mean  $\mu_i = \beta_{0,i} + \beta_{1,i}z + \beta_{2,i}z^2$  as a quadratic function of the maximum temperature  $z$ . Other covariates could have been included if they were available, and regressions on  $\sigma_i$  and  $\gamma_i$  constructed, although we did not pursue that here. For observations that fall below the threshold (black points in left panels of Fig. 4), as these observations are censored it could be argued that the level of the covariates at the threshold should be used when evaluating the likelihood contribution. This would beneficially reduce computational costs. However, as the covariates are still available for censored observations, they still provide valuable information to estimate the regression coefficients.

Using mean residual life plots (e.g. Beirlant et al. 2004, Section 1.2.2) we set the marginal 90-th empirical quantile as the threshold (points above the threshold are blue in Fig. 4), which results in  $t = 642.9$  ( $n = 1796$ ) for NO,  $t = 139.5$  ( $n = 1799$ ) for NO<sub>2</sub>,  $t = 108$  ( $n = 1779$ ) for PM<sub>10</sub> and  $t = 45.2$  ( $n = 1809$ ) for SO<sub>2</sub>. These thresholds are comparable with those in Falk et al. (2019, Table 3). We specify all marginal prior distributions  $\Pi(\mu_i, \sigma_i, \gamma_i) \propto 1/\sigma_i$ , for  $i = 1, \dots, 4$  as in Section 4.1, and dependence parameter prior distributions as in Section 3.2 with  $\Pi(\kappa) = \text{NegBin}(\kappa = 3 | m_{NB} = 6, \sigma_{NB} = 8)$  to allow for higher degree polynomial modelling of  $h_{\kappa-1}(w; \eta)$  if required. For the univariate analysis we implement an MCMC sampler with 300k iterations and retain the final 50k iterations for inference; for the multivariate analysis we retain the final 20k iterations from a chain of length



**Fig. 4** Univariate analyses of (top row to bottom row) NO, NO<sub>2</sub>, PM<sub>10</sub> and SO<sub>2</sub> data relative to the maximum daily temperatures ( $x$ -axes). Left column: Scatterplots of each pollutant versus maximum daily temperature, with observations above the threshold shown in blue. Second to fourth columns: Image plots of the univariate distribution of quantiles associated with probability  $p = 1/600$ ,  $1/1200$  and  $1/2400$  as a function of maximum daily temperature

50k. The univariate sampler ran in under 1 hour on a single-core 2.2 GHz Intel Core i7 processor on a MacBook Pro; the multivariate sampler took around 2.5 hours.

The image plots in Fig. 4 illustrate the (univariate) posterior distribution of the quantiles associated with the probabilities  $p = 1/600$ ,  $1/1200$  and  $1/2400$  (left to right) as a function of the maximum daily temperature, which correspond to pollutant levels that would be expected to be reached once every 5, 10 and 20 winters. Restricting our interpretation to the range of observed temperature, we note that extreme PM<sub>10</sub> quantiles (third row) appear higher for lower temperatures rather than high. Indeed, the main sources of PM<sub>10</sub> pollution include combustion engines (both diesel and petrol) and combustion for energy production in households. Accordingly, when maximum temperatures are low, one may expect household heating systems to work at higher capacity and an increase in the use of cars rather than less energy consumptive transport methods such as walking or cycling.

Extreme quantile behaviour for NO<sub>2</sub> (second row) as function of maximum temperature appears similar to PM<sub>10</sub> although the larger range of the quantiles ( $y$ -axis) reduces the visual curvature. This behaviour can be explained by the fact that nitrogen


dioxide, mainly emitted by power generation, industrial and traffic sources, is an important constituent of particulate matter. Sulfur dioxide (fourth row) forms either naturally (decomposition or combustion of organic matter) or due to human activity (smelting sulfur-containing mineral ores). Accordingly, it is not surprising to observe decreasing extreme  $\text{SO}_2$  quantiles across temperature levels (fourth row) as maximum temperature moves from 0 to 20°C. The decreasing quantile behaviour as temperature gets cold might be due to a reduction in human activity. The posterior mean of the tail indices  $\gamma_{\text{NO}_2}$ ,  $\gamma_{\text{NO}_2}$ ,  $\gamma_{\text{PM}_{10}}$  and  $\gamma_{\text{SO}_2}$  and their corresponding 90% credible interval are 0.050 (0.003, 0.138), 0.186 (0.055, 0.337), 0.054 (0.004, 0.148) and 0.064 (0.005, 0.167) respectively.

