



HAL
open science

Semi-parametric modeling of excesses above high multivariate thresholds with censored data

Anne Sabourin

► **To cite this version:**

Anne Sabourin. Semi-parametric modeling of excesses above high multivariate thresholds with censored data. 2014. hal-01088587

HAL Id: hal-01088587

<https://hal.science/hal-01088587>

Preprint submitted on 28 Nov 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

SEMI-PARAMETRIC MODELING OF EXCESSES ABOVE HIGH MULTIVARIATE THRESHOLDS WITH CENSORED DATA

Anne Sabourin ¹

¹ Institut Mines-Télécom, Télécom ParisTech, CNRS LTCI
37-38, rue Dareau, 75014 Paris, FRANCE
anne.sabourin@telecom-paristech.fr

ABSTRACT. How to include censored data in a statistical analysis is a recurrent issue in statistics. In multivariate extremes, the dependence structure of large observations can be characterized in terms of a non parametric *angular measure*, while marginal excesses above asymptotically large thresholds have a parametric distribution. In this work, a flexible semi-parametric Dirichlet mixture model for angular measures is adapted to the context of censored data and missing components. One major issue is to take into account censoring intervals overlapping the extremal threshold, without knowing whether the corresponding hidden data is actually extreme. Further, the censored likelihood needed for Bayesian inference has no analytic expression. The first issue is tackled using a Poisson process model for extremes, whereas a data augmentation scheme avoids multivariate integration of the Poisson process intensity over both the censored intervals and the failure region above threshold. The implemented MCMC algorithm allows simultaneous estimation of marginal and dependence parameters, so that all sources of uncertainty other than model bias are captured by posterior credible intervals. The method is illustrated on simulated and real data.

Multivariate extremes; censored data; data augmentation; semi-parametric Bayesian inference; MCMC algorithms.

1. INTRODUCTION

Data censoring is a commonly encountered problem in multivariate statistical analysis of extreme values. A ‘censored likelihood’ approach makes it possible to take into account partially extreme data (non concomitant extremes): coordinates that do not exceed some large fixed threshold are simply considered as left-censored. Thus, the possibly misleading information carried by non-extreme
5 coordinates is ignored, only the fact that they are not extreme is considered (Smith (1994); Ledford and Tawn (1996); Smith et al. (1997), see also Thibaud and Opitz (2013) or Huser et al. (2014)). However, there are other situations where the original data is incomplete. For example, one popular way to obtain large sample sizes
10 in environmental sciences in general and in hydrology in particular, is to take into account data reconstructed from archives, which results in a certain amount of left- and right-censored data, and missing data. As an example, what originally motivated this work is a hydrological data set consisting of daily water discharge
15 recorded at four neighboring stations in the region of the Gardons, in the south of France. The extent of systematic recent records is short (a few decades) and varies from one station to another, so that standard inference using only ‘clean’ data is unfeasible (only 3 uncensored multivariate excesses of ‘large’ thresholds - fixed after preliminary uni-variate analysis- are recorded). Historical information

is available, starting from the 17th century, a large part of it being censored: only major floods are recorded, sometimes as an interval data (*e.g.* ‘the water level exceeded the parapet but the Mr. X’s house was spared’). These events are followed by long ‘blank’ periods during which the previous record was not exceeded.
5 Uni-variate analysis for this data set has been carried on by Neppel et al. (2010) but a multivariate analysis of extremes has never been accomplished, largely due to the complexity of the data set, with multiple censoring.

While modeling multivariate extremes is a relatively well marked out path when ‘exact’ (non censored) data are at stake, many fewer options are currently available
10 for the statistician working with censored data. The aim of the present paper is to provide a flexible framework allowing multivariate inference in this context. Here, the focus is on the methodology and the inferential framework is mainly tested on simulated data with a censoring pattern that resembles that of the real data. A detailed analysis of the hydrological data raises other issues, such as,
15 among others, temporal dependence and added value of the most ancient data. These questions are addressed in a separate paper, intended for the hydrological community (Sabourin and Renard, 2014)¹.

Under a standard assumption of multivariate regular variation (see Section 2), the distribution of excesses above large thresholds is characterized by parametric
20 marginal distributions and a non-parametric dependence structure that is independent from threshold. Since the family of admissible dependence structures is, by nature, too large to be fully described by any parametric model, non-parametric estimation has received a great deal of attention in the past few years (Einmahl et al., 2001; Einmahl and Segers, 2009; Guillotte et al., 2011). To the best of
25 my knowledge, the non parametric estimators of the so-called *angular measure* (which characterizes the dependence structure among extremes) are only defined with exact data and their adaptation to censored data is far from straightforward.

For applied purposes, it is common practice to use a parametric dependence model. A widely used one is the Logistic model and its asymmetric and nested ex-
30 tensions (Gumbel, 1960; Coles and Tawn, 1991; Stephenson, 2009, 2003; Fougères et al., 2009). In the logistic family, censored versions of the likelihood are readily available, but parameters are subject to non linear constraints and structural modeling choices have to be made *a priori*, *e.g.*, by allowing only bi-variate or tri-variate dependence between closest neighbors.

35 One semi-parametric compromise consists in using mixture models, built from a potentially infinite number of parametric components, such as the Dirichlet mixture model (DM), first introduced by Boldi and Davison (2007). They have shown that it can approach arbitrarily well any valid angular measure for extremes. A re-parametrized version of the DM model (Sabourin and Naveau, 2014), allows
40 for consistent Bayesian inference - thus, a straightforward uncertainty assessment using posterior credible sets - with a varying number of mixture components *via* a *reversible-jumps* algorithm. The approach is appropriate for data sets of moderate dimension (typically, $d \approx 5$).

The purpose of the present work is to adapt the DM model to the case of cen-
45 sored data. The difficulties are two-fold: First, from a modeling perspective, when the censoring intervals overlap the extremal thresholds (determined by preliminary analysis), one cannot tell whether the event must be treated as extreme. The

¹preprint available on <https://hal.archives-ouvertes.fr/hal-01087687>

proposed approach here consists in reformulating the *Peaks-over-threshold* (POT) model originally proposed by Boldi and Davison (2007) and Sabourin and Naveau (2014), in terms of a *Poisson model*, in which the censored regions overlapping the threshold have a well-defined likelihood. The second challenge is numerical and algorithmic: for right-censored data above the extremal threshold (not overlapping it), the likelihood expression involves integrals of a density over rectangular regions, which have no analytic expression. The latter issue is tackled within a data augmentation framework, which is implemented as an extension of Sabourin and Naveau (2014)'s algorithm for Dirichlet mixtures.

An additional issue addressed in this paper concerns the separation between marginal parameters estimation and estimation of the dependence structure. Performing the two steps separately is a widely used approach, but it boils down to neglecting marginal uncertainty, which confuses uncertainty assessment about joint events such as probabilities of failure regions. It also goes against the principle of using regional information together with the dependence structure to improve marginal estimation, which is the main idea of the popular *regional frequency analysis* in hydrology. In this paper, simultaneous inference of marginal and dependence parameters in the DM model is performed, which amounts in practice to specifying additional steps for the marginal parameters in the MCMC sampler.

The rest of this paper is organized as follows: Section 2 recalls the necessary background for extreme values modeling. The main features of the Dirichlet mixture model are sketched. This POT model is then reformulated as a Poisson model, which addresses the issue of variable threshold induced by the fluctuating marginal parameters. Censoring is introduced in Section 3. In this context, the Poisson model has the additional advantage that censored data overlapping threshold have a well defined likelihood. The lack of analytic expression for the latter is addressed by a data augmentation scheme described in Section 4. The method is illustrated by a simulation study in Section 5: marginal performance in the DM model and in an independent one (without dependence structure) are compared, and the predictive performance of the joint model in terms of conditional probabilities of joint excesses is investigated. The model is also fitted to the hydrological data. Section 6 concludes. Most of the technicalities needed for practical implementation, such as computation of conditional distributions, or details concerning the data augmentation scheme and its consistency are relegated to the appendix.

35

2. MODEL FOR THRESHOLD EXCESSES

2.1. Dependence structure model: angular measures. In this paper, the sample space is the d -dimensional Euclidean space \mathbb{R}^d , endowed with the Borel σ -field. In what follows, bold symbols denote vectors and, unless otherwise mentioned, binary operators applied to vectors are defined component-wise. Let $(\mathbf{Y}_t)_{t \in \mathbb{N}}$ be independent, identically distributed (*i.i.d.*) random vectors in \mathbb{R}^d , with joint distribution \mathbf{F} and margins F_j , $1 \leq j \leq d$. The joint behavior of large observations is best expressed in terms of standardized data. Namely, define

$$\mathbf{X}_t = (-1/\log(F_1(Y_{1,t})), \dots, -1/\log(F_d(Y_{d,t}))).$$

Then the $X_{j,t}$'s have unit-Fréchet distribution, $\mathbb{P}(X_{j,t} \leq x) = e^{-1/x}$, $x > 0$. It is mathematically convenient to switch to pseudo-polar coordinates,

$$R = \sum_{j=1}^d X_j \text{ (radial component)}, \quad \mathbf{W} = \frac{1}{R} \mathbf{X} \in \mathbf{S}_d \text{ (angular component)},$$

where $\mathbf{S}_d = \{\mathbf{x} : x_j \geq 0, \sum_{j=1}^d x_j = 1\}$ is the unit simplex. The radial variable R corresponds to the ‘amplitude’ of the data whereas the angular component \mathbf{W} characterizes their ‘direction’. Asymptotic theory (Resnick, 1987; Beirlant et al., 2004; Coles, 2001) tells us that, under mild assumptions on \mathbf{F} (namely, belonging to a multivariate maximum domain of attraction), an appropriate model, commonly referred to as a multivariate *Peaks-over-threshold* (POT) model, for (R, \mathbf{W}) over high radial thresholds r_0 , is

$$\mathbb{P}(R > r, \mathbf{W} \in A \mid R > r_0) = \frac{r}{r_0} H(A), \quad r_0 > r, A \subset \mathbf{S}_d, \quad (2.1)$$

where H is the so-called ‘angular probability measure’ (called the ‘angular measure’ in the remainder of this paper). The angular measure is thus the limiting distribution of the angle, given that the radius is large. Concentration of H ’s mass in the middle of the simplex indicates strong dependence at extreme levels, whereas Dirac masses only on the vertices characterizes asymptotic independence. This paper focuses on the case where H is concentrated on the interior of the simplex, so that all the variables are asymptotically dependent.

Because of the standardization to unit Fréchet, a probability measure H on \mathbf{S}_d is a valid angular measure if and only if $\int_{\mathbf{S}_d} w_j dH(\mathbf{w}) = \frac{1}{d}$ ($1 \leq j \leq d$). This moments constraint is the only condition on H , so that the angular measure has no reason to be part of any particular parametric family.

2.2. Dirichlet mixture angular measures. In this paper, the angular measure H is modeled by a Dirichlet mixture distribution (Boldi and Davison, 2007; Sabourin and Naveau, 2014). A Dirichlet distribution can be characterized by a shape $\nu \in \mathbb{R}^+$ and a center of mass $\boldsymbol{\mu} \in \mathbf{S}_d$, so that its density with respect to the $d - 1$ dimensional Lebesgue measure $d\mathbf{w} = dw_1 \cdots dw_{d-1}$, is

$$\text{diri}_{\nu, \boldsymbol{\mu}}(\mathbf{w}) = \frac{\Gamma(\nu)}{\prod_{j=1}^d \Gamma(\nu \mu_j)} \prod_{j=1}^d w_j^{\nu \mu_j - 1} \quad (\mathbf{w} \in \mathbf{S}_d). \quad (2.2)$$

A parameter for a k -mixture is of the form

$$\psi = ((p_1, \dots, p_k), (\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_k), (\nu_1, \dots, \nu_k)),$$

with weights $p_m > 0$, such that $\sum_{m=1}^k p_m = 1$. This is summarized by writing $\psi = (p_{1:k}, \boldsymbol{\mu}_{1:k}, \nu_{1:k})$. The corresponding mixture density is

$$h_\psi(\mathbf{w}) = \sum_{m=1}^k p_m \text{diri}_{\nu, \boldsymbol{\mu}_m}(\mathbf{w}). \quad (2.3)$$

the moments constraint is satisfied if and only if

$$\sum_{m=1}^k p_m \boldsymbol{\mu}_m = (1/d, \dots, 1/d),$$

which, in geometric terms, means that the center of mass of the $\boldsymbol{\mu}_{1:k}$ ’s, with weights $p_{1:k}$, must lie at the center of the simplex. As established by Boldi and Davison (2007) and mentioned in the introduction, the family of Dirichlet mixture densities satisfying the moments constraint is weakly dense in the space of admissible angular measure. In addition, in a Bayesian framework, Sabourin and Naveau (2014) have

shown that the posterior is weakly consistent under mild conditions. These two features put together make the Dirichlet mixture model an adequate candidate for modeling the angular components of extremes.

2.3. Model for margins. The above model for excesses concerns standardized versions \mathbf{X}_t of the data \mathbf{Y}_t involving marginal cumulative distribution function F_j ($1 \leq j \leq d$), which have to be estimated. As a consequence of uni-variate extreme value theory (Pickands, 1975), uni-variate excesses above large thresholds v_j ($1 \leq j \leq d$) are approximately distributed according to a Generalized Pareto distribution with parameters ξ_j (shape) and σ_j (scale parameter),

$$P(Y_j > y \mid Y_j > v_j) \approx_{v_j \rightarrow \infty} \left(1 + \xi_j \frac{y - v_j}{\sigma_j}\right)^{-1/\xi_j}.$$

A widely used method to model the largest excesses is as follows: Define a high multivariate threshold $\mathbf{v} = (v_1, \dots, v_d)$ and call ‘marginal excess’ any $Y_{j,t} > v_j$. Then, marginal excesses above v_j are modeled as generalized Pareto random variables with parameters ξ_j and σ_j . The marginal parameters are gathered into a $(2d)$ -dimensional vector

$$\chi = (\log(\sigma_1), \dots, \log(\sigma_d), \xi_1, \dots, \xi_d) \in \mathbb{R}^{2d}.$$

Let $F_j^{\mathbf{y}}$ denote the j^{th} marginal distribution conditionally on Y_j not exceeding v_j , and let $\zeta_j = \mathbf{P}(Y_j > v_j)$ denote the probability of excursion above v_j . The j^{th} marginal model ($1 \leq j \leq d$) is thus

$$\begin{aligned} F_j^{(x)}(y) &= \mathbf{P}(Y_{j,t} \leq y \mid \xi_j, \sigma_j) \\ &= \begin{cases} 1 - \zeta_j \left(1 + \xi_j \frac{y - v_j}{\sigma_j}\right)^{-1/\xi_j} & (y \geq v_j), \\ (1 - \zeta_j) F_j^{\mathbf{y}}(y) & (y < v_j). \end{cases} \end{aligned} \quad (2.4)$$

It is common practice (Coles and Tawn, 1991; Davison and Smith, 1990) to use an empirical estimate $\hat{\zeta} = (\hat{\zeta}_1, \dots, \hat{\zeta}_d)$ for the vector of probabilities of marginal excursion, and to ignore any estimation error, so that $\hat{\zeta}$ is identified to ζ in the sequel.

