



HAL
open science

Multivariate sparse clustering for extremes

Nicolas Meyer, Olivier Wintenberger

► **To cite this version:**

Nicolas Meyer, Olivier Wintenberger. Multivariate sparse clustering for extremes. 2021. hal-02904347v2

HAL Id: hal-02904347

<https://hal.science/hal-02904347v2>

Preprint submitted on 17 Mar 2021 (v2), last revised 6 Dec 2022 (v3)

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.

Multivariate sparse clustering for extremes

Nicolas Meyer^{*1} and Olivier Wintenberger^{†1}

¹*Sorbonne Université, LPSM, F-75005, Paris, France*

March 17, 2021

Abstract

Identifying directions where extreme events occur is a major challenge in multivariate extreme value analysis. In this paper, we use the concept of sparse regular variation introduced by Meyer and Wintenberger (2021) to infer the tail dependence of a random vector \mathbf{X} . This approach relies on the Euclidean projection onto the simplex which better exhibits the sparsity structure of the tail of \mathbf{X} than the standard methods. Our procedure based on a rigorous methodology aims at capturing clusters of extremal coordinates of \mathbf{X} . It also includes the identification of a threshold above which the values taken by \mathbf{X} are considered as extreme. We provide an efficient and scalable algorithm called MUSCLE and apply it on numerical experiments to highlight the relevance of our findings. Finally we illustrate our approach with wind speed data and financial return data.

Keywords: Euclidean projection onto the simplex, model selection, multivariate extremes, regular variation, sparse regular variation

1 Introduction

The aim of this article is to study the tail dependence of a random vector $\mathbf{X} \in [0, \infty)^d$. In this context it is customary to assume that \mathbf{X} is regularly varying (Resnick (1987), Beirlant et al. (2006), Resnick (2007), Hult and Lindskog (2006)). Under this assumption we have

$$\mathbb{P}((|\mathbf{X}|/t, \mathbf{X}/|\mathbf{X}|) \in \cdot \mid |\mathbf{X}| > t) \xrightarrow{w} \mathbb{P}((Y, \Theta) \in \cdot), \quad t \rightarrow \infty, \quad (1.1)$$

where \xrightarrow{w} denotes weak convergence, and where Θ is a random vector on the positive unit sphere $\{\mathbf{x} \in [0, \infty)^d : |\mathbf{x}| = 1\}$ independent of the random variable Y which satisfies $\mathbb{P}(Y > y) = y^{-\alpha}$, $y > 1$, for $\alpha > 0$. The parameter α is called the *tail index*. It highlights the intensity of the extremes. The smaller this index is, the heaviest the tail of \mathbf{X} is likely to be. The random vector Θ is called the spectral vector and its distribution $\mathbb{P}(\Theta \in \cdot)$ the spectral measure. Its support indicates the directions supported by the large events. The subspaces of the positive unit sphere on which the spectral vector puts mass correspond to the directions where large events are likely to appear. Note that the choice of the norm in Equation (1.1) is arbitrary. In this article we choose the ℓ^1 -norm and thus focus on the simplex $\mathbb{S}_+^{d-1} := \{\mathbf{x} \in [0, \infty)^d : x_1 + \dots + x_d = 1\}$.

In order to study the support of the spectral measure it is convenient to partition the simplex in terms of the nullity of some coordinates (Chautru (2015), Goix et al. (2017), Simpson et al. (2019)). For $\beta \subset \{1, \dots, d\}$ the subspaces C_β are defined as

$$C_\beta = \{\mathbf{x} \in \mathbb{S}_+^{d-1} : x_i > 0 \text{ for } i \in \beta, x_i = 0 \text{ for } i \notin \beta\}. \quad (1.2)$$

^{*}*meyer@math.ku.dk* (corresponding author)

[†]*olivier.wintenberger@upmc.fr*

An illustration of these subsets in dimension 3 is given in Figure 1. This partition highlights the extremal structure of \mathbf{X} . For a given $\beta \subset \{1, \dots, d\}$ the inequality $\mathbb{P}(\Theta \in C_\beta) > 0$ implies that the marginals X_j , $j \in \beta$, are likely to take simultaneously large values while the ones for $j \in \beta^c$ are of smaller order. Hence the identification of clusters of directions β which concentrate the mass of the spectral measure brings out groups of coordinates which can be large together.

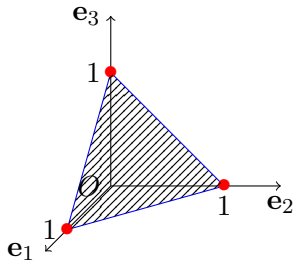


Figure 1: The subspaces C_β in dimension 3 for the ℓ^1 -norm. In red the vectors \mathbf{e}_1 , \mathbf{e}_2 , and \mathbf{e}_3 . In blue the subsets $C_{\{i,j\}}$ for $i, j \in \{1, 2, 3\}$, $i \neq j$. The shaded part corresponds to the interior of the simplex.

Highlighting such groups of coordinates is at the core of several recent papers on multivariate extremes, all of them relying on some hyperparameters (Chiapino and Sabourin (2016), Goix et al. (2017), Chiapino et al. (2019), Simpson et al. (2019)). This approach faces a crucial issue: the difference of support between Θ and $\mathbf{X}/|\mathbf{X}|$. Indeed, the spectral measure is likely to place mass on low-dimensional subspaces C_β , $\beta \neq \{1, \dots, d\}$. We say that this measure is *sparse* when the number of coordinates in the associated clusters β is small. Conversely, the distribution of the self-normalized vector $\mathbf{X}/|\mathbf{X}|$ only concentrates on the central subspace $C_{\{1, \dots, d\}}$ when \mathbf{X} models real-world data. Besides, all the existing approaches in the literature rely on nonstandard regular variation for which $\alpha = 1$ and all marginals are tail equivalent, possibly after a standardization. However, sparsity arises all the more for standard regular variation (1.1). In this case it is possible that the marginals of \mathbf{X} are not tail equivalent and thus that the support of the spectral measure is included in \mathbb{S}_+^{r-1} for $r \ll d$. This is the approach we use in this article. For a comparison of standard and nonstandard regular variation we refer to Resnick (2007), Section 6.5.6.