To date, the link between extreme levels of multiple pollutants and human health has not been well explored, although a multi-pollutant approach to assess the health risks associated with air pollution has been emerging (e.g. Dominici et al. 2010; Wesson et al., 2010). In the following we restrict our attention to the three pollutants with the largest tail indices:  $\text{NO}_2$ ,  $\text{SO}_2$  and  $\text{PM}_{10}$ .

Figure 6 shows information regarding extremal dependence between the three pollutant pairs:  $\text{NO}_2/\text{SO}_2$ ,  $\text{NO}_2/\text{PM}_{10}$  and  $\text{SO}_2/\text{PM}_{10}$ . The estimated basic sets demonstrate weak dependence between the pairs  $\text{NO}_2/\text{SO}_2$  and  $\text{SO}_2/\text{PM}_{10}$  and stronger dependence between  $\text{NO}_2$  and  $\text{PM}_{10}$ . Further, from the pairwise scatterplots it is apparent that neither the coldest (blue dots) nor the warmest (red dots) daily maximum temperatures appear to induce the largest pollutant levels. As discussed above,  $\text{NO}_2$  is a key constituent of  $\text{PM}_{10}$  and so it is realistic to expect the observed strong dependence between these pollutants. Similarly, the sources of  $\text{PM}_{10}$  (combustion engines) and  $\text{SO}_2$  (natural or smelting mineral ores) differ, explaining the independence between the extremes of these pollutants to some extent.


Figure 5 illustrates extreme quantiles regions corresponding to events that would expect to occur once every 5, 10 and 20 winters (left to right panels) when the daily maximum temperature is fixed to the minimum, median and maximum observed daily maximum temperatures (blue to red shading). These values are (−6.3 °C, 8.6 °C, 22.2 °C), (−5.1 °C, 8.6 °C, 22.2 °C) and (−5.1 °C, 8.7 °C, 22.2 °C) for the pairs  $\text{NO}_2/\text{SO}_2$ ,  $\text{NO}_2/\text{PM}_{10}$  and  $\text{SO}_2/\text{PM}_{10}$  respectively. The top panels of Fig. 5 exhibit a small quadratic influence of maximum daily temperature on the joint levels of  $\text{NO}_2$  and  $\text{SO}_2$ , and similarly for  $\text{SO}_2$  and  $\text{PM}_{10}$  (bottom row). It is apparent that extreme levels of  $\text{SO}_2$  are reduced with warmer weather. This phenomena can be loosely understood by the fact that sulfur dioxide is an aerosol which cools the planet by reflecting some of the sun's energy back into space. As such, large levels of  $\text{SO}_2$  are more likely to be associated with cold temperatures. Overall the estimated extreme quantile regions capture the behaviour of the data well, with the  $\text{NO}_2/\text{PM}_{10}$  pair exhibiting the strongest level of dependence. As expected, as the probability  $p$  of an event decreases the quantile regions become larger and reach higher pollutant levels. Despite having observations for 17 consecutive winters, our method is able to provide quantile levels and credible regions for events with probability of occurrence  $p$  at arbitrarily low levels.