2.4. Joint inference in a Poisson model. When it comes to simultaneous estimation of the margins and of the angular measure, the angular model (2.1) for radial excesses is difficult to handle, because a radial failure region $r > r_0$ on the Fréchet scale (*i.e.*, in terms of \mathbf{X} 's) corresponds to a complicated shaped failure region on the original scale, which depends on the marginal parameters and, accordingly, potentially contains a varying number of data points. It seems more reasonable to use a failure region which is fixed on the original scale (in terms of \mathbf{Y} 's). Further, a common criticism towards radial failure regions (Ledford and Tawn, 1996) is that the marginal Pareto model is not valid near the axes of the positive orthant. Last but not least, censoring occurs along the directions of the Cartesian coordinate system, which prevents using the polar model (2.1) as it is. To address these issues, the statistical model for threshold excesses developed in this paper uses a ‘rectangular’ threshold. Also, it will be very convenient (see Section 3.3) to adopt a Poisson process representation of extremes (see *e.g.* Coles and Tawn, 1991) as an alternative to the POT model (2.1), with a ‘censored likelihood’ near the axes.

Poisson model. Under the same condition of domain of attraction as above, the point process formed by time-marked, standardized and suitable re-scaled data converges in distribution to a Poisson process (see *e.g.* Resnick, 1987, 2007; Coles and Tawn, 1991),

$$\sum_{t=1}^n \mathbb{1}_{\left(\frac{t}{n}, \frac{\mathbf{x}_t}{n}\right)} \xrightarrow{w} \text{PP}(\ell \otimes \lambda), \quad (2.5)$$

in the space of point measures on $([0, 1] \times \mathbf{E})$, where $\mathbf{E} = [0, \infty]^d \setminus \{0\}$. The temporal component ℓ of the limiting intensity measure denotes the Lebesgue measure on \mathbb{R} and λ , the so-called *exponent measure*, is homogeneous of order -1 , and is related to the angular measure H via

$$d\lambda(r, \mathbf{w}) = \frac{d}{r^2} dr dH(\mathbf{w}). \quad (2.6)$$

From a statistical perspective, consider a failure region $A_{\mathbf{v}} = \mathbf{E} \setminus [0, \mathbf{v}]$, where \mathbf{v} is the high multivariate threshold introduced in section 2.3 and $[0, \mathbf{v}] = [0, v_1] \times \cdots \times [0, v_d] \setminus \{0\}$. Call ‘excess above \mathbf{v} ’ any point \mathbf{Y}_t in $A_{\mathbf{v}}$, as opposed to marginal excesses $Y_{j,t} > v_j$. The Fréchet re-scaled multivariate threshold is

$$\mathbf{u} = \mathbf{T}(\mathbf{v}) = -1/\log(1 - \zeta)$$

and does not depend on χ . Consider the re-scaled region on the Fréchet scale

$$A_{\mathbf{u},n} = \frac{1}{n} \mathbf{T}(A_{\mathbf{v}}) = [0, \infty]^d \setminus [0, \frac{u_1}{n}] \times \cdots \times [0, \frac{u_d}{n}],$$

and denote $A_{\mathbf{u}} = A_{\mathbf{u},1}$. Applying the marginal transformations

$$\mathcal{T}_j^{\chi}(y) = -1/\log\left(F_j^{(\chi)}(y)\right),$$

the marginal variables $X_{j,t} = \mathcal{T}_j^{\chi}(Y_{j,t})$ have unit Fréchet distribution, as required in (2.5). The point process $\mathbf{N} = \sum_{t=1}^n \mathbb{1}_{\left(\frac{t}{n}, \frac{\mathbf{x}_t}{n}\right)}$ composed of the excesses $\mathbf{X}_t \in A_{\mathbf{u}}$ (*i.e.* $\mathbf{Y}_t \in A_{\mathbf{u}}$) is modeled according to the limit in (2.5),

$$\sum_{t=1}^n \mathbb{1}_{\left(\frac{t}{n}, \frac{\mathbf{x}_t}{n}\right)} \sim \text{PP}(\ell \otimes \lambda) \text{ on } [0, 1] \times A_{\mathbf{u},n},$$

where λ is of the form (2.6), with angular component H written as a Dirichlet mixture of the form (2.3).

Joint likelihood of uncensored data. Let $\theta = (\chi, \psi)$ be the parameter for the joint model. As explained at the beginning of this section, the model likelihood needs to be expressed in Cartesian coordinates. The density of an exponent measure λ with respect to the d -dimensional Lebesgue measure $d\mathbf{x} = dx_1 \cdots dx_d$, is (Coles and Tawn, 1991, Theorem 1)

$$\frac{d\lambda}{d\mathbf{x}}(\mathbf{x}) = d \cdot r^{-(d+1)} h(\mathbf{w}). \quad (2.7)$$

Denote by λ_{ψ} the exponent measure corresponding to the Dirichlet mixture h_{ψ} . Then, in the simplified case where the $\mathbf{Y}_{j,t}$ ’s are exactly observed and where the marginals F_j^{χ} ’s below threshold are known, the likelihood in the Poisson model over $A_{\mathbf{v}}$ is

$$\mathcal{L}_{\mathbf{v}}(\{\mathbf{y}_t\}_{1 \leq t \leq n}, \theta) \propto e^{-n \lambda_{\psi}(A_{\mathbf{u}})} \prod_{i=1}^{n_{\mathbf{v}}} \left\{ \frac{d\lambda_{\psi}}{d\mathbf{x}}(\mathbf{x}_{t_i}) \prod_{j: y_{j,t_i} > v_j} J_j^{\chi}(y_{j,t_i}) \right\}, \quad (2.8)$$

where $t_1, \dots, t_{n_{\mathbf{v}}}$ are the occurrence times of excesses $\mathbf{y}_{t_i} \in A_{\mathbf{v}}$, $x_{j,t_i} = T_j^X(y_{j,t_i})$, and the J_j^X are Jacobian terms resulting from marginal transformations T_j^X (see Appendix A for details).

3. CENSORED MODEL

5 3.1. **Causes of censoring.** The presence of censored observations is the result of two distinct causes: first, data are partially observed, which results in interval- or right-censoring, which we call *natural censoring*. In addition, observed data points that exceed at least one threshold in one direction do not necessarily exceed all thresholds, so that the marginal extreme value model does not apply. Following
 10 Ledford and Tawn (1996), those components are also considered as left-censored. This second censoring process is thus a consequence of an inferential framework which is designed for analyzing extreme values only, and we call it *inferential censoring*.

The total censoring process \mathcal{C} , which results from the juxtaposition of natural
 15 and inferential censoring, is assumed to be non informative. This means that (see also Gómez et al., 2004), if F is the marginal *c.d.f.* for Y_j and f is the marginal density, then Y_j 's distribution conditional upon having observed only the left and right censoring bounds (L, R) is $f(\cdot)/[F(R) - F(L)]$. This definition is easily extended to the multivariate case by replacing $F(R) - F(L)$ by the integral of
 20 the density over the censored directions (see *e.g.* Schnedler, 2005, for a proof of consistency of maximum censored likelihood estimators).

3.2. **Natural censoring.** Call ‘Natural censoring’ the one which occurs independently from the choice of an extreme threshold \mathbf{v} by the statistician. The observed process is denoted $\mathbf{O} = (\mathbf{O}_t)_t$, with $\mathbf{O}_t = (O_{1,t}, \dots, O_{d,t})$. One marginal observa-
 25 tion $O_{j,t}$ consists in a label $\kappa_{j,t}$ indicating presence or absence of censoring, together with the exact data $Y_{j,t}$ (if observed) or the censoring bounds $(L_{j,t}, R_{j,t})$ (where $L_{j,t}$ may be set to 0 in the case of left-censoring or missing data and $L_{j,t} = +\infty$ in the case of right-censored or missing data). In the sequel, $\kappa_{j,t} = 0$ (*resp.* 1, 2, 3) refer respectively to missing, exact, right- and left- censored data.

30 In this context, the ‘position’ of a marginal data $O_{j,t}$ with respect to the threshold is not necessarily well defined for censored data. Recall that we consider a case where the process of interest $(\mathbf{Y}_t)_{t \geq 0}$ is stationary, whereas the censoring bounds vary with time, as a result of external factors on the observation process. When the censoring interval overlaps the threshold, *i.e.*, $v_j \in (L_{j,t}, R_{j,t})$, $\kappa_j \in \{2, 3\}$,
 35 the statistician does not know if an excess occurred or not. This situation is described here as *$O_{j,t}$ marginally overlapping the threshold*. The different positions of $O_{j,t}$ with respect to v_j encountered in the data set of interest in this paper are summarized in Figure 1.

For a multivariate observation \mathbf{O}_t , if at least one coordinate marginally overlaps
 40 the threshold or is missing, and if the others are below the threshold, then the position of \mathbf{O}_t with respect to the *multivariate threshold* \mathbf{v} is undetermined. Indeed, \mathbf{O}_t is *below threshold* if all its marginals are below the corresponding marginal threshold, and *above threshold* (in the failure region) otherwise. In the undetermined case, \mathbf{O}_t is qualified as *globally overlapping the threshold*.

45 3.3. **Inferential censoring below threshold.** Since the marginal distributions $F_j^{\mathbf{y}}$'s, conditional upon not exceeding v_j , are unknown, the $X_{j,t}$'s such that $Y_{j,t} < v_j$

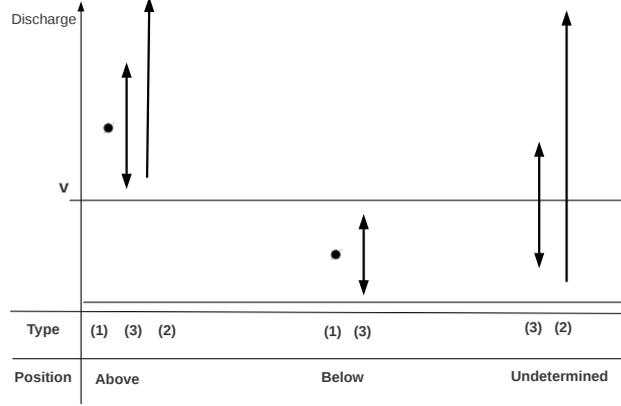


FIGURE 1. Position of marginal data points with respect to a marginal threshold v (horizontal line). Black dots: marginal data points of type $\kappa = 1$; vertical arrows: data of type $\kappa \in \{2, 3\}$.

are not available. Instead of attempting to estimate the F_j^y 's, one option is to censor the Fréchet-transformed components below threshold. More precisely, for a raw observation $O_{j,t} = (\kappa_{j,t}, Y_{j,t}, L_{j,t}, R_{j,t})$, let us denote by $C_{j,t}^X = (\tilde{\kappa}_{j,t}, X_{j,t}, \tilde{L}_{j,t}, \tilde{R}_{j,t})$ the corresponding 'Fréchet transformed' and censored one, and \mathbf{C}_t^X the multivariate observation $(C_{1,t}^X, \dots, C_{d,t}^X)$. The transformation $\mathbf{O}_t \mapsto \mathbf{C}_t^X$ is illustrated in Figure 2 in the bi-variate case.

In a nutshell, inferential censoring occurs when the censoring intervals are below threshold or marginally overlapping it, or when an exact data component $Y_{j,t} < v_j$ is recorded. A formal definition of C_j^X is as follows:

$$\begin{aligned}
 10 \quad \bullet \quad (\tilde{\kappa}_{j,t}, X_{j,t}) &= \begin{cases} (3, \text{NA}) & \text{if } \kappa_{j,t} = 1 \text{ and } Y_{j,t} < v_j, \\ (0, \text{NA}) & \text{if } \kappa_{j,t} = 2 \text{ and } L_{j,t} < v_j, \\ (\kappa_{j,t}, \mathcal{T}_j^X(Y_{j,t})) & \text{otherwise.} \end{cases} \\
 \bullet \quad \tilde{L}_{j,t} &= \begin{cases} 0 & \text{if } L_{j,t} < v_j, \\ \mathcal{T}_j^X(L_{j,t}) & \text{otherwise.} \end{cases} \\
 \bullet \quad \tilde{R}_{j,t} &= \begin{cases} u_j & \text{if } R_{j,t} < v_j, \\ \mathcal{T}_j^X(R_{j,t}) & \text{otherwise.} \end{cases}
 \end{aligned}$$

In the above definition, NA stands for a missing value and it is understood that $\mathcal{T}_j^X(\text{NA}) = \text{NA}$.