In this article we aim at providing a method which highlights the sparsity of the tail structure by exhibiting sparse clusters of extremal directions. By sparse clusters we mean groups of coordinates which contain a reduced number of directions compared to d . We refer to this method as *sparse clustering*. The statistical procedure we propose to achieve this clustering relies on the framework of Meyer and Wintenberger (2021) which allows to circumvent the estimation's issue that arises with the spectral measure. The idea is to replace the vector $\mathbf{X}/|\mathbf{X}|$ (resp. Θ) by $\pi(\mathbf{X}/t)$ (resp. $\mathbf{Z} = \pi(Y\Theta)$) where Y and t are the same as in (1.1) and where π denotes the Euclidean projection onto \mathbb{S}_+^{d-1} (Duchi et al. (2008), Kyrillidis et al. (2013), Condat (2016)), see Figure 2a. This substitution leads to the concept of *sparse regular variation*. A random vector \mathbf{X} is said to be sparsely regularly varying if

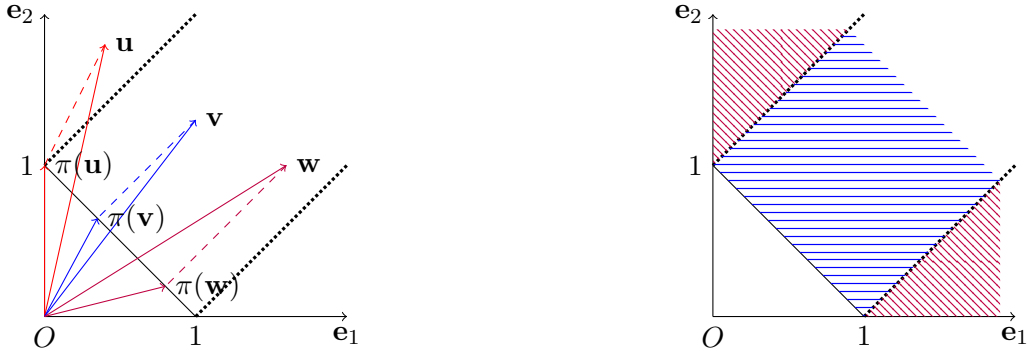
$$\mathbb{P}((|\mathbf{X}|/t, \pi(\mathbf{X}/t)) \in \cdot \mid |\mathbf{X}| > t) \xrightarrow{w} \mathbb{P}((Y, \mathbf{Z}) \in \cdot), \quad t \rightarrow \infty. \quad (1.3)$$

Meyer and Wintenberger (2021) proved that under mild assumption both concepts of regular variation (1.1) and (1.3) are equivalent (see Theorem 1 in their article). Similarly to the existing approaches with Θ , we are willing to capture the tail dependence of \mathbf{X} via the identification of the clusters β which satisfy $\mathbb{P}(\mathbf{Z} \in C_\beta) > 0$. We call such β 's the *extremal clusters*. They can be identified via the study of $\pi(\mathbf{X}/t)$ since the convergence $\mathbb{P}(\pi(\mathbf{X}/t) \in C_\beta \mid |\mathbf{X}| > t) \rightarrow \mathbb{P}(\mathbf{Z} \in C_\beta)$ holds for any $\beta \subset \{1, \dots, d\}$ (see Meyer and Wintenberger (2021), Proposition 2). The latter convergence mainly holds because the subsets $\pi^{-1}(C_\beta)$ do not have empty interior, see Figure 2b.

This encourages to consider the estimators $T_n(\beta) := \sum_{j=1}^n \mathbb{1}\{\pi(\mathbf{X}_j/u_n) \in C_\beta, |\mathbf{X}_j| > u_n\}$ where $\mathbf{X}_1, \dots, \mathbf{X}_n$ is a sample of iid sparsely regularly varying random vectors. These estimators depend on

a threshold u_n satisfying $u_n \rightarrow \infty$. It turns out that these estimators often overestimate the total number of extremal clusters β . We call the clusters which satisfy $T_n(\beta) > 0$ and $\mathbb{P}(\mathbf{Z} \in C_\beta) = 0$ the *biased clusters*. The approach we propose to reduce this bias relies on model selection and is similar to the minimization criterion developed by Akaike (1973). It is achieved for a fixed threshold for which Equation (1.3) becomes a good approximation. This is the first step of our procedure, which we call the *bias selection*.

The second step then consists in extending the procedure in order to automatically select an appropriate threshold. We call this step the *threshold selection*. Several authors have pointed out that choosing an optimal threshold is a challenging task in practice (see for instance Rootzén et al. (2006)). This issue is however tackled in a few articles (Stărică (1999), Abdous and Ghoudi (2005), Lee et al. (2015), Kiriliouk et al. (2019), Wan and Davis (2019), see also the review on marginals threshold selection by Caeiro and Gomes (2015)). It turns out that for sparse regular variation the choice of such a threshold is closely related to the clusters β on which $\pi(\mathbf{X}/t)$ places mass. Therefore our approach consists in extending the bias selection by including u_n as a parameter to tune. To the best of our knowledge, our work is the first one which simultaneously tackles this issue with the study of tail dependence.



(a) Three vectors and their image by π . The dotted lines partition the space depending on the localization of the projected vectors: \mathbf{e}_1 , \mathbf{e}_2 , or the interior of the simplex.

(b) The preimages of the subsets C_β by π . In purple $\pi^{-1}(C_{\{1\}})$ and $\pi^{-1}(C_{\{2\}})$, and in blue $\pi^{-1}(C_{\{1,2\}})$.

Figure 2: Euclidean projection onto the simplex \mathbb{S}_+^1 .

Outline of the paper The paper is organized as follows. Section 2 introduces the theoretical background on sparse regular variation and threshold selection that is needed throughout the paper. We particularly insist on the influence of the threshold t on the sparsity of the projected vector $\pi(\mathbf{X}/t)$. We end the section by a presentation of our procedure to detect the extremal clusters of a random vector \mathbf{X} . In particular we introduce the penalized criterion on which the model selection relies. Section 3 details the methodology of our approach. We develop the two steps of the model selection, the bias selection and the threshold selection. We gather all asymptotic results in Section 4 and we prove asymptotic normality for the estimators $T_n(\beta)$. We also derive in this section a test procedure based on the score to identify biased clusters. In Section 5 we illustrate our findings on different numerical results and compare our approach with the existing procedures of Goix et al. (2017) and Simpson et al. (2019). Finally we illustrate our approach on financial and environmental data in Section 6.