In terms of the pollutant thresholds suggested by European emission regulation, the top-left panel of Fig. 5 shows that when the the daily maximum temperature is fixed to the maximum, we expect that  $\text{NO}_2$  individually exceeds approximatively



**Fig. 5** Posterior mean estimates (solid line) and 90% credible regions for the extreme bivariate quantiles associated with the probabilities  $p = 1/600$ ,  $1/1200$  and  $1/2400$  (left to right) for three maximum daily temperature levels: minimum temperature = blue, median temperature = purple, maximum temperature = red

$200 \mu\text{g}/\text{m}^3$  (i.e. the short-term limit threshold) once every 5 winters ( $p = 1/600$ ) on average. However, this event is as equally probable as the joint event that both  $\text{NO}_2$  and  $\text{PM}_{10}$  each exceed approximately  $200 \mu\text{g}/\text{m}^3$ . Such a concentration for  $\text{PM}_{10}$  is four times higher than the individual limit threshold for this pollutant (see Section 1). Similarly, for the bottom-left panel of Fig. 5, when the daily maximum temperature is fixed to the median temperature, we expect that  $\text{PM}_{10}$  individually exceeds approximately  $200 \mu\text{g}/\text{m}^3$  on average once every 5 winters. However, this event is as equally probable as the joint event that both  $\text{SO}_2$  and  $\text{PM}_{10}$  each exceed approximately  $200 \mu\text{g}/\text{m}^3$ . This concentration for  $\text{SO}_2$  is almost twice the individual limit threshold for this pollutant (see Section 1). These results indicate both the strong dependence between harmful pollutants at extreme levels in Milan, and that



**Fig. 6** Posterior mean estimate (dotted line) and 90% credible regions (grey) for the inverse of the angular basic density (left column), the basic set  $S$  (middle column), and and observed data (right column) with temperature dependent data colouring (from cold = blue to warm = red). Top to bottom: The pairs of pollutants shown are  $\text{NO}_2/\text{SO}_2$ ,  $\text{NO}_2/\text{PM}_{10}$  and  $\text{SO}_2/\text{PM}_{10}$

limit threshold alerts for poor air quality in this city are likely to be issued for multiple pollutants simultaneously.

## 6 Discussion

We have presented a new method for estimating extreme quantile regions that is responsive to varying levels of extremal dependence, and comes with natural measures of uncertainty, both for model parameters and extreme quantile regions, under the Bayesian paradigm. The method was able to outperform the existing (EdHK) approach of Einmahl et al. (2013) which does not provide any measure of uncertainty. This methodology provides a useful general tool for long-term analysis of the air

pollution at the extreme level. We explored this in Section 5 for assessing and quantifying the health risks associated with multiple extreme pollutant exposures (Dominici et al. 2010; Wesson et al. 2010).

The methods developed in this paper have been incorporated into accessible functions in the R Package `ExtremalDep`. This package is freely available from the CRAN repository, see <https://CRAN.R-project.org/package=ExtremalDep>. The R code for the simulations and the multivariate analysis of real air pollution data is available online at the address [https://www.borisberanger.com/zip/BPS\\_2019.zip](https://www.borisberanger.com/zip/BPS_2019.zip).

**Acknowledgements** The authors are grateful to Andrea Krajina for sharing the code for the frequentist estimation of bivariate extreme quantile regions and her valuable suggestions and help. SAP is supported by the Bocconi Institute for Data Science and Analytics (BIDSA), Italy and PRIN 2015 research grant *Modern Bayesian Nonparametric Methods*. SAS and BB are supported by the Australian Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS; CE140100049) and the Australian Research Council Discovery Project scheme (FT170100079). The authors are also indebted to the Associate Editor and two anonymous reviewers for their careful reading of the manuscript and their constructive remarks

## Appendix: The Extremal- $t$ model with restriction to the positive reals

Consider a student- $t$  distribution restricted to  $(0, \infty)$ . Using Beirlant et al. (2004, p.59) the norming constants required in Eq. 2.1 are

$$a_n = n^{1/v} \left( \frac{2\Gamma(\frac{v+1}{2})v^{\frac{v}{2}-1}}{\sqrt{\pi}\Gamma(\frac{v}{2})} \right)^{1/v} \quad \text{and} \quad b_n = 0.$$