15 In the end, observations globally overlapping threshold have their marginal lower bounds $\tilde{L}_{j,t}$ set to zero if $L_{j,t} < v_j$. The interest of using a Poisson model becomes clear at this point. Indeed, censored observations \mathbf{C}_t^X obtained from observations \mathbf{O}_t globally overlapping threshold correspond to events of the kind 'The observation at time t belongs to $[\mathbf{0}, \tilde{\mathbf{R}}_t]$ ', which, by contrapositive, means 'No point is observed outside of $[\mathbf{0}, \tilde{\mathbf{R}}_t]$ between t and $t + 1$ '. This is written in terms of the Poisson process \mathbf{N} as

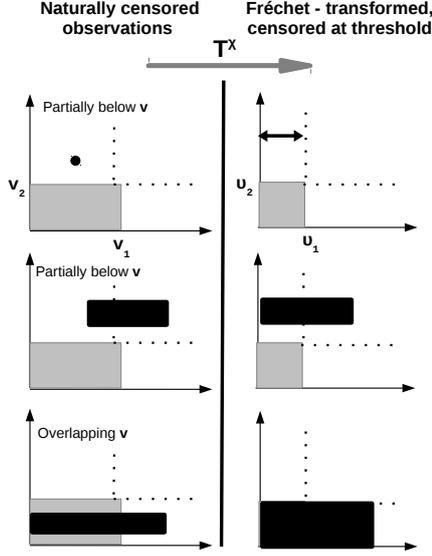


FIGURE 2. Example in the two-dimensional case of marginal transformation and censoring below threshold, for data of marginal types (1, 1), (upper panel) (3, 1) (middle panel) and (3, 3) (lower panel). In these three cases, observations are represented in black, respectively by a dot, an horizontal arrow and a rectangle. The Grey areas represent respectively the multivariate threshold \mathbf{v} (left side) and the Fréchet transformed one \mathbf{u} (right side). The two upper panels correspond to observations above threshold, while the lower panel shows an undetermined observation.

$$\mathbf{N} \left(\left(\left[\frac{t}{n}, \frac{t+1}{n} \right] \times \left[\mathbf{0}, \frac{\tilde{\mathbf{R}}_t}{n} \right]^c \right) \right) = 0, \quad (3.1)$$

which is a measurable event with respect to \mathbf{N} . The overlapping observations thus have a well defined likelihood in a Poisson model, as detailed in the next section, whereas they could not be taken into account in Sabourin and Naveau (2014)'s POT model.

- 5 3.4. **Poisson likelihood with censored and missing data.** Due to the combination of natural and inferential censoring, the data set (from which missing days are excluded) is decomposed into data in the failure region, data overlapping threshold and data below threshold. Let $n_{\mathbf{v}}$, $n'_{\mathbf{v}}$ and $n''_{\mathbf{v}}$ be the respective number of observations in each category. The number of non missing days is thus

$$n_{\text{obs}} = n_{\mathbf{v}} + n'_{\mathbf{v}} + n''_{\mathbf{v}},$$

- 10 and the number of 'determined' data (*i.e.* not overlapping \mathbf{v}) is

$$n_{\text{det}} = n_{\mathbf{v}} + n''_{\mathbf{v}}.$$

The $n_{\mathbf{v}}$ Fréchet-transformed observations $\{\mathbf{C}_{t_i}^{\mathbf{x}}\} (i \in \{1, \dots, n_{\mathbf{v}}\})$ correspond to events of the kind

$$\mathbf{X}_{t_i} \in [\tilde{\mathbf{L}}_{t_i}, \tilde{\mathbf{R}}_{t_i}] \quad (\text{rectangular region})$$

where $\tilde{R}_{j,t_i} = \tilde{L}_{j,t_i} = X_{j,t_i}$ in the case of exact data.

Observations overlapping threshold correspond to events of the kind (3.1) introduced in the previous section. When a limited number \mathcal{I}' of right censoring bounds $\mathbf{R}_t = (R_{1,t}, \dots, R_{d,t})$ are present, it is convenient to classify these overlapping

5 events accordingly, writing $n'_v = \sum_{i=1}^{\mathcal{I}'} n'_i$ where n'_i is the number of observations with right censoring bound \mathbf{R}_{t_i} . Since the Poisson process is temporally exchangeable, there is no loss of generality in assuming that the latter observations occur at consecutive dates $(t'_i, \dots, t'_i + n'_i - 1)$. The region $\mathbf{E} \setminus [\mathbf{0}, \tilde{\mathbf{R}}_{t'_i}]$ is ‘missed’ by the Fréchet re-scaled process \mathbf{X}_t during this time period.

10 With these notations, the censored likelihood in the Poisson model may be written

$$\begin{aligned} \mathcal{L}_v(\mathbf{O}, \theta) = & \exp \left[-n_{\text{det}} \lambda_\psi(A_{\mathbf{u}}) - \sum_{i=1}^{\mathcal{I}'} n'_i \lambda_\psi(A'_i) \right] \times \dots \\ & \dots \prod_{i=1}^{n_v} \left\{ \int_{[\tilde{\mathbf{L}}_{t_i}, \tilde{\mathbf{R}}_{t_i}]} \frac{d\lambda_\psi}{d\mathbf{x}} d\ell_i(\mathbf{x}) \prod_{j: y_{j,t_i} > v_j} J_j^X(y_{j,t_i}) \right\}, \end{aligned} \quad (3.2)$$

where notation ‘ $d\ell_i(\mathbf{x})$ ’ in the integral terms is a shorthand for ‘the Lebesgue measure of dimension equal to that of $[\tilde{\mathbf{L}}_{t_i}, \tilde{\mathbf{R}}_{t_i}]$ ’ when the latter is greater than one, or ‘the Dirac mass at $\mathbf{x}_{t_i} = \tilde{\mathbf{L}}_{t_i} = \tilde{\mathbf{R}}_{t_i}$ ’ for exact data. Compared with the uncensored likelihood (2.8), n has been replaced with n_{obs} , the exponential term for the non overlapping data follows from

$$\exp \left(-\frac{n_{\text{det}}}{n_{\text{obs}}} \lambda_\psi(A_{\mathbf{u}, n_{\text{obs}}}) \right) = \exp \left(-n_{\text{det}} \lambda_\psi(A_{\mathbf{u}}) \right),$$

and a similar argument yields the additional terms $\exp(-n'_i \lambda_\psi(A'_i))$ for overlapping data.

At this stage, the model has been entirely specified. The remaining issue concerns the treatment of the integral terms

$$\int_{[\tilde{\mathbf{L}}_{t_i}, \tilde{\mathbf{R}}_{t_i}]} \frac{d\lambda_\psi}{d\mathbf{x}} d\ell_i(\mathbf{x}) \quad (3.3)$$

and the exponential terms

$$\exp \left[-n_{\text{det}} \lambda_\psi(A_{\mathbf{u}}) \right] \quad \text{and} \quad \exp \left[-n'_i \lambda_\psi(A'_i) \right], \quad (3.4)$$

which have no analytic expression, as they require integrating λ_ψ over rectangular regions. First, the dimension of numerical integration can be reduced as far as ‘missing coordinates’ are involved, because partial integration of λ_ψ over $[0, \infty]$ in one direction has an exact expression. The model is stable under marginalization, in the sense that the obtained marginal exponent measures correspond again to Dirichlet mixtures on a lower dimensional simplex (see Appendix B for details).

30 However, no closed form is available for the integral in the remaining censored directions, nor for the exponent measures of $A_{\mathbf{u}}$ or the A'_i ’s. This problem is tackled in the next section *via* a data augmentation method.

4. DATA AUGMENTATION

4.1. **Background.** In a Bayesian context, one major objective is to generate parameter samples approximately distributed according to the posterior. In classical MCMC algorithms, the value of the likelihood is needed to define the transition

kernel. Evaluating the integrated likelihood $\mathcal{L}_{\mathbf{v}}(\mathbf{O}, \theta)$ at each iteration of the algorithm seems unmanageable: The dimension of integration can grow up to d , for each observation, and the shape of the integrand varies from one iteration to another, which is not favorable to standard quadrature methods. In particular, large or low values of the shape parameters ν_m in ψ induce concentration of the integrand around the centers $\boldsymbol{\mu}_m$ or unboundedness at the simplex boundaries. Instead, data augmentation methods (see *e.g.* Tanner and Wong, 1987; Van Dyk and Meng, 2001) treat missing or partially observed data as additional parameters, so that the numerical integration step is traded against an increased dimension of the parameter space. In this section, $[\cdot]$ denotes the distribution of a random quantity as well as its density with respect to some appropriate reference measure. Proportionality between σ -finite measures is denoted by \propto . Thus, $[\theta]$ is the prior density and $[\theta | \mathbf{O}] \propto [\theta] \mathcal{L}_{\mathbf{v}}(\mathbf{O}, \theta)$ is the posterior.

The main idea is to define an augmentation space \mathcal{Z} , and a probability measure $[\cdot | \mathbf{O}]_+$, on the augmented space $\Theta \times \mathcal{Z}$, conditional on the observations, which may be sampled using classical MCMC methods, and which is consistent with the ‘objective distribution’ on Θ . That is, the posterior on Θ must be obtained by marginalization of $[\cdot | \mathbf{O}]_+$,

$$[\theta | \mathbf{O}] = \int_{\mathcal{Z}} [z, \theta | \mathbf{O}]_+ dz. \quad (4.1)$$

In the sequel, the invariant measure (or its density) $[z, \theta | \mathbf{O}]_+$ is referred to as the *augmented posterior*.

4.2. Data augmentation in the Poisson model. In our context, finding an augmentation random variable Z_+ with easily manageable conditional distributions, such that the augmented posterior $[z, \theta | \mathbf{O}]_+$ satisfy (4.1), is far from straightforward, mainly due to the exponential terms (3.4) in the likelihood.

Instead, an intermediate variable \mathbf{Z} is introduced, which plays the role of a proposal distribution in the MCMC algorithm. \mathbf{Z} is defined conditionally to θ , through the *augmented likelihood* $[\mathbf{z}, \mathbf{O} | \theta]$, so that the full conditionals $[\mathbf{Z} | \theta, \mathbf{O}]$ can be directly simulated as block proposals in a *Metropolis-within-Gibbs* algorithm (Tierney, 1994). To ensure the marginalization condition (4.1), the augmented posterior $[\mathbf{Z}, \theta | \mathbf{O}]_+$ is *not* the same as $[\mathbf{Z}, \theta | \mathbf{O}]$. Instead it has a density of the form

$$[\mathbf{z}, \theta | \mathbf{O}]_+ \propto [\mathbf{z}, \mathbf{O} | \theta] [\theta] \varphi(\mathbf{z}), \quad (4.2)$$

where $\varphi(\mathbf{z})$ is any weight function allowing to enforce the consistency condition (4.1). The $[\theta]$ terms cancel out and the latter condition is equivalent to

$$\mathcal{L}_{\mathbf{v}}(\mathbf{O}, \theta) \propto \int [\mathbf{z}, \mathbf{O} | \theta] \varphi(\mathbf{z}) d\mathbf{z}. \quad (4.3)$$

In the end, a posterior sample from $[\theta | \mathbf{O}]$ is simply obtained by ignoring the \mathbf{Z} -components from the one produced with the ‘augmented’ Markov chain.

In our case, the augmented data \mathbf{Z} consist of two parts, $\mathbf{Z} = (\mathbf{Z}_{\text{above}}, \mathbf{Z}')$. The first one, $\mathbf{Z}_{\text{above}} = \{\mathbf{Z}_{t_i}\}_{i \leq n_{\mathbf{v}}}$, accounts for integral terms (3.3), while the second one, $\mathbf{Z}' = (\mathbf{Z}'_{\mathbf{u}}, \{\mathbf{Z}'_i\}_{i \leq \mathcal{I}'})$, accounts for the exponential terms (3.4). The \mathbf{Z}_{t_i} ’s have an intuitive interpretation, which is standard in data augmenting methods: they replace the censored X_{j, t_i} ’s. On the contrary, the \mathbf{Z}'_i ’s and the $\mathbf{Z}'_{\mathbf{u}}$ are just a

computational trick accounting for the exponential terms: they are the points of independent Poisson processes defined on ‘nice’ radial sets containing the failure regions of interest A'_i ’s and $A_{\mathbf{u}}$, and φ is a smoothed version of an indicator function of the failure regions.

5 The augmented likelihood factorizes as

$$[\mathbf{z}, \mathbf{O} \mid \theta] = \prod_{i=1}^{n_{\mathbf{v}}} \{[\mathbf{z}_{t_i}, \mathbf{O}_{t_i} \mid \theta]\} [\mathbf{z}'_{\mathbf{u}} \mid \psi] \prod_{i=1}^{\mathcal{I}'} [\mathbf{z}'_i \mid \psi], \quad (4.4)$$

and the weight function φ is of the form

$$\varphi(\mathbf{z}) = \varphi_{\mathbf{u}}(\mathbf{z}'_{\mathbf{u}}) \prod_{i=1}^{\mathcal{I}'} \varphi'_i(\mathbf{z}'_i).$$

A precise definition of $\mathbf{Z}_{\text{above}}$ is given in Appendix C.1, together with the expression of the corresponding contributions to the augmented likelihood, $[\mathbf{z}_{t_i}, \mathbf{O}_{t_i} \mid \theta]$. The full conditionals $[\mathbf{Z}_{\text{above}} \mid \mathbf{O}, \theta]$ are derived in C.2. The augmentation Poisson processes \mathbf{Z}' , together with the weight φ are defined in Appendix C.3, and the compatibility condition (4.3) is proved to hold in Appendix C.4.

4.3. Implementation of a MCMC algorithm on the augmented space.

This section describes only the main features of the algorithm, more details are provided in Appendix D. The present algorithm is an extension of Sabourin and Naveau (2014)’s one, who proposed a *Metropolis-within-Gibbs* algorithm to sample the posterior distribution of the angular measure in a POT framework (2.1), within the Dirichlet mixture model (2.3). The number of components in the Dirichlet mixture is not fixed in their model, and the MCMC allows *reversible-jumps* between parameters sub-spaces of fixed dimension, each corresponding to a fixed number of components in the Dirichlet mixture. Their algorithm approximates the posterior distribution $[\psi \mid \mathcal{W}]$, where ψ is the DM parameter and \mathcal{W} is an angular data set consisting of the angular components W_t of data X_t that have been normalized to Fréchet margins in a preliminary step.

In contrast, the present algorithm handles ‘raw’ data (not standardized to Fréchet margins), with censoring of various types as described in Section 3, so as to approach the full posterior distribution (marginal and dependence parameters) $[\theta \mid \mathbf{O}]$. In practice, this amounts to allowing two additional move types in the *Metropolis-within-Gibbs* sampler: *marginal moves* (updating the marginal parameters) and *augmentation moves* updating the augmentation data \mathbf{Z} described above. The reversible jumps and the moves updating the DM parameters keeping the dimension constant are unchanged.