2 Preliminaries and MUSCLE algorithm

2.1 Notation

We introduce some standard notation that is used throughout the paper. Symbols in bold such as $\mathbf{x} \in \mathbb{R}^d$ are column vectors with components denoted by $x_j, j \in \{1, \dots, d\}$. Operations and relationships involving such vectors are meant componentwise. If $\mathbf{x} = (x_1, \dots, x_d)^\top \in \mathbb{R}^d$ is a vector, then $\text{Diag}(x_1, \dots, x_d)$ denotes the diagonal matrix whose diagonal is \mathbf{x} . We also denote by Id_s the identity matrix of \mathbb{R}^s . We define $\mathbb{R}_+^d := \{\mathbf{x} \in \mathbb{R}^d : x_1 \geq 0, \dots, x_d \geq 0\}$, $\mathbf{0} := (0, \dots, 0)^\top \in \mathbb{R}^d$, and $\mathbf{1} := (1, \dots, 1)^\top \in \mathbb{R}^d$. For $j = 1, \dots, d$, \mathbf{e}_j denotes the j -th vector of the canonical basis of \mathbb{R}^d . If $\mathbf{x} \in \mathbb{R}^d$ and $I := \{i_1, \dots, i_r\} \subset \{1, \dots, d\}$, then \mathbf{x}_I denotes the vector $(x_{i_1}, \dots, x_{i_r})$ of \mathbb{R}^r . For $p \in (1, \infty]$ we denote by $|\cdot|_p$ the ℓ^p -norm in \mathbb{R}^d . The ℓ^1 -norm will simply be denoted by $|\cdot|$. For $d \geq 1$ we denote by \mathcal{P}_d the power set of $\{1, \dots, d\}$ and by \mathcal{P}_d^* the set $\mathcal{P}_d \setminus \{\emptyset\}$. Finally we write $A \stackrel{d}{=} B$ when the random variables A and B have the same distribution.

2.2 Sparse regular variation

We consider a sparsely regularly varying random vector $\mathbf{X} \in \mathbb{R}_+^d$ and focus on its angular component $\pi(\mathbf{X}/t)$:

$$\mathbb{P}(\pi(\mathbf{X}/t) \in \cdot \mid |\mathbf{X}| > t) \xrightarrow{w} \mathbb{P}(\mathbf{Z} \in \cdot), \quad t \rightarrow \infty, \quad (2.1)$$

where $\pi : \mathbb{R}_+^d \rightarrow \mathbb{S}_+^{d-1}$ denotes the Euclidean projection onto the simplex. We refer to [Duchi et al. \(2008\)](#), [Kyrillidis et al. \(2013\)](#), and [Condat \(2016\)](#) for a theoretical and algorithmic study of this function. Both limit components of (2.1) satisfies the relation $\mathbf{Z} \mid Y > r \stackrel{d}{=} \pi(r\mathbf{Z})$ for all $r > 1$. The choice of this particular angular vector is justified by the many properties enjoyed by the projection π . The main one is that the vector $\pi(\mathbf{X}/t)$ may put mass in every subspace C_β even if \mathbf{X} does not. This is a key difference with the self-normalized vector $\mathbf{X}/|\mathbf{X}|$ which shares the same sparsity properties as \mathbf{X} , and therefore always concentrates on the interior $C_{\{1, \dots, d\}}$ of the simplex.

Our aim is to infer the behavior of the angular vector \mathbf{Z} by focusing on the probabilities $p^*(\beta) := \mathbb{P}(\mathbf{Z} \in C_\beta)$ for $\beta \in \mathcal{P}_d^*$. We are willing to identify the extremal directions of \mathbf{X} , i.e. the clusters of coordinates β for which are $p^*(\beta) > 0$. We define the set

$$\mathcal{S}^*(\mathbf{Z}) := \{\beta : p^*(\beta) > 0\}, \quad (2.2)$$

and denote by s^* its cardinality. [Meyer and Wintenberger \(2021\)](#) proved that for any β we have

$$\mathbb{P}(\pi(\mathbf{X}/t) \in C_\beta \mid |\mathbf{X}| > t) \rightarrow p^*(\beta), \quad t \rightarrow \infty. \quad (2.3)$$

This convergence allows one to study the behavior of \mathbf{Z} on the subsets C_β via the one of $\pi(\mathbf{X}/t)$. This main difference between the vector \mathbf{Z} and Θ comes from the topological properties of the subspaces $\{|\mathbf{x}| > 1, \pi(\mathbf{x}) \in C_\beta\}$. Indeed, while the latter are of positive Lebesgue measure the subsets $\{|\mathbf{x}| > 1, \mathbf{x}/|\mathbf{x}| \in C_\beta\}$ associated to $\mathbf{X}/|\mathbf{X}|$ are of zero Lebesgue measure (see [Figure 2b](#) for an illustration of the subspaces $\pi^{-1}(C_\beta)$ in dimension 2). This topological difference entails that it is easier to estimate the tail dependence of \mathbf{X} with \mathbf{Z} rather than with Θ .

The aim of this paper is to build a statistical procedure to identify the extremal clusters $\beta \in \mathcal{S}^*(\mathbf{Z})$. In a statistical setting we consider Equation (2.3) as an approximation for t high enough. Hence we need to provide a method to select an optimal threshold t . This issue is all the more challenging since the directions on which $\pi(\mathbf{X}/t)$ places mass vary with t , as explained in the following section.

Example 1 (Discrete spectral measure). For $\beta \in \mathcal{P}_d^*$, we denote by $\mathbf{e}(\beta)$ the sum $\sum_{j \in \beta} \mathbf{e}_j$. Then the vector $\mathbf{e}(\beta)/|\beta|$ belongs to the simplex \mathbb{S}_+^{d-1} , with $|\beta|$ corresponding the length of the cluster β . We

consider the following family of discrete distributions on the simplex:

$$S = \sum_{\beta \in \mathcal{P}_d^*} c(\beta) \delta_{\mathbf{e}(\beta)/|\beta|}, \quad (2.4)$$

where $(c(\beta))_\beta$ is a probability vector on \mathbb{R}^{2^d-1} (see Segers (2012), Example 3.3). Meyer and Wintenberger (2021) proved that in this case we have $\mathbf{Z} = \Theta$ a.s. and that the family of distribution in (2.4) is the only possible discrete distributions for \mathbf{Z} . We have $\mathcal{S}^*(\mathbf{Z}) = \{\beta : c(\beta) > 0\}$.