Applying the conditional tail dependence function framework of Nikoloulopoulos et al. (2009), the exponent function can be written as:

$$V(x, y) = \lim_{n \rightarrow \infty} x^{-v} \mathbb{P}(Z_2 \leq a_n y | Z_1 \leq a_n x) + y^{-v} \mathbb{P}(Z_1 \leq a_n x | Z_2 \leq a_n y),$$

where  $(Z_1, Z_2)^\top$  follows a centred bivariate- $t$  distribution on  $(0, \infty)^2$  with unit variance, correlation  $\rho$  and degree of freedom  $v$ . The conditional distribution of  $Z_2 | Z_1 = z_1$  is a truncated  $t$  distribution on  $(0, \infty)$  with mean  $\rho z_1$ , variance  $(v + z_1^2)(1 - \rho^2)/(v + 1)$  and  $v + 1$  degrees of freedom. As a consequence, we obtain

$$\mathbb{P}(Z_2 \leq a_n y | Z_1 \leq a_n x) = \frac{T_{1,v+1} \left( \sqrt{\frac{v+1}{1-\rho^2}} \frac{a_n(y-\rho x)}{\sqrt{v+a_n^2 x^2}} \right) - T_{1,v+1} \left( -\rho \sqrt{\frac{v+1}{1-\rho^2}} \frac{a_n x}{\sqrt{v+a_n^2 x^2}} \right)}{1 - T_{1,v+1} \left( -\rho \sqrt{\frac{v+1}{1-\rho^2}} \frac{a_n x}{\sqrt{v+a_n^2 x^2}} \right)},$$

and

$$\lim_{n \rightarrow \infty} \mathbb{P}(Z_2 \leq a_n y | Z_1 \leq a_n x) = \frac{T_{1,v+1} \left( \sqrt{\frac{v+1}{1-\rho^2}} \left( \frac{y}{x} - \rho \right) \right) - T_{1,v+1} \left( -\rho \sqrt{\frac{v+1}{1-\rho^2}} \right)}{1 - T_{1,v+1} \left( -\rho \sqrt{\frac{v+1}{1-\rho^2}} \right)}.$$



Identical calculations can be applied to the second term in the exponent function. Transforming the margins to unit-Fréchet margins allows expression of the exponent function as

$$V(x, y) = \frac{1}{1 - T_{1, v+1} \left( -\rho \sqrt{\frac{v+1}{1-\rho^2}} \right)} \times \left\{ \frac{1}{x} \left[ T_{1, v+1} \left( \sqrt{\frac{v+1}{1-\rho^2}} \left( \frac{y}{x} - \rho \right) \right) - T_{1, v+1} \left( -\rho \sqrt{\frac{v+1}{1-\rho^2}} \right) \right] + \frac{1}{y} \left[ T_{1, v+1} \left( \sqrt{\frac{v+1}{1-\rho^2}} \left( \frac{x}{y} - \rho \right) \right) - T_{1, v+1} \left( -\rho \sqrt{\frac{v+1}{1-\rho^2}} \right) \right] \right\}.$$

It is easy to verify that  $\lim_{y \rightarrow 0} \partial/\partial x V(x, y) = 0$  and  $\lim_{x \rightarrow 0} \partial/\partial y V(x, y) = 0$  which implies  $H(\{0\}) = H(\{1\}) = 0$ . Finally, note that due to the form of  $V(x, y)$ , taking the double partial derivative with respect to  $x$  and  $y$  is equivalent to the double partial derivative of the exponent function of the Extremal- $t$  model multiplied by a scaling term. Hence the angular density on the interior of the 2-dimensional unit simplex is as given.