5. SIMULATIONS AND REAL CASE EXAMPLE

Keeping in mind the application, the aim of this section is to verify that the algorithm provides reasonable estimates with data sets that ‘resemble’ the particular one motivating this work.

After a description of the experimental setting, an example of results obtained with a single random data set is given, then a more systematic study is conducted over 50 independent data sets. Finally, a brief description of the results obtained with the original hydrological data is provided. The latter part is kept short

because, as mentioned in introduction, from a hydrological point of view, a full discussion of the benefits brought by historical information, as well as treatment of temporal dependence in the case of heavily censored data is needed and will be the subject of a separate paper.

- 5 **5.1. Experimental setting.** For this simulation study, the marginal shape parameters are constrained to be equal to each other, $\xi_1 = \dots = \xi_d := \xi$, in accordance with the *regional frequency analysis* hydrological framework (Hosking and Wallis, 2005), where the different gauging stations under study are relatively close to each other (in the same watershed). Also, the dimension is set to $d = 4$, as it is
10 for the hydrological data set of interest.

Preliminary likelihood maximization (with respect to the marginal parameters) is performed on the hydrological data, again assuming independence between locations for the sake of simplicity and imposing a common shape parameter (this latter hypothesis not being rejected by a likelihood ratio test). Then, data sets are
15 simulated according to the model for excesses (2.1), with marginal parameters and threshold excess probabilities (for daily data) approximately equal to the inferred ones, *i.e.*

$$\zeta \approx (0.021, \dots, 0.021), \quad \xi = 0.4, \quad \log(\sigma) = (4.8, 4.6, 5.9, 5.1),$$

for a total number of days $n = 148\,401$, which is the total number of days in the original data. The four-variate dependence structure is chosen as a Dirichlet
20 mixture distribution h_ψ , where

$$\psi : \quad \mathbf{p} = (0.25, 0.25, 0.5), \quad \boldsymbol{\mu} = \begin{pmatrix} 0.1 & 0.7 & 0.1 \\ 0.1 & 0.1 & 0.4 \\ 0.1 & 0.1 & 0.4 \\ 0.7 & 0.1 & 0.1 \end{pmatrix}, \quad \boldsymbol{\nu} = (70, 50, 80).$$

A radial threshold for simulation on the Fréchet scale and the number of points simulated above the latter are respectively set to $r_s = -1/\log(1 - \max(\zeta))$, $n_{\text{rad.exc}} = n * 4/r_s$. The remaining $n - n_{\text{rad.exc}}$ points are arbitrarily scattered below the radial threshold, so that the proportion of radial excesses is $4/r_s$, the
25 exponent measure of the region $\{\mathbf{x} \in \mathbb{R}^4 : \|\mathbf{x}\|_1 > r_s\}$.

Afterwards, the data set is censored following the real data's censoring pattern, *i.e.* censoring occurs on the same days and on the same locations (here, a location is a coordinate $j \in \{1, \dots, 4\}$) as for the real data, with same censoring bounds (on the Fréchet scale), so that the data are observed only if the
30 censoring threshold is exceeded. Finally, in order to account for the loss of information resulting from time dependence within the real data (whereas the simulated data are time independent), only n_{eff} data out of n are kept for inference, where $n_{\text{eff}} = \lfloor n/\text{mean cluster size} \rfloor = 118\,911$ (see Section 6 for an explanation about clusters). The vast majority of data points (real as well as simulated) are either
35 censored or below the multivariate threshold: in the real data set, only 125 multivariate excesses above threshold are recorded, among which only 3 have all their coordinates of type 1 (exact data). In such a context, a simplified inferential framework in which censored data would be discarded is not an option: only 3 points would be available for inference. For simulated data, the threshold is arbitrarily
40 set to the same value as the one defined for real data, *i.e.* $\mathbf{v} = (300, 320, 520, 380)$.

The MCMC sampler described in Section 4.3 and Appendix D is run for each simulated data set, yielding a parameter sample, which approaches the posterior distribution given the censored data. To allow comparison with the default space-independent framework in terms of marginal estimates, such as probabilities of a marginal excess, Bayesian inference is also performed in the independent model defined as follows: the marginal models are the same as in (2.4), the $Y_{j,t}, 1 \leq j \leq 4$ are assumed to be independent while the shape parameters $\xi_j, 1 \leq j \leq 4$ are, again, equal to each other. A MCMC sampler is straightforwardly implemented for this independent model, following the same pattern as defined in the *marginal moves* for the full model (see Appendix D). In the sequel, the full Poisson model with Dirichlet mixture dependence structure is referred to as the *DM Poisson model*, or the *dependent model*, as opposed to the independent model.

5.2. Illustration of the method with one simulated data set. The censoring pattern described above yields a censored data set which resembles the real data set in terms of average number of exact (uncensored) coordinates in each observation, as shown in Figure 3: most of the extracted data have only one exact coordinate.

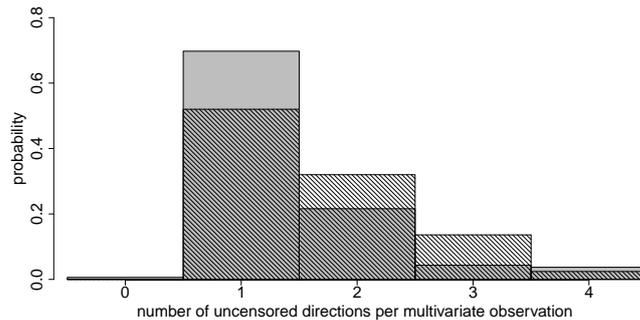


FIGURE 3. Average number of exactly observed coordinates above threshold in each recorded multivariate excess in the real data set (hatched bars) and in a simulated data set (Gray bars).

In addition to data above threshold, the number of threshold-overlapping blocks (made of data which position with respect to the threshold is undetermined, see Section 3) is $\mathcal{I}' = 39$ in this simulated data set, with block sizes varying between 1 and 39 845, and a total number of ‘undetermined’ days $n_{\mathcal{V}}' = 112\,676$. To fix ideas, the bi-variate projections onto the six pairs $(i < j)$ of a simulated data set are displayed in Figure 4.

Let us turn to results obtained with this particular data set. A ‘flat’ prior was specified for the DM Poisson model parameters and MCMC proposals for the DM parameter were set in a similar way as in Sabourin and Naveau (2014), which resulted in satisfactory convergence diagnostics after 10^6 iterations, see Appendix D.6 for details. Figure 5 displays the posterior predictive for bi-variate marginalization’s of the angular density, obtained *via* equation (B.2), together with the true density and posterior credible sets around the estimates. The estimated density captures reasonably well the features of the true one and the posterior quantiles are rather concentrated around the true density. This is all the more satisfying

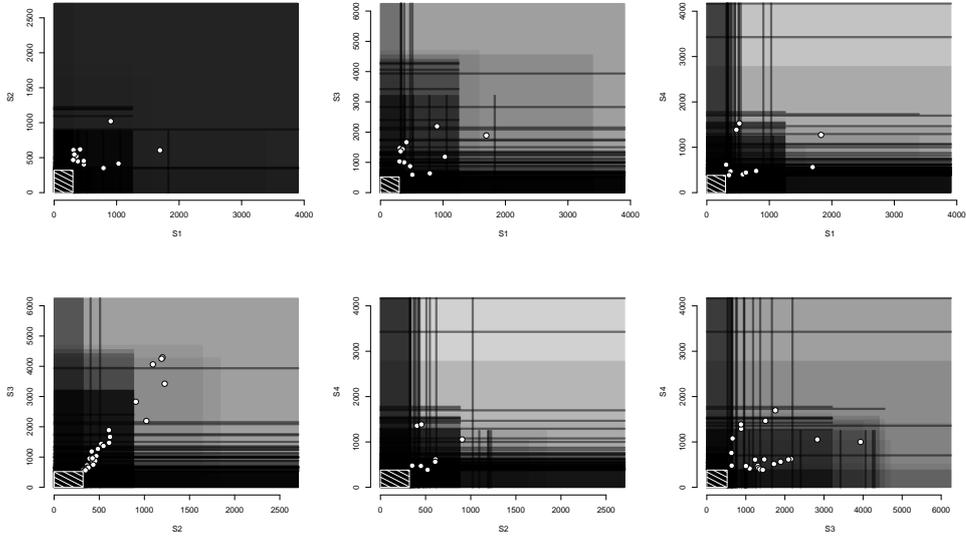


FIGURE 4. Bi-variate projection of the simulated data set after censoring. White points correspond to pairs for which both coordinates are observed. Superimposed Gray rectangles (resp. segments) represent pairs for which both coordinates (resp. one coordinate) are (is) censored. The white striped rectangle is the region below the multivariate threshold \mathbf{v} .

that, at first view (Figure 4), the censored data used for inference seem to convey little information about the distribution of the angular components.

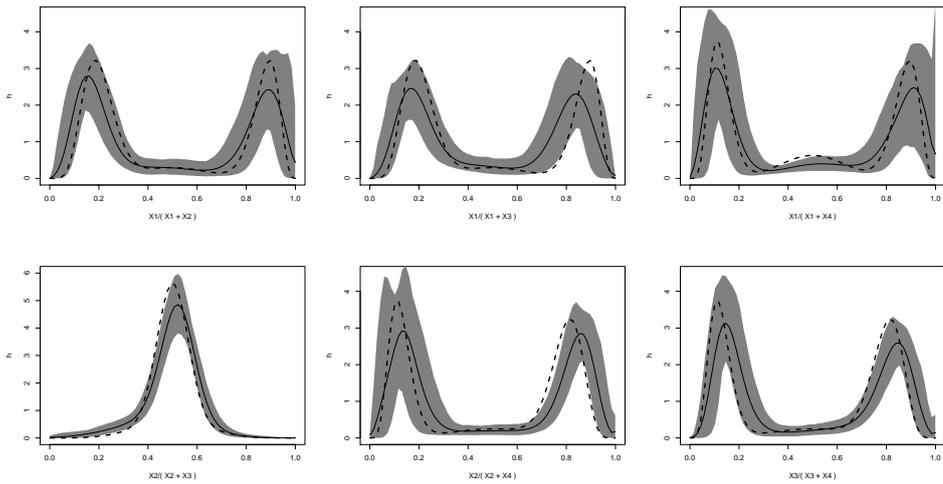


FIGURE 5. Bi-variate angular predictive densities (thick black line) with posterior credible sets (Gray regions) corresponding to the posterior 0.05 – 0.95 quantiles, together with the true angular density (dashed line).

In risk analysis, especially in hydrology, return level plots (*i.e.* quantile plots) are used to summarize the marginal behavior of extremes. The return level Q for a return period T at location j , when marginal data are distributed according to F_j^X and where there is no temporal dependence, is usually defined as the $1 - 1/T$ -
5 quantile of F_j^X . Figure 6 compares the return levels obtained both in the dependent and the independent models, together with the true return levels. The posterior estimates in the dependent model are very close to the truth, relatively to to the size of the credible intervals. In contrast, estimation in the independent model under-estimates the return levels, and the true curve lies outside the posterior
10 quantiles at two locations out of four. This single example is however not enough to conclude that the dependence structure improves significantly marginal estimation. The absence of significant improvement (or deterioration) of marginal estimates is indeed one of the conclusions of the next subsection.

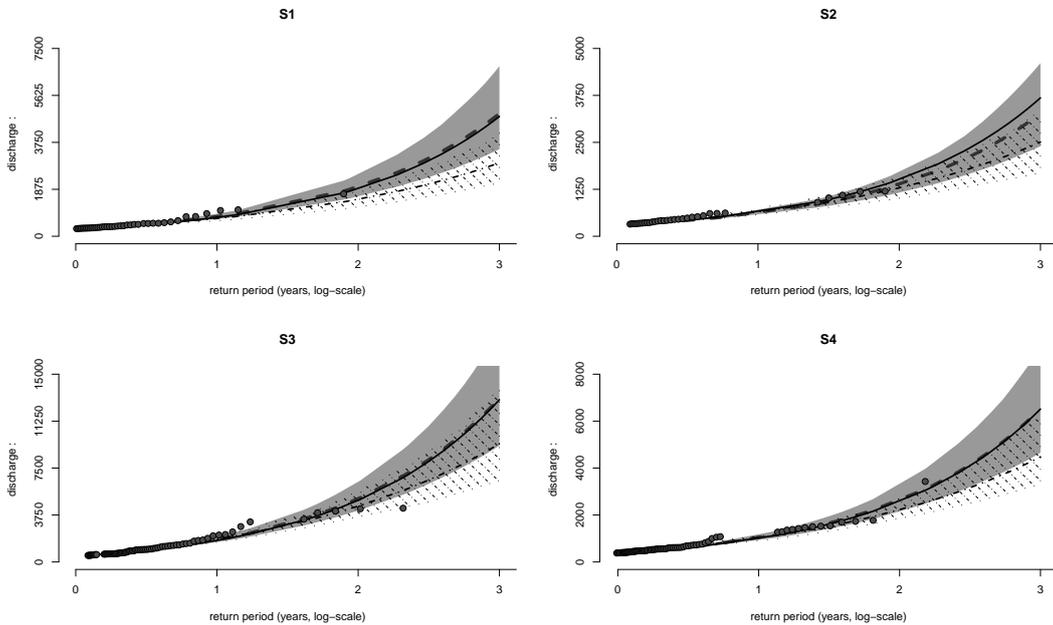


FIGURE 6. Return level plots: Quantile versus logarithm (base 10) of the return period at the four locations. Grey points: empirical return level of observed threshold excesses (simulated data); Black line: true curves; Gray dashed line and shaded area: posterior mean and 0.05 – 0.95-quantiles in the dependent model with Dirichlet mixture angular structure; Black dash-dotted line and area: *idem* in the independent model.

5.3. **Simulation study with 50 data sets.** The aim of this section is to verify
15 that the posterior distribution of the DM Poisson model parameters is reasonably informative, even with censored data. The procedure described above is applied to generate independently 50 data sets.