If we choose $c(\beta) = 1/d$ for all β of length 1 then the spectral measure becomes $S = d^{-1} \sum_{j=1}^d \delta_{\mathbf{e}_j}$, which corresponds to asymptotic independence (see for instance Ledford and Tawn (1996), Heffernan and Tawn (2004), de Haan and Ferreira (2006), Section 6.2).

2.3 Impact of the threshold on the bias selection

In a statistical context we want to study the tail behavior of n iid random vectors. It is then common to choose a threshold $u_n \rightarrow \infty$, or a number $k = k_n$ of vectors which have the largest norms satisfying $k \rightarrow \infty$ and $k/n \rightarrow 0$ as $n \rightarrow \infty$ (de Haan and Ferreira (2006), Beirlant et al. (2006), Resnick (2007)). The most common choice for k in a multivariate setting is

$$k = n\mathbb{P}(|\mathbf{X}| > u_n). \quad (2.5)$$

The smaller k the closer we are from the theoretical framework. However, we need to keep a substantial number of extreme data to correctly learn their tail structure. No rigorous methodology for choosing k has been obtained in a multivariate framework yet.

For $t > 0$, let us denote by π_t the Euclidean projection onto the positive sphere $\{\mathbf{x} \in \mathbb{R}_+^d : x_1 + \dots + x_d = t\}$. For a vector $\mathbf{v} \in \mathbb{R}_+^d$ with ℓ^1 -norm $|\mathbf{v}|$ the number of null coordinates of the projected vector $\pi_t(\mathbf{v})$ strongly depends on the choice of t . Indeed, if t is close to $|\mathbf{v}|$ then $\pi_t(\mathbf{v})$ has almost only non-null coordinates (as soon as \mathbf{v} itself has non-null coordinates). On the contrary, if $t \ll |\mathbf{v}|$ then the vector $\pi_t(\mathbf{v})$ becomes sparser. The impact of the threshold t on the sparsity of projected vectors is illustrated in Figure 3.

For a large threshold t only extreme data are selected but many vectors are close to this threshold. This implies that these vectors are projected on subsets C_β with large $|\beta|$'s. The projected vectors are thus not very sparse. On the other hand if we select a low threshold then we move away from the extreme region. In this case the largest vectors are projected on subsets C_β with small $|\beta|$'s, i.e. the projected vectors are very sparse. Thus we have to make a balanced choice between providing a sparse structure for the data and staying in the extreme region.

These considerations entail that the bias selection and the threshold selection have to be achieved accordingly. For practical reasons it is often more convenient to focus on the number of exceedances k rather than on the threshold t . We observe on Figure 3 that the projected vector $\pi(\mathbf{X}/t)$ is all the more sparse that the threshold t is away from the vector \mathbf{X} . If the level k is fixed, the sparsest representation of $\pi(\mathbf{X}_j/t)$ is thus obtained for $t = |\mathbf{X}_{(k+1)}|$, where $|\mathbf{X}_{(j)}|$ denotes the j th largest norm among the sample $\mathbf{X}_1, \dots, \mathbf{X}_n$.

The identification of an appropriate number of extreme values k and of the set $\mathcal{S}^*(\mathbf{Z})$ is then achieved via a model selection procedure. For a textbook treatment of model selection, see Massart (2007). Our approach consists in fitting a multinomial model to the data. This is achieved by comparing the Kullback-Leibler divergence between the data and the theoretical model (Kullback and Leibler (1951)), see Section 2.4. This leads to a model selection procedure which relies on the minimization of a penalized maximum likelihood similarly to Akaike's criterion (Akaike (1973)). Since the latter only holds for a constant sample size, we have to adapt the standard approach to an extreme setting where the number of extremes vary. Therefore we include the non-extreme values in the model and separate the data into an *extreme* group and a *non-extreme* one. The procedure then provides a level k for which this separation is optimal.

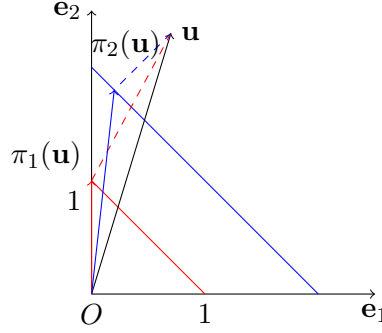


Figure 3: Impact of the threshold on sparsity. The image of the vector \mathbf{u} is $\pi_1(\mathbf{u}) = (0, 1)$ with the threshold $z = 1$ while it is $\pi_2(\mathbf{u}) > \mathbf{0}$ with the threshold $z = 2$. The sparsity decreases when the threshold increases.

2.4 Algorithm: MULTivariate Sparse CLustering for Extremes (MUSCLE)

Our procedure aims at providing the tail structure of \mathbf{X} based on a sample $\mathbf{X}_1, \dots, \mathbf{X}_n$ of iid sparsely regularly varying random vectors with the same distribution as \mathbf{X} . In order to identify the set $\mathcal{S}^*(\mathbf{Z})$ defined in (2.2) we provide suitable estimators for the family of probabilities $\mathbf{p}^* := (p^*(\beta))_{\beta \in \mathcal{P}_d^*}$. Therefore we rewrite Equation (2.3) as

$$p_n(\beta) := \mathbb{P}(\pi(\mathbf{X}/u_n) \in C_\beta \mid |\mathbf{X}| > u_n) \rightarrow p^*(\beta), \quad n \rightarrow \infty,$$

where u_n is a threshold satisfying $u_n \rightarrow \infty$. We define the estimators

$$T_n(\beta) = \sum_{j=1}^n \mathbb{1}\{\pi(\mathbf{X}_j/u_n) \in C_\beta, |\mathbf{X}_j| > u_n\}, \quad \beta \in \mathcal{P}_d^*,$$

and consider the random vector $\mathbf{T}_n = (T_n(\beta))_{\beta \in \mathcal{P}_d^*}$ whose components are in the decreasing order (see Section 3). This leads to the empirical version of the set $\mathcal{S}^*(\mathbf{Z})$ given by

$$\hat{\mathcal{S}}_n := \{\beta : T_n(\beta) > 0\}. \quad (2.6)$$

We denote by \hat{s}_n its cardinality. For a given integer level k satisfying (2.5) the vector \mathbf{T}_n follows a multinomial distribution denoted by \mathbf{P}_k with parameters k and $\mathbf{p}_n := (p_n(\beta))_{\beta \in \mathcal{P}_d^*}$. The bias between the clusters β in $\mathcal{S}^*(\mathbf{Z})$ and in $\hat{\mathcal{S}}_n$ entails that $s^* \leq \hat{s}_n$ (see also Section 4.2). Note that for sake of simplicity we omit the dependence in u_n of the quantities \mathbf{T}_n , \mathbf{p}_n , and \hat{s}_n .