## References

- Antoniano-Villalobos, I., Walker, S.G.: Bayesian nonparametric inference for the power likelihood. *J. Comput. Graph. Stat.* **22**, 801–813 (2013)
- Beirlant, J., Goegebeur, Y., Segers, J., Teugels, J.: *Statistics of Extremes: Theory and Applications*. Wiley, Chichester (2004)
- Beranger, B., Padoan, S.A.: Extreme dependence models. In: *Extreme Value Modeling and Risk Analysis*, pp. 325–352. Chapman and Hall/CRC (2015)
- Bienvenue, A., Robert, C.Y.: Likelihood inference for multivariate extreme value distributions whose spectral vectors have known conditional distributions. *Scand. J. Stat.* **44**(1), 130–149 (2017)
- Bremnes, J.B.: Probabilistic forecasts of precipitation in terms of quantiles using nwp model output. *Mon. Weather. Rev.* **132**(1), 338–347 (2004a)
- Bremnes, J.B.: Probabilistic wind power forecasts using local quantile regression. *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology* **7**(1), 47–54 (2004b)
- Cai, J., Einmahl, J.H.J., de Haan, L.: Estimation of extreme risk regions under multivariate regular variation. *Ann. Stat.* **39**, 1803–1826 (2011)
- Cheng, Y., Kan, H.: Effect of the interaction between outdoor air pollution and extreme temperature on daily mortality in Shanghai, China. *J. Epidemiol.* **22**(1), 28–36 (2012)
- Clapp, L.J., Jenkin, M.E.: Analysis of the relationship between ambient levels of o3, no2 and no as a function of nox in the UK. *Atmos. Environ.* **35**(36), 6391–6405 (2001)
- Coles, S.G., Pan, F.: The analysis of extreme pollution levels: A case study. *J. Appl. Stat.* **23**, 333–348 (1996)
- Cooley, D., Thibaud, E., Castillo, F., Wehner, M.F.: A nonparametric method for producing isolines of bivariate exceedance probabilities. [arXiv:1710.05248](https://arxiv.org/abs/1710.05248) (2017)
- Dahlhaus, R.: Graphical interaction models for multivariate time series. *Metrika* **51**(2), 157–172 (2000)
- de Haan, L., Ferreira, A.: *Extreme Value Theory: An Introduction*. Springer, New York (2006)
- De Sario, M., Katsouyanni, K., Michelozzi, P.: Climate change, extreme weather events, air pollution and respiratory health in Europe. *Europ. Resp. J.* **42**(3), 826–843 (2013)
- Dominici, F., Peng, R.D., Barr, C.D., Bell, M.L.: Protecting human health from air pollution: shifting from a single-pollutant to a multi-pollutant approach. *Epidemiology (Cambridge Mass.)* **21**(2), 187 (2010)