Marginal predictive performance. The model's ability to estimate of the probability of a marginal excess is first investigated. Large thresholds (V_1, \dots, V_4) are

specified so that their true marginal probability of exceedance is $P_0 = 1/(10 * 365)$, an approximate ten years return level. The quantities of interest are the posterior distributions of $\Delta_j(\theta) = \mathbb{P}_\theta(Y_j > V_j)$, which we hope to be concentrated around P_0 . Each posterior samples $(\theta_\iota)_{\iota \in 1..N}$ issued by a MCMC algorithm is transformed into a series of exceedance probabilities, $(\delta_j(\theta_\iota))_{\iota \in 1..N}$ which empirical distribution $\hat{F}_j = \frac{1}{N} \sum_{\iota=1}^N \mathbb{1}_{\delta_j(\theta_\iota)}(\cdot)$ approximates the posterior distribution of Δ_j . The performance of the posterior may then be investigated in terms of posterior quadratic loss,

$$\text{QL}(\hat{F}_j) = (\mathbb{E}_{\hat{F}_j}(\Delta_j) - P_0)^2 + \text{Var}_{\hat{F}_j}(\Delta_j).$$

This loss corresponds to the the predictive model choice criterion (PMCC) (Laud and Ibrahim, 1995). In the framework of scoring rules (Gneiting and Raftery, 2007) the PMCC is not a ‘proper score’ in general, but it is so when the true probability distribution of the quantity of interest Δ_j is a Dirac mass, which is the case here (Dirac mass at P_0).

After 1.10^6 MCMC iterations, at least one chain (out of six chains run in parallel for each simulated data set) passed the Heidelberger and Welch’s stationarity tests (Heidelberger and Welch, 1983) at level 10^{-4} , for each data set. Table 1 gathers, for the margins $j \in \{1, \dots, 4\}$, the mean and standard deviation of the QL scores, normalized by the (squared) true probability P_0^2 for readability,

$$\overline{\text{QL}}_j = \frac{\text{QL}(\hat{F}_j)}{P_0^2}.$$

$\overline{\text{QL}}_j$	1	2	3	4
mean	0.55	0.17	0.22	0.32
standard error	0.64	0.14	0.18	0.30
first quartile	0.14	0.08	0.08	0.10
third quartile	0.77	0.20	0.32	0.45

TABLE 1. Normalized scores $\overline{\text{QL}}_j$ for the marginal probability of an excess: mean and standard deviation, first and third quartiles over the 50 data sets. Column j corresponds to an excess at location j ($1 \leq j \leq 4$).

Although the variability of the scores (standard deviations) is relatively high, normalized third quartile less than one indicate that the posterior distribution concentrates in regions where the probability of a marginal excess is of the same order of magnitude as the true probability.

In order to verify that introducing a rather complex dependence structure model does not deteriorate the marginal estimates, the same quadratic loss score is computed with samples issued from the independent model. In view of Figure 7, displaying box-plots of the scores computed in both models (dependent *versus* independent), there is no significant difference between the two models in terms of marginal estimation.

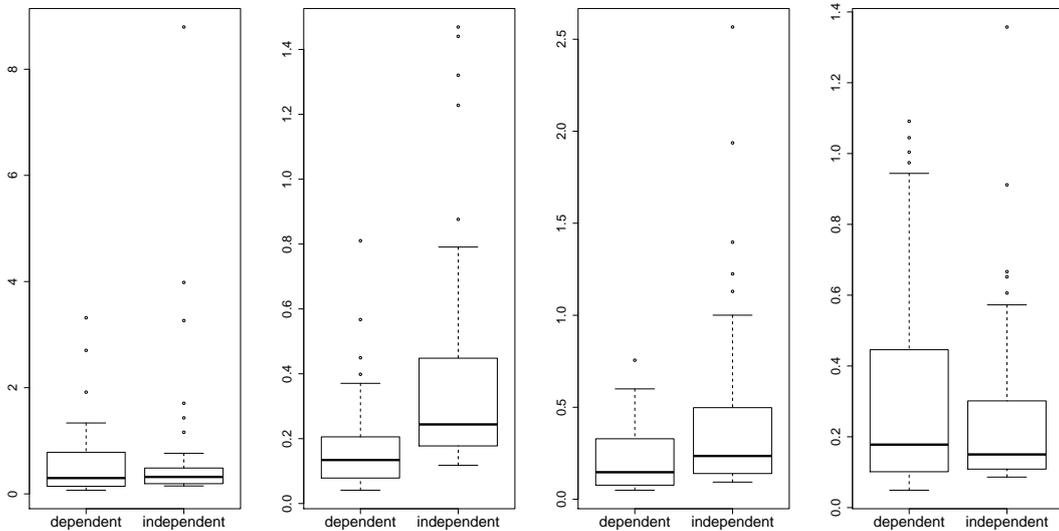


FIGURE 7. Normalized QR scores in the dependent Dirichlet-Poisson model versus the independent model, at the four different locations, for 50 simulated datasets

Estimation of the angular measure. To assess the performance of the DM Poisson model in terms of estimation of the dependence structure of extremes, a similar scoring procedure as above is followed, with quantities of interest Δ_j defined as the probability of a *joint* excess of the V_i 's, given a marginal excess of V_j ,

$$\tilde{\Delta}_j = \mathbb{P}(Y_1 \geq V_1, \dots, Y_4 \geq V_4 | Y_j \geq V_j).$$

These quantities do not have an explicit expression in the DM model, but are easily approached by standard Monte-Carlo sampling. Lacking a reference model in this context (the independent one is obviously unable to predict these quantities), only the scores in the dependent model are available. Table 2 summarizes the results in terms of normalized scores

$$\tilde{QL}_j = \frac{QL(\tilde{\Delta}_j)}{P_j^2},$$

where P_j is the true conditional probability of a joint excess.

A comforting fact is that the dependence structure seems to be even better estimated than the marginal distributions of extremes, in view of Tables 1 and 2.

5.4. Real data. Analyzing the hydrological data set presented in the introduction requires an additional declustering step: a multivariate run-declustering scheme was implemented. A cluster starts when one component exceeds the threshold. It and ends when, during τ consecutive days, all components are below their respective threshold, or (when censoring is present) with undertermined position, so that the observer can not ascertain that an excess occurred. The lag parameter $\tau = 3$ was set by considering stability regions of the estimates, and additional physical characteristics of the hydrological catchment, see Sabourin and Renard (2014) for details. After-wise the model is fitted to the extracted four-variate

$\tilde{Q}L_j$	1	2	3	4
mean	0.06	0.18	0.25	0.05
standard error	0.05	0.11	0.13	0.05
first quartile	0.02	0.11	0.17	0.01
third quartile	0.10	0.23	0.33	0.06

TABLE 2. Normalized QL scores $\tilde{Q}L_j$ for conditional probabilities of a joint excess: mean, standard deviation, first and third quartiles over the 50 simulated datasets. Column j corresponds to conditioning upon an excess at location j .

cluster maxima. Again, 6 parallel MCMC's of length 10^6 are run, with satisfactory convergence diagnostic after a $2 \cdot 10^5$ burn-in period. The posterior mean estimates of marginal parameters are close to the maximum likelihood estimates mentioned in subsection 5.1. Figure 8 displays bi-variate versions of the four-variate posterior predictive angular measure, together with point-wise 90% posterior credibility sets.

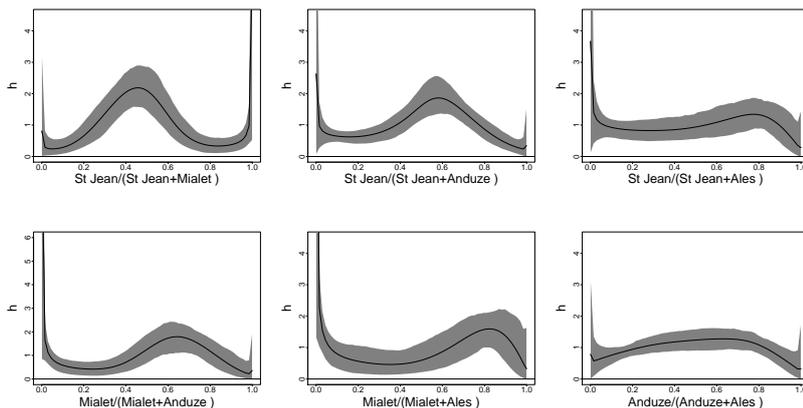


FIGURE 8. Bi-variate angular measure for the hydrological data set: posterior predictive density (black line) and posterior 0.05 and 0.95 quantiles.

The inferred dependence structure is rather complex, which fosters the use of such a flexible semi-parametric model. The posterior credible bounds of the angular density are relatively narrow around the posterior predictive in the central regions of the simplex (which is the segment $[0, 1]$ for bi-variate data) and indicate that asymptotic dependence is present. On the other hand, high levels (even unboundedness) of the predictive density near certain edges indicates a ‘weakly asymptotically dependent regime’, for the considered pair, *i.e.* a regime where one component may be large while the other one is not. Possibly, some ‘true’ asymptotic angular mass is concentrated on these edges, which translates in the Dirichlet mixture model (which only allows mass on the topological interior of the simplex) into high densities near the edges. The observed widening of posterior credible regions near the edges where the density is unbounded is not surprising : an unbounded density corresponds to Dirichlet parameters with components $\nu\mu_i < 1$,

for which a small variation of the parameter value induces a large variation of the density near the edges.

6. CONCLUSION

In this work, a flexible semi-parametric Bayesian inferential scheme is implemented to estimate the joint distribution of excesses above multivariate high thresholds, when the data are censored. A simulation example is designed on the same pattern as a real case borrowed from hydrology. Although the tuning of the MCMC algorithm requires some care, taking into account all kinds of observations for various censoring bounds allows to obtain satisfactory estimates, despite the loss of information relative to the angular structure induced by the censoring process. In particular, accurate enough estimations of quantities of interest such as marginal or conditional probabilities of an excess of a large threshold can be obtained. The main methodological novelty consists in taking advantage of the conditioning and marginalizing properties of the Dirichlet distributions, in order to simulate augmentation data which ‘replace’ the missing ones. Also, exponential terms in the likelihood with no explicit expressions are handled by sampling well chosen functionals of augmentation Poisson processes. This new inferential framework opens the road to statistical analysis of the extremes of data sets that would otherwise have been deemed unworkable.

ACKNOWLEDGMENTS

Part of this work has been supported by the EU-FP7 ACQWA Project (www.acqwa.ch), by the PEPER-GIS project, by the ANR-MOPERA project, by the ANR-McSim project and by the MIRACCLE-GICC project. The author would like to thank Benjamin Renard for providing the hydrological data that motivated this work and for his useful comments, and Anne-Laure Fougères and Philippe Naveau for their advice and interesting discussions we had during the writing of this paper.

APPENDIX A. POISSON LIKELIHOOD OF UNCENSORED DATA

Consider a Dirichlet mixture density $h = h_\psi$ as in (2.3). The density of the exponent measure in Cartesian coordinates is, using (2.7) and the expression of the Dirichlet density (2.2),

$$\frac{d\lambda_\psi(\mathbf{x})}{d\mathbf{x}} = d \sum_{m=1}^k \left\{ \frac{p_m \Gamma(\nu_m)}{\prod_{j=1}^d \Gamma(\nu_m \mu_{j,m})} \prod_{j=1}^d x_j^{\nu_m \mu_{j,m} - 1} \left(\sum_{j=1}^d x_j \right)^{-(\nu_m + 1)} \right\}. \quad (\text{A.1})$$

The likelihood expression (2.8) is simply that of a Poisson process on the region $[0, 1] \times A_{\mathbf{u}, n}$. Recall that, for a Poisson process with intensity μ on a region A , the likelihood of n points (Z_1, \dots, Z_n) observed in A is proportional to $e^{-\mu(A)} \prod_{i=1}^n \frac{d\mu}{dz}(Z_i)$.

The exponential term $e^{-n\lambda_\psi(A_{\mathbf{u}})}$ in (2.8) follows from the homogeneity property of λ_ψ ,

$$e^{-\ell([0,1]) \cdot \lambda_\psi(A_{\mathbf{u}, n})} = e^{-n\lambda_\psi(A_{\mathbf{u}})}.$$

The terms J_j^X are the inverse Jacobians of the marginal transformations $T_j^X : Y_j \mapsto X_j$, *i.e.*

$$J_j^X(y_j) = \sigma_j^{-1} (\zeta_j)^{-\xi_j} x_j^2 e^{\frac{1}{x_j}} \left[1 - e^{\frac{-1}{x_j}} \right]^{1+\xi_j},$$

where $x_j = \mathcal{T}_j^X(y_j)$,

APPENDIX B. INTEGRATION OF THE EXPONENT MEASURE ALONG DIRECTIONS OF MISSING COMPONENTS

The exact expression for the integral of the exponent measure's density in (3.2) along the axes $[0, \infty]$ corresponding to missing coordinates is given below. Let $\mathcal{D} = \{j_1, \dots, j_r\}$ be the non missing coordinates ($r < d$). Integrating $\frac{d\lambda_\psi}{d\mathbf{x}}$ over $(0, \infty)$ in the missing directions $\mathcal{D}_0 = i_1, \dots, i_{d-r}$ yields the marginal density of λ with respect to the Lebesgue measure on the vector space spanned by \mathcal{D} , *i.e.* $\frac{\partial \lambda_\psi(\mathbf{x})}{\partial x_{j_1} \cdots \partial x_{j_r}}$. With a DM angular measure, the integral has an analytic expression, which is, using (2.2) and (2.3),

$$\begin{aligned} \frac{\partial \lambda_\psi(\mathbf{x})}{\partial x_{j_1} \cdots \partial x_{j_r}} &= \int_{\{\mathbf{z}: z_j = x_j (j \in \mathcal{D}), z_i \in \mathbb{R}^+ (i \in \mathcal{D}_0)\}} \frac{d\lambda_\psi}{d\mathbf{x}}(\mathbf{z}) dz_{i_1}, \dots, dz_{i_{d-r}} \\ &= r \sum_{m=1}^k \left(\frac{p_m^0 \Gamma(\nu_m^0)}{\prod_{j \in \mathcal{D}} \Gamma(\nu_m^0 \mu_{j,m}^0)} \prod_{j \in \mathcal{D}} x_j^{\nu_m^0 \mu_{j,m}^0 - 1} \left(\sum_{j \in \mathcal{D}} x_j \right)^{-(\nu_m^0 + 1)} \right), \end{aligned} \quad (\text{B.1})$$

with

$$\begin{aligned} \nu_m^0 &= \nu_m (1 - \sum_{i \in \mathcal{D}_0} \mu_{i,m}), \quad \boldsymbol{\mu}_m^0 = (1 - \sum_{i \in \mathcal{D}_0} \mu_{i,m})^{-1} \boldsymbol{\mu}_m, \\ p_m^0 &= \frac{d}{r} (1 - \sum_{i \in \mathcal{D}_0} \mu_{i,m}) p_m. \end{aligned} \quad (\text{B.2})$$