The bias selection consists in comparing the distribution \mathbf{P}_k with a theoretical multinomial model \mathbf{M}_k whose probability vector is given by $(p_1, \dots, p_s, p, \dots, p, 0, \dots, 0)^\top \in \mathbb{R}^{2^d-1}$, with $p_1 \geq \dots \geq p_s > p$ and $r - s$ components p satisfying $p_1 + \dots + p_1 + (r - s)p = 1$. The parameters p_j model the probability that \mathbf{Z} belongs to the associated subsets C_β while the parameter p models the probability that a biased cluster appears. We denote by \mathbf{p} the vector $(p_1, \dots, p_s)^\top \in \mathbb{R}^s$. The identification of the extremal clusters β in $\mathcal{S}^*(\mathbf{Z})$ is then done by choosing the model \mathbf{M}_k which best fits the sample \mathbf{T}_n . To this end, recall that the Kullback-Leibler divergence between two discrete distributions F and G is given by

$$KL(F||G) = \mathbb{E}_{\xi \sim F} \left[\log \left(\frac{F(\xi)}{G(\xi)} \right) \right] = \sum_{x \in \Xi} x \log(F(x)/G(x)),$$

where Ξ denotes the support of the distribution F (Kullback and Leibler (1951)). In order to minimize the divergence $KL(\mathbf{P}_k||\mathbf{M}_k)$ we estimate it by $KL(\mathbf{P}_k||\mathbf{M}_k)|_{\hat{\mathbf{p}}}$, where $\hat{\mathbf{p}}$ denotes the maximum-likelihood

estimator of \mathbf{p} (see [Akaike \(1973\)](#)). This boils down to finding s which minimizes the penalized log-likelihood

$$-\log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n) + (s + 1), \quad (2.7)$$

where $L_{\mathbf{M}_k}$ denotes the likelihood of the model \mathbf{M}_k , see [Section 3.1](#) for more details. We denote by $\hat{s}(k)$ the parameter s for which the minimum is reached.

The second step consists in including the level k as a parameter of the model in order to tackle the threshold selection. This is done by considering the vector \mathbf{T}'_n which models the behavior of the large vectors \mathbf{X}_j 's on the subsets C_β as well as the non-extreme ones. We assume that this vector follows a multinomial distribution \mathbf{P}'_n . We consider an extended multinomial model \mathbf{M}'_n with probability vector given by $(q'p'_1, \dots, q'p'_{s'}, q'p', \dots, q'p', 0, \dots, 0, 1 - q')^\top \in \mathbb{R}^{2^d}$ with $p'_1 \geq \dots \geq p'_{s'} > p'$ and $r' - s'$ components $q'p'$ satisfying the relation $p'_1 + \dots + p'_{s'} + (r' - s')p' = 1$. Here q' models the proportion of extreme vectors. We denote by \mathbf{p}' the vector $(p'_1, \dots, p'_{s'}, q')^\top \in \mathbb{R}^{s'+1}$. As in the bias selection we estimate the Kullback-Leibler divergence $KL(\mathbf{P}'_n || \mathbf{M}'_n)$ by $KL(\mathbf{P}'_n || \mathbf{M}'_n)|_{\hat{\mathbf{p}}'}$, where $\hat{\mathbf{p}}'$ denotes the maximum-likelihood estimator of \mathbf{p}' . Thus we choose the model \mathbf{M}'_n whose parameters minimize the following penalized log-likelihood:

$$\frac{1}{k} \left(-\log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n) + (s + 1) - k \log(1 - k/n) \right), \quad (2.8)$$

see [Section 3](#). This leads to the choice of the estimators \hat{k} and $\hat{s}(\hat{k})$. Then we define $\hat{\mathcal{S}}^*$ as the set gathering the $\hat{s}(\hat{k})$ clusters corresponding to the largest $T_n(\beta)$'s. Finally we consider the probability vector $\hat{\zeta}$ defined by

$$\hat{\zeta}(\beta) := \frac{T_n(\beta)/\hat{k}}{\sum_{\gamma \in \hat{\mathcal{S}}^*} T_n(\gamma)/\hat{k}},$$

for $\beta \in \hat{\mathcal{S}}^*$ and 0 elsewhere, as an estimator of \mathbf{p}^* . Our procedure entails the following parameter-free algorithm called MUSCLE for MULTivariate Sparse CLustering for Extremes.

Algorithm 1: MULTivariate Sparse CLustering for Extremes (MUSCLE)

Data: A sample $\mathbf{X}_1, \dots, \mathbf{X}_n \in \mathbb{R}_+^d$ and a range of values K for the level

Result: A list $\hat{\mathcal{S}}^*$ of clusters β and the associated probability vector $\hat{\zeta}$.

for $k \in K$ **do**

 Compute $u_n = |\mathbf{X}|_{(k+1)}$ the $(k + 1)$ -th largest norm;

 Assign to each $\pi(\mathbf{X}_j/u_n)$ the subsets C_β it belongs to;

 Compute \mathbf{T}_n ;

 Compute the minimizer $\hat{s}(k)$ which minimizes the criterion given in [Equation \(2.7\)](#);

end

Minimize \hat{k} of [\(2.8\)](#) plugging in the minimal value in [\(2.7\)](#);

Output: $\hat{\mathcal{S}}^* = \{\text{the clusters } \beta \text{ associated to the } T_{n,1}, \dots, T_{n,\hat{s}(\hat{k})}\}$ and $\hat{\zeta}$ as above.

3 Methodology

We develop in this section our methodology to estimate the set $\mathcal{S}^*(\mathbf{Z})$. We consider a sample of sparsely regularly varying random vectors $\mathbf{X}_1, \dots, \mathbf{X}_n$ with generic distribution \mathbf{X} . Our model selection relies on some asymptotic results for which it is necessary to have $k_n p_n(\beta) \rightarrow \infty$ when $n \rightarrow \infty$ for some clusters β and a level $k = k_n$. Thus for a given level k we define the set

$$\mathcal{S}_k^\infty = \{\beta \in \mathcal{P}_d^* : k_n p_n(\beta) \rightarrow \infty \text{ when } n \rightarrow \infty\}, \quad (3.1)$$

and denote by s_∞ its cardinality. In what follows we fix a large n and work under the event $\mathcal{S}_k^\infty = \widehat{\mathcal{S}}_n$. We refer to Section 4.2 for more insights on this assumption.