- Einmahl, J.H.J., de Haan, L., Krajina, A.: Estimating extreme bivariate quantile regions. *Extremes* **16**, 121–145 (2013)
- Falk, M., Padoan, S.A., Wisheckel, F.: Generalized Pareto copulas: A key to multivariate extremes. *Journal of Multivariate Analysis*. To Appear (2019)
- Fan, Y., Sisson, S.A.: Reversible jump Markov chain Monte Carlo. In: Brooks, S.P., Gelman, A., Jones, G., Meng, X.-L., Fan, Y., Sisson, S.A. (eds.), pp. 69–94. Chapman & Hall/CRC (2011)
- Friederichs, P., Thorarindottir, T.L.: Forecast verification for extreme value distributions with an application to probabilistic peak wind prediction. *Environmetrics* **23**(7), 579–594 (2012)
- Garthwaite, P.H., Fan, Y., Sisson, S.A.: Adaptive optimal scaling of Metropolis–Hastings algorithms using the Robbins–Monro process. *Commun. Stat.: Theory Methods* **45**(17), 5098–5111 (2016)
- Guerreiro, C., Ortiz, A.G., de Leeuw, F., Viana, M., Horálek, J.: Air quality in Europe – 2016 Report. Publications Office of the European Union (2016)
- Haario, H., Saksman, E., Tamminen, J.: An adaptive Metropolis algorithm. *Bernoulli* **7**, 223–242 (2001)
- Heffernan, J.E., Tawn, J.A.: A conditional approach for multivariate extreme values. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **66**(3), 497–546 (2004). With discussions and reply by the authors
- Huser, R., Davison, A.C., Genton, M.G.: Likelihood estimators for multivariate extremes. *Extremes* **19**(1), 79–103 (2016)
- Jagger, T.H., Elsner, J.B.: Modeling tropical cyclone intensity with quantile regression. *Int. J. Climatol.: A J. R. Meteorol. Soc.* **29**(10), 1351–1361 (2009)
- Katsouyanni, D.K., Pantazopoulou, A., Touloumi, G., Tselepidaki, I., Moustris, K., Asimakopoulou, D., Pouloupoulou, G., Trichopoulos, D.: Evidence for interaction between air pollution and high temperature in the causation of excess mortality. *Arch. Environ. Health: Int. J.* **48**(4), 235–242 (1993). PMID: 8357272
- Ledford, A.W., Tawn, J.A.: Statistics for near independence in multivariate extreme values. *Biometrika* **83**, 169–187 (1996)
- Marcon, G., Padoan, S.A., Antoniano-Villalobos, I.: Bayesian inference for the extremal dependence. *Electron. J. Stat.* **10**, 3310–3337 (2016)
- Marcon, G., Naveau, P., Padoan, S.: A semi-parametric stochastic generator for bivariate extreme events. *Statistics* **6**(1), 184–201 (2017)
- Martins, L.D., Wikuats, C.F.H., Capucim, M.N., de Almeida, D.S., da Costa, S.C., Albuquerque, T., Carvalho, V.S.B., de Freitas, E.D., de FÁtima Andrade, M., Martins, J.A.: Extreme value analysis of air pollution data and their comparison between two large urban regions of South America. *Weather Clim. Extremes* **18**, 44–54 (2017)
- Naess, A., Karpa, O.: Statistics of extreme wind speeds and wave heights by the bivariate acer method. *J. Offshore Mech. Arctic Eng.* **137**(2), 021602 (2015)
- Nikolouloupoulos, A.K., Joe, H., Li, H.: Extreme value properties of multivariate  $t$  copulas. *Extremes* **12**(2), 129–148 (2009)
- Northrop, P.J., Attalides, N.: Posterior propriety in Bayesian extreme value analyses using reference priors. *Stat. Sin.* **26**(2), 721–743 (2016)
- Padoan, S.A., Ribatet, M., Sisson, S.A.: Likelihood-based inference for max-stable processes. *J. Am. Stat. Assoc.* **105**, 263–277 (2010)
- Porter, W.C., Heald, C.L., Cooley, D., Russell, B.: Investigating the observed sensitivities of air-quality extremes to meteorological drivers via quantile regression. *Atmos. Chem. Phys.* **15**(18), 10349–10366 (2015)
- Prescott, P., Walden, A.: Maximum likelihood estimation of the parameters of the three-parameter generalized extreme-value distribution from censored samples. *J. Stat. Comput. Simul.* **16**(3–4), 241–250 (1983)
- Roberts, G.O., Gelman, A., Gilks, W.R.: Weak convergence and optimal scaling of random walk Metropolis algorithms. *Ann. Appl. Probab.* **7**(1), 110–120 (1997)
- Roberts, S.: Interactions between particulate air pollution and temperature in air pollution mortality time series studies. *Environ. Res.* **96**(3), 328–337 (2004)
- Sisson, S., Coles, S.: Modelling dependence uncertainty in the extremes of Markov chains. *Extremes* **6**(4), 283–300 (2003). (2005)
- Smith, R.L.: Multivariate threshold methods. In: Galambos, J., Lechner, J., Simiu, E., Smith, R.L. (eds.), pp. 225–248. Kluwer (1994)
- Vettori, S., Huser, R., Genton, M.G.: Bayesian modeling of air pollution extremes using nested multivariate max-stable processes. [arXiv:1804.04588](https://arxiv.org/abs/1804.04588) (2018)

- Wang, H.J., Li, D.: Estimation of extreme conditional quantiles. In: *Extreme Value Modeling and Risk Analysis*, pp. 313–331. Chapman and Hall/CRC (2015)
- Wesson, K., Fann, N., Morris, M., Fox, T., Hubbell, B.: A multi-pollutant, risk-based approach to air quality management: Case study for Detroit. *Atmos. Pollut. Res.* **1**(4), 296–304 (2010)
- World Health Organization: *Air Quality Guidelines: Global Update 2005*, 3rd edn. World Health Organization ISBN 9289021926 (2006)
- Yu, K., Lu, Z., Stander, J.: Quantile regression: applications and current research areas. *J. R. Stat. Soc.: Series D (The Statistician)* **52**(3), 331–350 (2003)

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.