This is the spectral measure associated with another angular DM distribution on \mathbf{S}_r with parameter $\psi^0 = (\nu_{1:k}^0, \boldsymbol{\mu}_{1:k}^0, p_{1:k}^0)$. The censored likelihood (3.2) can thus be re-written as

$$\begin{aligned} \mathcal{L}_{\mathbf{v}}(\mathbf{O}, \theta) &= \exp \left[-n_{\det} \lambda_\psi(A_{\mathbf{u}}) - \sum_{i=1}^{\mathcal{I}'} n'_i \lambda_\psi(A'_i) \right] \times \cdots \\ &\cdots \prod_{i=1}^{n_{\mathbf{v}}} \left\{ \int_{[\tilde{\mathbf{L}}_{t_i}^0, \tilde{\mathbf{R}}_{t_i}^0]} \frac{\partial^{r(i)} \lambda_\psi}{\partial x_{j_1(t_i)} \cdots \partial x_{j_r(i)(t_i)}}(\mathbf{x}) d\ell_i^0 \prod_{j: Y_{j,t_i} > v_j} J_j^X(y_{j,t_i}) \right\}, \end{aligned} \quad (\text{B.3})$$

where integration is performed in the censored, non-missing directions, and where ℓ_i^0 is the Lebesgue measure on the corresponding subspace of \mathbb{R}^d and $[\tilde{\mathbf{L}}_{t_i}^0, \tilde{\mathbf{R}}_{t_i}^0]$ are the original bounds $[\tilde{\mathbf{L}}_{t_i}, \tilde{\mathbf{R}}_{t_i}]$, modulo the subspace of missing components.

APPENDIX C. DATA AUGMENTATION DETAILS

Here is detailed the construction of augmentation data $\mathbf{Z} = (\mathbf{Z}_{\text{above}}, \mathbf{Z}')$, first introduced in Section 4.2.

C.1. **Definition of augmentation the variable $\mathbf{Z}_{\text{above}} = \{\mathbf{Z}_{t_i}\}_{i \leq n_v}$.** Consider a Fréchet-transformed, censored observation $\mathbf{C}_{t_i}^X = (C_1^X, \dots, C_d^X)$, with

$$C_j^X = (\tilde{\kappa}_{j,t_i}, X_{j,t_i}, \tilde{L}_{j,t_i}, \tilde{R}_{j,t_i}),$$

as in Section 3.3. Let $\mathcal{D}_c(i) = (j'_1(i), \dots, j'_c(i))$ be the censored coordinates in observation $\mathbf{C}_{t_i}^X$. The latent variables $\mathbf{Z}_{t_i} = (Z_{j'_1,t_i}, \dots, Z_{j'_c,t_i})$ are defined so as to
 5 ‘replace’ those coordinates: More formally, let

$$\bar{\mathbf{X}}_{t_i} | \psi \sim \frac{1}{\lambda_\psi(A_{\mathbf{u}})} \mathbb{1}_{A_{\mathbf{u}}}(\cdot) \lambda_\psi(\cdot), \quad (\text{C.1})$$

be an uncensored d -dimensional variable with Fréchet margins and dependence structure given by λ_ψ on $A_{\mathbf{u}}$. Then, \mathbf{Z}_{t_i} is defined through its joint distribution with the $\{X_{j,t_i} : j \notin \mathcal{D}_c(i)\}$, conditionally on θ ,

$$\left(\mathbf{Z}_{t_i}, \{X_{j,t_i} : j \notin \mathcal{D}_c(i)\} \right) | \theta \stackrel{\text{distribution}}{=} \bar{\mathbf{X}}_{t_i} | \psi.$$

10 Then, conditionally on the observation \mathbf{O}_{t_i} ,

$$\begin{aligned} [\mathbf{Z}_{t_i} | \mathbf{O}_{t_i}, \theta] &= [(\bar{X}_{j'_1,t_i}, \dots, \bar{X}_{j'_c,t_i}) | \mathbf{O}_{t_i}, \theta] \\ &= \left[(\bar{X}_{j'_1,t_i}, \dots, \bar{X}_{j'_c,t_i}) \middle| \bar{\mathbf{X}}_{t_i} \in [\tilde{\mathbf{L}}_{t_i}, \tilde{\mathbf{R}}_{t_i}], \psi \right]. \end{aligned} \quad (\text{C.2})$$

The contribution of the ‘augmented’ data point $(\mathbf{Z}_{t_i}, \mathbf{O}_{t_i})$ to the augmented likelihood (4.4) is

$$[\mathbf{z}_{t_i}, \mathbf{O}_{t_i} | \theta] = \frac{d\lambda_\psi(\bar{\mathbf{x}}_{t_i})}{dx} \prod_{j:\kappa_{j,t_i}=1} J_j^X(y_{j,t_i}), \quad (\text{C.3})$$

where

$$\bar{x}_{j,t_i} = \begin{cases} \mathcal{T}_j^X(y_{j,t_i}) & \text{if } \tilde{\kappa}_{j,t_i} = 1, \\ z_{j,t_i} & \text{otherwise.} \end{cases}$$

15 *Remark.* With missing components $\mathcal{D}_0(i) = \{j : \kappa_{j,t_i} = 0\} \neq \emptyset$, integration in the direction $\mathcal{D}_0(i)$ can be performed analytically (see Appendix B), which reduces the dimension of the augmented data. Indeed, in such a case, the corresponding Z_{j,t_i} ’s need not be included in \mathbf{Z}_{t_i} , the uncensored variable $\bar{\mathbf{X}}_{t_i}$ is defined on the quotient spaces $\mathbf{E}/\mathcal{D}_0(i)$ and its distribution is proportional to the exponent measure
 20 λ_{ψ^0} defined by equations (B.1) and (B.2), with density $\frac{\partial^{r(i)} \lambda_\psi}{\partial x_{j_1(t_i)} \dots \partial x_{j_{r(i)}(t_i)}}(\cdot)$. as in equation (B.3), so that

$$[\mathbf{z}_{t_i}, \mathbf{O}_{t_i} | \theta] = \frac{\partial^{r(i)} \lambda_\psi}{\partial x_{j_1(t_i)} \dots \partial x_{j_{r(i)}(t_i)}}(\bar{\mathbf{x}}_{t_i}) \prod_{j:\kappa_{j,t_i}=1} J_j^X(y_{j,t_i}), \quad (\text{C.4})$$

C.2. **Full conditional distribution of augmented data $[\mathbf{Z}_{\text{above}} | \mathbf{O}, \theta]$.** The full conditionals $[Z_{j,t_i} | \{Z_{s,t_i}\}_{s \neq j}, \mathbf{O}_{t_i}, \theta]$ are functions of truncated Beta distributions that can easily be sampled in a Gibbs step of the algorithm, as shown below.

In the remaining of this subsection, we omit the temporal index t_i . If ψ is a mixture of k Dirichlet distributions, as in (2.3), then, for any bounded, continuous function g defined on $[\tilde{L}_j, \tilde{R}_j]$, the conditional expectation of $g(Z_j)$ is, up to a multiplicative constant,

$$\begin{aligned}
\mathbb{E} [g(Z_j) | \mathbf{O}, \{Z_r\}_{r \neq j}, \theta] &= \mathbb{E} \left[g(\bar{X}_j) \mid \bar{X}_j \in [\tilde{L}_j, \tilde{R}_j], \bar{X}_r = x_r (r \neq j), \psi \right] \\
&= \int_{\tilde{L}_j}^{\tilde{R}_j} g(x_j) h_\psi \left(\frac{\mathbf{x}}{\sum_r x_r} \right) \left(\sum_r x_r \right)^{-(d+1)} dx_j \\
&= \sum_{m=1}^k p_m \underbrace{\int_{\tilde{L}_j}^{\tilde{R}_j} g(x_j) h_{\psi,m} \left(\frac{\mathbf{x}}{\sum_r x_r} \right) \left(\sum_r x_r \right)^{-(d+1)} dx_j}_{I_m},
\end{aligned} \tag{C.5}$$

(see equation (C.1) for the definition of \bar{X}_j).

Each term I_m ($m \leq k$ for a mixture of k components) is

$$\begin{aligned}
I_m &= \gamma_m \int_{\tilde{L}_j}^{\tilde{R}_j} g(x_j) \prod_{r \leq d} x_r^{\nu_m \mu_{r,m} - 1} \left(\sum_{r \leq d} x_r \right)^{-(d+1) - (\sum_{r \leq d} (\nu_m \mu_{r,m} - 1))} dx_j \\
&= \gamma_m \int_{\tilde{L}_j}^{\tilde{R}_j} g(x_j) \prod_{r \leq d} x_r^{\nu_m \mu_{r,m} - 1} \left(\sum_{r \leq d} x_r \right)^{-(\nu_m + 1)} dx_j \\
&= \gamma_m \rho_j \int_{\tilde{L}_j}^{\tilde{R}_j} g(x_j) x_j^{\nu_m \mu_{j,m} - 1} (s_j + x_j)^{-\nu_m - 1} dx_j,
\end{aligned}$$

where $\gamma_m = \frac{\Gamma(\nu_m)}{\prod_{r=1}^d \Gamma(\nu_m \mu_{r,m})}$, $s_j = \sum_{r \neq j} x_r$ and $\rho_j = \prod_{r \neq j} x_r^{\nu_m \mu_{r,m} - 1}$. Changing variable with $u = x_j / (x_j + s_j)$, the integration bounds are

$$R'_j = \frac{\tilde{R}_j}{s_j + \tilde{R}_j}, \quad L'_j = \frac{\tilde{L}_j}{s_j + \tilde{L}_j},$$

5 and we have

$$\begin{aligned}
I_m &= \gamma_m \rho_j \int_{L'_j}^{R'_j} g \left(\frac{s_j u}{1-u} \right) \left(\frac{s_j u}{1-u} \right)^{\nu_m \mu_{j,m} - 1} \left(s_j + \frac{s_j u}{1-u} \right)^{-\nu_m - 1} \frac{s_j}{(1-u)^2} du \\
&= \gamma_m \rho_j s_j^{-\nu_m (1 - \mu_{j,m}) - 1} \int_{L'_j}^{R'_j} g \left(\frac{s_j u}{1-u} \right) u^{\nu_m \mu_{j,m} - 1} (1-u)^{\nu_m (1 - \mu_{j,m})} du
\end{aligned}$$

One recognizes in the integrand the unnormalized density of a Beta random variable $U_m \sim \text{beta}(a_m, b_m)$, with

$$a_m = \nu_m \mu_{j,m}, \quad b_m = \nu_m (1 - \mu_{j,m}) + 1;$$

Let $\text{IB}_{a,b}(x, y)$ denote the incomplete Beta function (*i.e.* the integral of the Beta density) between truncation bounds x and y . The missing normalizing constant in
10 the integrand is

$$\begin{aligned}
D_m &= \frac{\Gamma(a_m + b_m)}{\Gamma(a_m) \Gamma(b_m) \text{IB}_{a_m, b_m}(L'_j, R'_j)} \\
&= \frac{\Gamma(\nu_m)}{\Gamma(\nu_m \mu_{j,m}) \Gamma(\nu_m (1 - \mu_{j,m}))} \frac{1}{\text{IB}_{a_m, b_m}(L'_j, R'_j) (1 - \mu_{j,m})}
\end{aligned}$$

Finally, we have

$$I_m = C_m \cdot D_m \int_{L'_j}^{R'_j} g\left(\frac{s_j u}{1-u}\right) u^{\nu_m \mu_{j,m} - 1} (1-u)^{\nu_m (1-\mu_{j,m})} du$$

with

$$C_m = (1 - \mu_{j,m}) \frac{\Gamma(\nu_m (1 - \mu_{j,m}))}{\prod_{s \neq j} \Gamma(\nu_m \mu_{s,m})} \text{IB}_{a_m, b_m}(L'_j, R'_j) \rho_j s_j^{-\nu_m (1 - \mu_{j,m}) - 1}, \quad (\text{C.6})$$

so that that the conditional expectation (C.5) is that of a mixture distribution,

$$\mathbb{E} [g(Z_j) | \mathbf{O}, \{Z_{r, r \neq j}\}] = \sum_{m=1}^j p'_m \mathbb{E} \left[g\left(\frac{s_j U_m}{1 - U_m}\right) \right]$$

with weights $(p'_m)_{m \leq k}$,

$$p'_m = p_m C_m, \quad (\text{C.7})$$

where C_m is given by equation (C.6).

- 5 As a conclusion, the conditional variable $[Z_j | \mathbf{O}, Z_{s \neq j}, \theta]$ is a mixture distribution of k components

$$\left(p'_m, V_{j,m} = \frac{s_j U_m}{1 - U_m} \right)_{1 \leq m \leq k}. \quad (\text{C.8})$$

In presence of missing coordinates, (C.8) still holds, up to replacing the mixture parameters $(\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\nu})$ with $(\mathbf{p}^0, \boldsymbol{\mu}^0, \boldsymbol{\nu}^0)$ as in equation (B.2).