3.1 Bias selection

In this section we fix a level $k = k_n$. Recall that the distribution of \mathbf{T}_n is denoted by \mathbf{P}_k and that the multinomial model \mathbf{M}_k has $2^d - 1$ outcomes adding up to k and probability vector $(p_1, \dots, p_s, p, \dots, p, 0, \dots, 0)^\top$. Recall also that \mathbf{p} denotes the vector $(p_1, \dots, p_s)^\top \in \mathbb{R}^s$. In order to estimate the Kullback-Leibler divergence between \mathbf{P}_k and \mathbf{M}_k we compute the maximum likelihood estimators of the latter model. The likelihood $L_{\mathbf{M}_k}$ is given by

$$L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{x}) = \frac{k!}{\prod_{i=1}^{2^d-1} x_i!} \prod_{i=1}^s p_i^{x_i} \prod_{i=s+1}^r \left(\frac{1 - \sum_{j=1}^s p_j}{r - s} \right)^{x_i} \mathbf{1}_{\{x_j = 0, j \geq r+1\}},$$

for any vector $\mathbf{p} \in [0, 1]^s$ such that $p_1 + \dots + p_s \leq 1$, and $\mathbf{x} \in \mathbb{R}_+^{2^d-1}$ adding up to k . It is maximal for a vector \mathbf{x} whose components are ordered in the decreasing order. Hence we define $T_{n,1} = \max_\beta T_n(\beta)$ and

$$T_{n,j} = \max \{T_n(\beta), \beta \in \mathcal{P}_d^* \setminus \{T_{n,1}, \dots, T_{n,j-1}\}\}, \quad j = 2, \dots, 2^d - 1.$$

We consider the associated order for the vectors \mathbf{p}^* and \mathbf{p}_n . Finally, we fix the number of positive parameters r to be equal to the number of clusters that appear \widehat{s}_n and work under this event.

The optimization of the log-likelihood

$$\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n) = \log(k!) - \sum_{i=1}^{2^d-1} \log(T_{n,i}!) + \sum_{i=1}^s T_{n,i} \log(p_i) + \left(\sum_{i=s+1}^r T_{n,i} \right) \log \left(\frac{1 - \sum_{j=1}^s p_j}{r - s} \right) \quad (3.2)$$

yields the maximum likelihood estimators $\widehat{p}_j := T_{n,j}/k$ for $1 \leq j \leq s$.

The last parameter remaining is s which needs to be fitted. It is achieved by computing the Kullback-Leibler divergence between the true distribution \mathbf{P}_k and the model \mathbf{M}_k ,

$$KL(\mathbf{P}_k \parallel \mathbf{M}_k) = \mathbb{E} \left[\log \left(\frac{L_{\mathbf{P}_k}(\mathbf{T}_n)}{L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n)} \right) \right] = \mathbb{E}[\log L_{\mathbf{P}_k}(\mathbf{T}_n)] - \mathbb{E}[\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n)]. \quad (3.3)$$

This quantity must be seen as a function of \mathbf{p} . In particular the first term of the right-hand side is constant with respect to the parameter \mathbf{p} . Regarding the second term, Equation (3.2) yields

$$\mathbb{E}[\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n)] = \log(k!) - \mathbb{E} \left[\sum_{i=1}^{2^d-1} \log(T_{n,i}!) \right] + k \sum_{i=1}^s p_{n,i} \log(p_i) + k \left(\sum_{i=s+1}^r p_{n,i} \right) \log \left(\frac{1 - \sum_{j=1}^s p_j}{r - s} \right), \quad (3.4)$$

for any vector $\mathbf{p} \in [0, 1]^s$ such that $p_1 + \dots + p_s \leq 1$. Following the approach of Akaike (1973) we select the multinomial model which minimizes the Kullback-Leibler divergence between the true distribution and the model \mathbf{M}_k evaluated at the maximum likelihood $\widehat{\mathbf{p}}$:

$$KL(\mathbf{P}_k \parallel \mathbf{M}_k)|_{\mathbf{p}=\widehat{\mathbf{p}}} = \mathbb{E}[\log L_{\mathbf{P}_k}(\mathbf{T}_n)] - \mathbb{E}[\log L_{\mathbf{M}_k}(\widehat{\mathbf{p}}; \mathbf{T}_n)]|_{\mathbf{p}=\widehat{\mathbf{p}}}. \quad (3.5)$$

We establish in Section 4 some asymptotic results which lead to the following approximation for the expectation of the estimator in Equation (3.5):

$$\begin{aligned} \mathbb{E}[KL(\mathbf{P}_k \parallel \mathbf{M}_k)|_{\mathbf{p}=\widehat{\mathbf{p}}}] &\approx \mathbb{E}[\log L_{\mathbf{P}_k}(\mathbf{T}_n)] - \mathbb{E}[\log L_{\mathbf{M}_k}(\widehat{\mathbf{p}}; \mathbf{T}_n)] + \mathbb{E}[\chi^2(s+1)] \\ &\approx \mathbb{E}[\log L_{\mathbf{P}_k}(\mathbf{T}_n)] - \mathbb{E}[\log L_{\mathbf{M}_k}(\widehat{\mathbf{p}}; \mathbf{T}_n)] + (s+1), \end{aligned}$$

where $\chi^2(s+1)$ denotes a chi-squared distribution with $s+1$ degrees of freedom. We defer the technical calculations to Appendix C.1. The first term of the right-hand side is constant with respect to the parameter \mathbf{p} . Moving back to Equation (3.3), we estimate the quantity $\mathbb{E}[\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n)]$ with $-\log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n) + (s+1)$. Therefore for a given level k the bias selection procedure consists in choosing the parameter $\hat{s}(k)$ which minimizes this penalized log-likelihood. This provides $\hat{s}(k)$ extremal clusters β on which the distribution of the vector \mathbf{Z} places mass.

3.2 Threshold selection

The second step is to consider $k = k_n$ as a parameter which has to be estimated and tuned. It is therefore necessary to consider all observations $\mathbf{X}_1, \dots, \mathbf{X}_n$ and not only the extreme ones. We consider a vector \mathbf{T}'_n which models the repartition of the \mathbf{X}_j 's on the subsets C_β and with the last components $T'_{n,2^d}$ corresponding to the number of non-extreme values of the sample. We assume that this vector follows a multinomial distribution \mathbf{P}'_n with parameter n and probability vector

$$\mathbf{p}'_n = (q_n p_{n,1}, \dots, q_n p_{n,2^d-1}, 1 - q_n)^\top \in \mathbb{R}^{2^d}, \quad (3.6)$$

where the parameter $1 - q_n$ is associated to the non-extreme values. We make the following assumption on the quantity q_n .

Assumption 1. $nq_n \log(n) \rightarrow 0$ when $n \rightarrow \infty$.

Similarly to Section 3.1 we aim to compare the distribution \mathbf{P}'_n with the theoretical multinomial model \mathbf{M}'_n with parameter \mathbf{p}' defined in Section 2.4 by computing the Kullback-Leibler divergence between both distributions. The interpretation of the model \mathbf{M}'_n is the same as for \mathbf{M}_k . The extra-parameter q' models the proportion of extreme values taken among the data.

We are willing to select the model \mathbf{M}'_n which minimizes the Kullback-Leibler divergence between \mathbf{P}'_n and \mathbf{M}'_n given by

$$KL(\mathbf{P}'_n \parallel \mathbf{M}'_n) = \mathbb{E} \left[\log \left(\frac{L_{\mathbf{P}'_n}(\mathbf{T}'_n)}{L_{\mathbf{M}'_n}(\mathbf{p}'; \mathbf{T}'_n)} \right) \right] = \mathbb{E}[\log L_{\mathbf{P}'_n}(\mathbf{T}'_n)] - \mathbb{E}[\log L_{\mathbf{M}'_n}(\mathbf{p}'; \mathbf{T}'_n)], \quad (3.7)$$

where $L_{\mathbf{P}'_n}$ (resp. $\log L_{\mathbf{M}'_n}$) denotes the likelihood of the distribution \mathbf{P}'_n (resp. \mathbf{M}'_n). The log-likelihood $\log L_{\mathbf{M}'_n}(\mathbf{p}'; \mathbf{T}'_n)$ can be decomposed as

$$\begin{aligned} \log L_{\mathbf{M}'_n}(\mathbf{p}'; \mathbf{T}'_n) &= \log(n!) - \sum_{j=1}^{2^d} \log(T'_{n,j}!) + \sum_{j=1}^{s'} T'_{n,j} \log(q' p'_j) + \left(\sum_{j=s'+1}^{2^d-1} T'_{n,j} \right) \log(p' q') + T'_{n,2^d} \log(1 - q') \\ &= \log((n - T'_{n,2^d})!) - \sum_{j=1}^{2^d-1} \log(T'_{n,j}!) + \sum_{j=1}^{s'} T'_{n,j} \log(p'_j) + \log(p') \sum_{j=s'+1}^{2^d-1} T'_{n,j} \\ &\quad + \log \left(\frac{n!}{(n - T'_{n,2^d})!} \right) - \log(T'_{n,2^d}!) + (n - T'_{n,2^d}) \log(q') + T'_{n,2^d} \log(1 - q') \\ &= \log L_{\mathbf{M}_{n-T'_{n,2^d}}}(\mathbf{p}; \mathbf{T}'_{n,\{1, \dots, 2^d-1\}}) + \phi(n, q', T'_{n,2^d}), \end{aligned}$$

where

$$\phi(n, q', T'_{n,2^d}) = \log \left(\frac{n!}{(n - T'_{n,2^d})!} \right) - \log(T'_{n,2^d}!) + (n - T'_{n,2^d}) \log(q') + T'_{n,2^d} \log(1 - q').$$

Following the same ideas as in Section 3.1 and similarly to an AIC procedure we estimate the Kullback-Leibler divergence in Equation (3.7) by the estimator $KL(\mathbf{P}'_n || \mathbf{M}'_n) |_{\hat{\mathbf{p}}'}$ where $\hat{\mathbf{p}}'$ denotes the maximum likelihood estimator of \mathbf{p}' . Hence, we are willing to study the expectation of this estimator:

$$\mathbb{E}[KL(\mathbf{P}'_n || \mathbf{M}'_n) |_{\hat{\mathbf{p}}'}] = \mathbb{E}[\log L_{\mathbf{P}_n}(\mathbf{T}'_n)] + \mathbb{E} \left[\mathbb{E} \left[-\log L_{\mathbf{M}_{n-T'_{n,2^d}}}(\mathbf{p}; \mathbf{T}'_{n, \{1, \dots, 2^d-1\}}) | T'_{n, 2^d} \right] \Big|_{\hat{\mathbf{p}}} \right] - \mathbb{E} \left[\mathbb{E} \left[\phi(n, q', T'_{n, 2^d}) \right] \Big|_{\hat{\mathbf{p}}'} \right]. \quad (3.8)$$

We refer to Appendix C.2 for several calculations regarding the aforementioned quantity which lead to the following approximation:

$$\mathbb{E}[KL(\mathbf{P}'_n || \mathbf{M}'_n) |_{\hat{\mathbf{p}}'}] \approx C_n + D_n \frac{1}{k} \left(\mathbb{E}[-\log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n)] + (s+1) - k \log(1 - k/n) \right),$$

where C_n and $D_n > 0$ are constants depending on n . This implies that the penalized log-likelihood

$$\frac{1}{k} \left(-\log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n) + (s+1) - k \log(1 - k/n) \right) \quad (3.9)$$

provides up to some constants an approximation of the Kullback-Leibler divergence in Equation (3.7) evaluated at the maximum likelihood estimator $\hat{\mathbf{p}}'$. In practice we choose a large range of k (often between 0.5% and 10% of n) and we compute the value of (3.9) for these k and for $s = 1, \dots, \hat{s}_n$, where we recall that \hat{s}_n depends on the chosen level k . Finally we choose the couple $(\hat{k}, \hat{s}(\hat{k}))$ which minimizes this quantity in (3.9). We then consider the vector $\hat{\boldsymbol{\zeta}}$ defined in Section 2.4 as an estimator of vector \mathbf{p}^* .

Remark 1. While our procedure leads to the choice of a unique \hat{k} , we expect that this approach is not too sensitive to this choice. Therefore, it is relevant to plot the function $k \mapsto \hat{s}(k)$ on which a stability region around the optimal value \hat{k} should appear.

4 Asymptotic analysis of the support of extremes

In this section we establish asymptotic results which justify the approximations in the model selection procedure developed in Section 3. Note that in all what follows we study random variables with the condition $|\mathbf{X}| > u_n$.