- 10 **C.3. Augmentation Poisson process $\mathbf{Z}' = \{\{\mathbf{Z}'_i\}_{i \leq \mathcal{I}'}, \mathbf{Z}'_{\mathbf{u}}\}$ and weight function φ .** Let us define a region $E_{\mathbf{u}, n_{\text{det}}} = \{\mathbf{x} \in (\mathbb{R}^+)^d : \|\mathbf{x}\|_1 > \min_j(\frac{u_j}{n_{\text{det}}})\}$, so that $A_{\mathbf{u}, n_{\text{det}}} \subset E_{\mathbf{u}, n_{\text{det}}}$. Choose a multiplicative constant $\tau > 0$ and define a Poisson intensity measure $\lambda'(\cdot) = \tau \lambda_\psi(\cdot)$. The augmentation process $\mathbf{Z}'_{\mathbf{u}}$ is a Poisson processes which is defined together with $\varphi_{\mathbf{u}}$ by

$$\begin{cases} \mathbf{Z}'_{\mathbf{u}} \sim \text{PP}(\lambda') \text{ on } E_{\mathbf{u}, n_{\text{det}}}, \\ \varphi_{\mathbf{u}}(\mathbf{Z}'_{\mathbf{u}}) = (1 - 1/\tau)^{\mathbf{Z}'_{\mathbf{u}}(A_{\mathbf{u}, n_{\text{det}}})}, \end{cases} \quad (\text{C.9})$$

- 15 where $\mathbf{Z}'_{\mathbf{u}}(A_{\mathbf{u}, n_{\text{det}}})$ is the number of points forming $\mathbf{Z}'_{\mathbf{u}}$ which hit $A_{\mathbf{u}, n_{\text{det}}}$. Full justification and simulation details are given in the next subsection (Appendix C.4).

Let $\{\mathbf{x}'_{\mathbf{u}, s} = (r'_{\mathbf{u}, s}, \mathbf{w}'_{\mathbf{u}, s})\}_{s \leq N'_{\mathbf{u}}}$ be the points of $\mathbf{Z}'_{\mathbf{u}}$ in $E_{\mathbf{u}, n_{\text{det}}}$, the density of $\mathbf{Z}'_{\mathbf{u}}$ over $E_{\mathbf{u}, n_{\text{det}}}$, which contributes to the augmented likelihood (4.4), is

$$[\mathbf{Z}'_{\mathbf{u}} | \psi] = \frac{1}{N'_{\mathbf{u}}!} e^{-\frac{n_{\text{det}} \tau d}{\min_j \leq d u_j}} \prod_{s=1}^{N'_{\mathbf{u}}} \frac{\tau d}{(r'_{\mathbf{u}, s})^2} h_\psi(\mathbf{w}'_{\mathbf{u}, s}). \quad (\text{C.10})$$

- 20 The processes \mathbf{Z}'_i 's and the weights φ'_i 's are defined similarly, replacing n_{det} with n'_i and u_j with $\tilde{R}_{j, i}$ ($1 \leq i \leq \mathcal{I}'$).

C.4. Consistency of the augmentation model . The augmentation process \mathbf{Z}' and the weight function φ have been constructed with a hint towards using the general expression of the Laplace transform of a Poisson process to prove consistency of the augmented posterior.

Proposition 1. *Define the factors*

$$[\mathbf{z}_{t_i}, \mathbf{O}_{t_i} | \theta] \ (i \leq n_{\mathbf{v}}) \quad \text{and} \quad [\mathbf{z}'_{\mathbf{u}} | \psi], [\mathbf{z}'_i | \psi] \ (1 \leq i \leq \mathcal{I}')$$

composing the augmented likelihood (4.4), according to equations (C.3) and (C.10), and let the factors $\varphi_{\mathbf{u}}$, φ_i , $i \leq \mathcal{I}'$ of the weight function φ in (4.2) be defined by equation (C.9).

- 5 Then, the augmented posterior $[\mathbf{z}, \theta | \mathbf{O}]_+ \propto [\theta][\mathbf{z}, \mathbf{O} | \theta] \varphi(\mathbf{z})$ is consistent, in the sense that condition (4.3), whence (4.1), is satisfied.

Proof.

It is enough to show that, on the one hand,

$$\int [\mathbf{z}_{\text{above}}, \mathbf{O} | \theta] d\mathbf{z}_{\text{above}} = \prod_{i=1}^{n_{\mathbf{v}}} \left\{ \int_{[\tilde{\mathbf{L}}_{t_i}, \tilde{\mathbf{R}}_{t_i}]} \frac{d\lambda_{\psi}}{d\mathbf{x}} d\ell_i(\mathbf{x}) \prod_{j: y_{j,t_i} > v_j} J_j^X(y_{j,t_i}), \right\} \quad (\text{C.11})$$

- 10 and, on the other hand,

$$\begin{cases} \mathbb{E}[\varphi_{\mathbf{u}}(\mathbf{Z}'_{\mathbf{u}})] = \exp(-n_{\det} \lambda_{\psi}(A_{\mathbf{u}})), \\ \mathbb{E}[\varphi_i(\mathbf{Z}'_i)] = \exp(-n'_i \lambda_{\psi}(A'_i)) \end{cases} \quad (i \leq \mathcal{I}'). \quad (\text{C.12})$$

where the above expectations are taken with respect to $[\mathbf{Z}'_{\mathbf{u}} | \psi]$ and $[\mathbf{Z}'_i | \psi]$.

Establishing (C.11) is immediate: from definition (C.3) of the $[\mathbf{z}_{t_i}, \mathbf{O} | \theta]$'s composing $[\mathbf{z}_{\text{above}}, \mathbf{O} | \theta]$,

$$\int_{[\tilde{\mathbf{L}}_{t_i}, \tilde{\mathbf{R}}_{t_i}]} [\mathbf{z}_{t_i}, \mathbf{O}_{t_i} | \theta] d\mathbf{z}_{t_i} = \int_{[\tilde{\mathbf{L}}_{t_i}, \tilde{\mathbf{R}}_{t_i}]} \frac{d\lambda_{\psi}}{d\mathbf{x}} d\ell_i(\mathbf{x}) \prod_{j: y_{j,t_i} > v_j} J_j^X(y_{j,t_i}),$$

which yields (C.11) by taking the product over indices $1 \leq i \leq n_{\mathbf{v}}$.

- 15 It remains to show (C.12). To wit, $\varphi_{\mathbf{u}}$ is a smoothed version of the indicator function $\mathbb{1}_{\{\mathbf{N}'_{\mathbf{u}}(A_{\mathbf{u}, n_{\det}}) = 0\}}$, which expectancy is $\mathbb{P}(\mathbf{N}'_{\mathbf{u}}(A_{\mathbf{u}, n_{\det}}) = 0) = e^{-n_{\det} \lambda_{\psi}(A_{\mathbf{u}})}$, as soon as $\mathbf{N}'_{\mathbf{u}}$ is a Poisson process with intensity measure λ_{ψ} . For f a bounded, continuous function defined on a nice space E and N a point process on E , denote

$$N(f) = \int_E f dN = \sum_{i=1}^{N(E)} f(s_i).$$

- 20 Then, if N is a Poisson process $\text{PP}(\lambda)$ on E , the Laplace transform $\text{Lap}_N(f) \triangleq \mathbb{E}(e^{-N(f)})$ is (Resnick, 1987, Chap. 3)

$$\text{Lap}_N(f) = \exp\left(-\int_E (1 - e^{-f(s)}) d\lambda(s)\right).$$

Consider the region $E = E_{\mathbf{u}, n_{\det}}$ as in (C.9) and take

$$f_{\mathbf{u}}(\mathbf{x}) = -\log(1 - 1/\tau) \mathbb{1}_{A_{\mathbf{u}, n_{\det}}}(\mathbf{x}), \quad \text{so that} \quad 1 - e^{-f_{\mathbf{u}}} = \frac{1}{\tau} \mathbb{1}_{A_{\mathbf{u}, n_{\det}}}.$$

With these notations,

$$\varphi_{\mathbf{u}}(\mathbf{Z}'_{\mathbf{u}}) \triangleq \left(1 - \frac{1}{\tau}\right)^{\mathbf{Z}'_{\mathbf{u}}(A_{\mathbf{u}, n_{\det}})} = \exp(-\mathbf{Z}'_{\mathbf{u}}(f_{\mathbf{u}})),$$

whence

$$\begin{aligned}
\mathbb{E}(\varphi_{\mathbf{u}}(\mathbf{Z}'_{\mathbf{u}})) &= \text{Lap}_{\mathbf{Z}'_{\mathbf{u}}}(f_{\mathbf{u}}) \\
&= \exp\left(-\int_{E_{\mathbf{u},n_{\text{det}}}} (1 - e^{-f_{\mathbf{u}}}) d\lambda'\right) \\
&= \exp\left(-\int_{E_{\mathbf{u},n_{\text{det}}}} \frac{1}{\tau} \mathbb{1}_{A_{\mathbf{u},n_{\text{det}}}} d(\tau\lambda_{\psi})\right) \\
&= \exp(-\lambda_{\psi}(A_{\mathbf{u},n_{\text{det}}})) .
\end{aligned}$$

This shows the first equality in (C.12). The second one is derived with a similar argument. \square

The points of $\mathbf{Z}'_{\mathbf{u}}$ can easily be simulated (see Resnick, 1987, Chap.3): the number of points $N'_{\mathbf{u}}$ in $E_{\mathbf{u},n_{\text{det}}}$ is a Poisson random variable with mean equal to

$$\lambda'(E_{\mathbf{u},n_{\text{det}}}) = \frac{\tau d}{\min_j(u_j / n_{\text{det}})},$$

and each point has density in polar coordinates equal to $\frac{1}{\lambda'(E_{\mathbf{u},n_{\text{det}}})} \frac{\tau d}{r^2} h_{\psi}(\mathbf{w})$.

Remark. One may be tempted to define $\mathbf{Z}'_{\mathbf{u}}$ as a Poisson process with intensity λ_{ψ} on some $E \supset A_{\mathbf{u},n_{\text{det}}}$, and $\varphi(\mathbf{Z}_{\mathbf{u}})$ as the indicator $\mathbb{1}_{\mathbf{z}'_{\mathbf{u}}(A_{\mathbf{u},n_{\text{det}}})=0}$, with a similar definition for the φ'_i 's and the \mathbf{Z}'_i 's. As pointed out in the proof of Proposition 1, one would have $\mathbb{E}[\mathbb{1}_{\mathbf{z}'_{\mathbf{u}}(A_{\mathbf{u},n_{\text{det}}})=0}] = \exp(-n_{\text{det}} \lambda_{\psi}(A_{\mathbf{u}}))$, as required. However, even if this construction is valid in theory, it leads to a very large rate of rejection in the Metropolis algorithm: $\varphi(\mathbf{Z}'_{\mathbf{u}})$ has too much variability around its mean value and the proposal is systematically rejected each time a point in the augmentation process hits the failure region.

C.5. Expression of the augmented posterior. Recall from Section 4.2, equation (4.2), that the augmented posterior density to be sampled by the MCMC algorithm is

$$[\mathbf{z}, \theta | \mathbf{O}]_+ \propto [\theta][\mathbf{z}, \mathbf{O} | \theta] \varphi(\mathbf{z}).$$

Combining equations (C.3) and (C.9), and integrating out missing components as in Appendix B, the developed expression is

$$\begin{aligned}
[\mathbf{z}, \theta | \mathbf{O}]_+ &\propto [\theta] \left([\mathbf{z}'_{\mathbf{u}} | \psi] \prod_{i=1}^{I'} [\mathbf{z}'_i | \psi] \right) \cdot (1 - 1/\tau)^{\mathbf{z}'_{\mathbf{u}}(A_{\mathbf{u},n_{\text{det}}}) + \sum_{i=1}^{I'} \mathbf{z}'_i(A'_{i,n'_i})} \dots \\
&\dots \prod_{i=1}^{n_{\mathbf{v}}} \left\{ \frac{\partial^{r(i)} \lambda_{\psi}}{\partial x_{j_1(i)} \dots \partial x_{j_{r(i)}(i)}}(\bar{\mathbf{x}}_{t_i}) \prod_{j: Y_{j,t} > v_j} J_j^X(y_{j,t}) \right\}, \quad (\text{C.13})
\end{aligned}$$

where $[\mathbf{z}'_{\mathbf{u}} | \psi]$ and the $[\mathbf{z}'_i | \psi]$'s are given by equation (C.10).

APPENDIX D. MCMC ALGORITHM

The MCMC algorithm generates a sample $(\theta_{\iota}, \mathbf{Z}_{\iota})_{\iota=1, \dots, N}$ which distribution converges to the invariant distribution of the chain, which is the augmented posterior distribution $[\mathbf{Z}, \theta | \mathbf{O}]_+$, as defined in Section 4 and Appendix C.5. The quantity of interest here is the joint parameter θ , which is the concatenation of the marginal parameters and the dependence parameter: $\theta = (\chi, \psi)$. We recall that MCMC

algorithms aiming at sampling a quantity $\Delta \in E$ according to a density $\pi(\cdot)$ proceed typically as follows

- Start with any value $\delta(0) \in E$
- for $\iota \in \{1, \dots, N\}$
 - 5 (1) Generate δ^* according to a *proposal distribution* with density $q(\delta(\iota), \cdot)$
 - (2) Compute the *acceptance ratio*

$$\alpha = \frac{\pi(\delta^*) q(\delta^*, \delta(\iota))}{\pi(\delta(\iota)) q(\delta(\iota), \delta^*)},$$
and generate U , a uniform random variable on $[0, 1]$.
 - (3) If $U > \alpha$, ‘reject’ the proposal and set
$$\delta(\iota + 1) = \delta(\iota).$$
Otherwise (*i.e.* with probability α), set
$$\delta(\iota + 1) = \delta^*$$
 - 10 (4) Set $\iota = \iota + 1$, go to (1).
- Return $(\delta(\iota_{\min}), \dots, \delta(N))$,

where ι_{\min} is the length of the *burn-in period*, after which the chain is deemed to have reached a stationary behavior. In our case, the unnormalized posterior measure $[\mathbf{Z}, \theta | \mathbf{O}]_+$ (see equation (C.13)) on the augmented parameter space plays
15 the role of the objective density π above.

In a *Metropolis-within-Gibbs* MCMC, several proposal kernels q_1, \dots, q_T are defined, each of them corresponding to a *type of move*, which is randomly chosen among $\{1, \dots, T\}$ at each iteration ι and allows to modify some subset of components in δ alone. The algorithm developed in this paper builds on the MCMC
20 algorithm proposed by Sabourin and Naveau (2014), in which several *types of move* modifying the dependence structure (Dirichlet mixture parameter ψ) have been defined. Those are kept as they are in the present work, the novel part of which concerns the definition of *marginal moves* (modifying the marginal parameter χ) and *augmentation moves* (modifying the augmentation data \mathbf{Z}). Additional
25 notations distinguishing between the quantities appearing in the general MCMC algorithm above, according to the type of move, are omitted in the remainder of this section.

D.1. Starting values. In a preliminary step, likelihood optimization is performed in the independent model (the likelihood for one multivariate observation is the
30 product of d Pareto densities). This provides starting values for the marginal parameters as well as a Hessian matrix \mathcal{H} , that may be used as the inverse of a reference covariance matrix when updating the marginal parameters.

D.2. Marginal moves. The marginal parameter χ is updated as a block: The proposal is normal, with mean at $\chi(\iota)$ and co-variance matrix $\Sigma = \delta \mathcal{H}^{-1}$, where δ
35 is a scaling factor fixed by the user, that may typically be set around 0.5 and \mathcal{H} is the Hessian matrix computed in the preliminary step. Since the proposal density is symmetric, and since this move does not modify the dependence structure, neither the terms involving the proposal density, nor the point processes \mathbf{Z}' , appear in acceptance ratio. The augmented variables $\mathbf{Z}_{\text{above}}$ are left unchanged. If any augmented component \mathbf{Z}_{j,t_i} is outside of the candidate censoring interval $[\tilde{\mathbf{L}}_{t_i}^*, \tilde{\mathbf{R}}_{t_i}^*]$ (on
40 the new Fréchet scale) resulting from the modification of the marginal parameters,

the move is rejected, since in such a case, the candidate has augmented likelihood $[\mathbf{z}, \mathbf{O} | \theta^*] = 0$, *i.e.* $\alpha = 0$.

Otherwise, the uncensored Fréchet-transformed variables (such that $\tilde{\kappa}_{j,t_i} = 1$) are updated to $\mathbf{X}_{t_i}^* = \mathcal{T}_{j,t_i}^{\chi^*}(Y_{j,t_i})$. The acceptance ratio is

$$\alpha = \frac{[\chi^*]}{[\chi(\ell)]} \prod_{i=1}^{n_v} \left\{ \frac{\partial^{r(i)} \lambda_\psi}{\partial x_{j_1(i)} \cdots \partial x_{j_{r(i)}(i)}}(\bar{\mathbf{X}}_{t_i}^*) \left[\frac{\partial^{r(i)} \lambda_\psi}{\partial x_{j_1(i)} \cdots \partial x_{j_{r(i)}(i)}}(\bar{\mathbf{X}}_{t_i}(\ell)) \right]^{-1} \prod_{j: Y_{j,t_i} > v_j} \frac{J_j^{\chi^*}(Y_{j,t_i})}{J_j^{\chi(\ell)}(Y_{j,t_i})} \right\},$$

- 5 where $j_1, \dots, j_{r(i)}$ are the non-missing components in the censored observation $\mathbf{C}_{t_i}^X$ and $\bar{X}_j = X_j$ for uncensored components, $\bar{X}_{j,t_i} = Z_{j,t_i}$ otherwise (see Appendix B).

D.3. Augmentation moves for $\mathbf{Z}_{\text{above}}$. The augmented components $\{Z_{j,t_i}\} = \{\bar{\mathbf{X}}_{j,t_i} : \tilde{\kappa}_j \in \{2, 3\}\}$ (*c.f.* Section 4.2) are re-sampled, one coordinate at a time, from their exact conditional distribution given the other coordinates, as derived
10 in Appendix C.1. Since no other component of (\mathbf{Z}, θ) is modified, the proposal density equals the objective density, and the acceptance ratio is thus set to $\alpha = 1$.

D.4. Augmentation moves for \mathbf{Z}' . During this move, proposals

$$\mathbf{Z}'^* = \{\mathbf{Z}'_{\mathbf{u}}, \mathbf{Z}'_i, i \leq \mathcal{I}'\},$$

for the augmentation Poisson processes introduced in the end of Section 3.4, are sampled under their exact distribution,

$$q(\mathbf{Z}'(\ell), \mathbf{Z}'^*) = [\mathbf{Z}'^* | \psi(\ell)].$$

- 15 The latter is determined by their intensity measure $\lambda' = \tau \lambda_\psi$: the multiplicative constant τ and the sampling procedure have been described in Appendix C.3. The acceptance ratio is thus

$$\alpha = \frac{[\mathbf{Z}'^*, \theta(\ell) | \mathbf{O}]_+ [\mathbf{Z}'(\ell) | \psi(\ell)]}{[\mathbf{Z}'^*, \theta(\ell) | \mathbf{O}]_+ [\mathbf{Z}'^* | \psi(\ell)]}$$

all the terms cancel out except the ratio $\varphi(\mathbf{Z}'^*)/\varphi(\mathbf{Z}'(\ell))$, so that

$$\alpha = (1 - 1/\tau)^{[(\mathbf{Z}'_{\mathbf{u}})^*(A_{\mathbf{u},n_{\text{det}}}) - \mathbf{Z}'_{\mathbf{u}}(\ell)(A_{\mathbf{u},n_{\text{det}}})] + \sum_{i \leq \mathcal{I}'} [(\mathbf{Z}'_i)^*(A'_{i,n'_i}) - \mathbf{Z}'_i(\ell)(A'_{i,n'_i})]}.$$

D.5. Dependence moves. These types of moves allow to update $\psi(\ell)$. The only
20 difference between the present algorithm and what is described is Sabourin and Naveau (2014) is that not enough exact angular data are available to construct proposals for moving or splitting a Dirichlet mixture components $\boldsymbol{\mu}_m$. Indeed, most of the observations have at least one coordinate missing or censored, so that no ‘angle’ is available. Consequently, the latter proposal is a simple Dirichlet
25 distribution with mode at $\boldsymbol{\mu}_m(\ell)$, with re-centering parameter $0 < \epsilon < 0.5$,

$$q(\boldsymbol{\mu}_m(\ell), \cdot) = \text{diri}_{\frac{d}{\epsilon}, \gamma^*}(\cdot),$$

with $\gamma^* = (1 - \epsilon) \boldsymbol{\mu}_m + \epsilon(\frac{1}{d}, \dots, \frac{1}{d})$.

Each dependence move (except for a *shuffling move* which only affects the representation of the angular distribution, see Sabourin and Naveau (2014)) is systematically followed by an augmentation move updating the Poisson processes \mathbf{Z}' ,

which improves the chain’s mixing properties. This also avoids the computation of the ‘costly’ term involving the density $[\mathbf{z}'|\psi]$ (see equation (C.10)). Indeed, the acceptance ratio for the two consecutive moves (dependence move followed by a augmentation move) is

$$\alpha = \frac{[\psi^*]}{[\psi(\iota)]} \frac{q(\psi^*, \psi(\iota))}{q(\psi(\iota), \psi^*)} \times \dots$$

$$\dots (1 - 1/\tau) \left\{ [\mathbf{Z}'_{\mathbf{u}}(A_{\mathbf{u}, n_{\text{det}}}) - \mathbf{Z}'_{\mathbf{u}}(\iota)(A_{\mathbf{u}, n_{\text{det}}})] + \sum_{i \leq \mathcal{I}'} [\mathbf{Z}'_i(A'_{i, n'_i}) - \mathbf{Z}'_i(\iota)(A'_{i, n'_i})] \right\} \times \dots$$

$$\dots \prod_{i=1}^{n_{\mathbf{v}}} \left\{ \frac{\partial^{r(i)} \lambda_{\psi^*}}{\partial x_{j_1(i)} \cdots \partial x_{j_r(i)}(i)}(\bar{\mathbf{X}}_{t_i}(\iota)) \left[\frac{\partial^{r(i)} \lambda_{\psi(\iota)}}{\partial x_{j_1(i)} \cdots \partial x_{j_r(i)}(i)}(\bar{\mathbf{X}}_{t_i}(\iota)) \right]^{-1} \right\}.$$

5 D.6. MCMC settings and convergence diagnostics in the simulation study.

For the simulation study, the prior on the Dirichlet mixture distributions is specified in a similar way as in Sabourin and Naveau (2014). The number k of mixture components has truncated geometric distribution, $[k] \propto (1 - \frac{1}{\lambda})^{k-1} \frac{1}{\lambda} \mathbb{1}_{[1, k_{\text{max}}]}(k)$ with upper bound $k_{\text{max}} = 10$ and mean parameter $\lambda = 4$. Also, for the sake of
10 simplicity, all the marginal parameters are assumed to be *a priori* independent, with normal distributions (after log-transformation of the scales). The shape parameter has standard normal distribution and the logarithms of the scales have mean and standard deviation both equal to 5.

As for the augmentation Poisson process data, the multiplicative constant τ
15 involved in the Poisson intensity is set to 50. It appeared that smaller values of τ (close to 1) considerably affected the mixing properties of the chains.

Convergence of the dependence parameters $\psi(\iota)$ can be monitored using functionals based on integration of the simulated densities against Dirichlet test functions (see Sabourin and Naveau, 2014, for details). To detect possible mixing
20 defects, six chains of 10^6 iterations each are run in parallel. Standard convergence diagnostic tests are implemented in R (Heidelberger and Welch, 1983; Gelman and Rubin, 1992), respectively testing for non-stationarity and poor mixing. For example, the stationarity test detects three non-stationary chains out of six for the simulated data set exemplified in Section 5.2. The mixing properties of the three
25 retained ones, as measured by a variance ratio inter/intra chains, are satisfactory enough: all the potential scale reduction factors (Gelman and Rubin, 1992) are below 1.1. The same is true of the marginal parameter component of the chains, $(\chi(\iota))_{\iota}$.

REFERENCES

- 30 Beirlant, J., Goegebeur, Y., Segers, J., and Teugels, J. (2004). *Statistics of extremes: Theory and applications*. John Wiley & Sons: New York.
- Boldi, M.-O. and Davison, A. C. (2007). A mixture model for multivariate extremes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(2):217–229.
- 35 Coles, S. (2001). *An introduction to statistical modeling of extreme values*. Springer Verlag.
- Coles, S. and Tawn, J. (1991). Modeling extreme multivariate events. *JR Statist. Soc. B*, 53:377–392.

- Davison, A. and Smith, R. (1990). Models for exceedances over high thresholds. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 393–442.
- Einmahl, J., de Haan, L., and Piterbarg, V. (2001). Nonparametric estimation of the spectral measure of an extreme value distribution. *The Annals of Statistics*, 29(5):1401–1423.
- Einmahl, J. and Segers, J. (2009). Maximum empirical likelihood estimation of the spectral measure of an extreme-value distribution. *The Annals of Statistics*, 37(5B):2953–2989.
- Fougères, A.-L., Nolan, J. P., and Rootzén, H. (2009). Models for dependent extremes using stable mixtures. *Scandinavian Journal of Statistics*, 36(1):42–59.
- Gelman, A. and Rubin, D. (1992). Inference from iterative simulation using multiple sequences. *Statistical science*, pages 457–472.
- Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.
- Gómez, G., Calle, M. L., and Oller, R. (2004). Frequentist and bayesian approaches for interval-censored data. *Statistical Papers*, 45(2):139–173.
- Guillette, S., Perron, F., and Segers, J. (2011). Non-parametric bayesian inference on bivariate extremes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- Gumbel, E. (1960). Distributions des valeurs extrêmes en plusieurs dimensions. *Publ. Inst. Statist. Univ. Paris*, 9:171–173.
- Heidelberger, P. and Welch, P. (1983). Simulation run length control in the presence of an initial transient. *Operations Research*, pages 1109–1144.
- Hosking, J. R. M. and Wallis, J. R. (2005). *Regional frequency analysis: an approach based on L-moments*. Cambridge University Press.
- Huser, R., Davison, A. C., and Genton, M. G. (2014). A comparative study of likelihood estimators for multivariate extremes. *arXiv preprint arXiv:1411.3448*.
- Laud, P. W. and Ibrahim, J. G. (1995). Predictive model selection. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 247–262.
- Ledford, A. and Tawn, J. (1996). Statistics for near independence in multivariate extreme values. *Biometrika*, 83(1):169–187.
- Neppel, L., Renard, B., Lang, M., Ayrat, P., Coeur, D., Gaume, E., Jacob, N., Payrastre, O., Pobanz, K., and Vinet, F. (2010). Flood frequency analysis using historical data: accounting for random and systematic errors. *Hydrological Sciences Journal—Journal des Sciences Hydrologiques*, 55(2):192–208.
- Pickands, J. I. (1975). Statistical inference using extreme order statistics. *the Annals of Statistics*, pages 119–131.
- Resnick, S. (1987). *Extreme values, regular variation, and point processes, volume 4 of Applied Probability. A Series of the Applied Probability Trust*. Springer-Verlag, New York.
- Resnick, S. (2007). *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. Springer Series in Operations Research and Financial Engineering.
- Sabourin, A. and Naveau, P. (2014). Bayesian dirichlet mixture model for multivariate extremes: A re-parametrization. *Computational Statistics & Data Analysis*, 71(0):542 – 567.

- Sabourin, A. and Renard, B. (2014). Combining regional estimation and historical floods: a multivariate semi-parametric peaks-over-threshold model with censored data.
- Schmedler, W. (2005). Likelihood estimation for censored random vectors. *Economic Reviews*, 24(2):195–217.
- 5 Smith, R. (1994). Multivariate threshold methods. *Extreme Value Theory and Applications*, 1:225–248.
- Smith, R., Tawn, J., and Coles, S. (1997). Markov chain models for threshold exceedances. *Biometrika*, 84(2):249–268.
- 10 Stephenson, A. (2003). Simulating multivariate extreme value distributions of logistic type. *Extremes*, 6(1):49–59.
- Stephenson, A. (2009). High-dimensional parametric modelling of multivariate extreme events. *Australian & New Zealand Journal of Statistics*, 51(1):77–88.
- Tanner, M. and Wong, W. (1987). The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association*, 82(398):528–540.
- 15 Thibaud, E. and Opitz, T. (2013). Efficient inference and simulation for elliptical pareto processes. *arXiv preprint arXiv:1401.0168*.
- Tierney, L. (1994). Markov chains for exploring posterior distributions. *the Annals of Statistics*, pages 1701–1728.
- 20 Van Dyk, D. and Meng, X. (2001). The art of data augmentation. *Journal of Computational and Graphical Statistics*, 10(1):1–50.