4.1 Consistency

We consider a sequence of iid sparsely regularly varying random vectors $\mathbf{X}_1, \dots, \mathbf{X}_n$ with generic distribution \mathbf{X} , a positive sequence $(u_n)_{n \in \mathbb{N}}$ and a level $k = k_n = n\mathbb{P}(|\mathbf{X}| > u_n)$. We discuss some convergence results for the random vector \mathbf{T}_n . The first one is consistency.

Proposition 1. *If $u_n \rightarrow \infty$ and $k_n = n\mathbb{P}(|\mathbf{X}| > u_n) \rightarrow \infty$, then the following convergence in probability holds:*

$$\frac{\mathbf{T}_n}{k_n} \rightarrow \mathbf{p}^*, \quad n \rightarrow \infty. \quad (4.1)$$

Proof. It suffices to prove that a univariate convergence holds for every cluster β . For $\epsilon > 0$, Markov's

inequality entails that

$$\begin{aligned}
\mathbb{P}\left(\left|\frac{T_n(\beta)}{k_n} - p_n(\beta)\right| > \epsilon\right) &\leq \frac{\mathbb{E}[|T_n(\beta)/k_n - p_n(\beta)|^2]}{\epsilon^2} \\
&\leq \frac{\mathbb{E}\left[\sum_{j=1}^n \mathbb{1}\{\pi(\mathbf{X}_j/u_n) \in C_\beta, |\mathbf{X}_j| > u_n\} - \mathbb{P}(\pi(\mathbf{X}_j/u_n) \in C_\beta, |\mathbf{X}_j| > u_n)\right]^2}{k_n^2 \epsilon^2} \\
&= \frac{n \operatorname{Var}(\mathbb{1}\{\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n\})}{k_n^2 \epsilon^2} \\
&\leq n \mathbb{P}(\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n) k_n^{-2} \epsilon^{-2}.
\end{aligned}$$

Since $n \mathbb{P}(\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n) k_n^{-2} \epsilon^{-2} \leq k_n^{-1} \epsilon^{-2} \rightarrow 0$ when $n \rightarrow \infty$ we obtain that $T_n(\beta)/k_n - p_n(\beta)$ converges to 0 in probability. The convergence $p_n(\beta) \rightarrow p^*(\beta)$ in Equation (2.3) concludes the proof. \square

Proposition 1 implies that if $p^*(\beta) = 0$, i.e. if \mathbf{Z} does not place mass on the subset C_β , then $T_n(\beta)/k_n$ becomes smaller and smaller as n increases. Actually as soon as the dimension d is large a lot of $T_n(\beta)$'s are even equal to 0 since the level k_n is far below the number of subsets $2^d - 1$.

4.2 The bias between $\widehat{\mathcal{S}}_n$ and $\mathcal{S}^*(\mathbf{Z})$

We study in this section the difference between $\mathcal{S}^*(\mathbf{Z})$ and the empirical set $\widehat{\mathcal{S}}_n$. The following lemma gives an asymptotic relation for the probability $\mathbb{P}(T_n(\beta) = 0)$.

Lemma 1. *For $\beta \in \mathcal{P}_d^*$, we have*

$$\frac{\log \mathbb{P}(T_n(\beta) = 0)}{-k_n p_n(\beta)} \rightarrow 1, \quad n \rightarrow \infty.$$

Proof. We have

$$\begin{aligned}
\mathbb{P}(T_n(\beta) = 0) &= \mathbb{P}(\pi(\mathbf{X}_j/u_n) \notin C_\beta \text{ or } |\mathbf{X}_j| \leq u_n, j = 1, \dots, n) \\
&= [1 - \mathbb{P}(\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n)]^n \\
&= \exp(n \log[1 - \mathbb{P}(\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n)]).
\end{aligned}$$

Since $\mathbb{P}(\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n) \rightarrow 0$ we obtain a Taylor expansion

$$n \log(1 - \mathbb{P}(\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n)) \sim -k_n p_n(\beta), \quad n \rightarrow \infty.$$

This concludes the proof. \square

Lemma 1 encourages to focus on the quantity $k_n p_n(\beta)$. Recall that we defined in (3.1) the set $\mathcal{S}_k^\infty = \{\beta \in \mathcal{P}_d^* : k_n p_n(\beta) \rightarrow \infty \text{ when } n \rightarrow \infty\}$ which contains $\mathcal{S}^*(\mathbf{Z})$. In particular we have the inequality $s^* \leq s_\infty$. A consequence of Lemma 1 is that

$$\mathbb{P}(\mathcal{S}_k^\infty \subset \widehat{\mathcal{S}}_n) = 1 - \mathbb{P}(\exists \beta \in \mathcal{S}_k^\infty, \beta \notin \widehat{\mathcal{S}}_n) \geq 1 - \sum_{\beta \in \mathcal{S}_k^\infty} \mathbb{P}(T_n(\beta) = 0) \rightarrow 1,$$

when $n \rightarrow \infty$. Consequently, since $\mathcal{S}^*(\mathbf{Z}) \subset \mathcal{S}_k^\infty$ we obtain the following proposition.

Proposition 2. *With probability converging to 1, we have the inclusions*

$$\mathcal{S}^*(\mathbf{Z}) \subset \mathcal{S}_k^\infty \subset \widehat{\mathcal{S}}_n. \quad (4.2)$$

These inclusions highlight the fact that the observations $T_n(\beta)$ tend to overestimate the number of clusters β in $\mathcal{S}^*(\mathbf{Z})$. It therefore justifies the assumption made for the bias selection that all extremal clusters appear in $\widehat{\mathcal{S}}_n$ but some extra clusters might appear because of a bias between the true probability \mathbf{p}^* and the pre-asymptotic one \mathbf{p}_n . We recall that for our study we work under the event $\mathcal{S}_k^\infty = \widehat{\mathcal{S}}_n$.

Remark 2. If a cluster β satisfies $k_n p_n(\beta) \rightarrow 0$, then Lemma 1 ensures that $\mathbb{P}(T_n(\beta) = 0) \rightarrow 1$. If this holds for all $\beta \in \mathcal{S}^*(\mathbf{Z})^c$, then we obtain that

$$\mathbb{P}(\mathcal{S}^*(\mathbf{Z})^c \subset \widehat{\mathcal{S}}_n^c) = 1 - \mathbb{P}(\exists \beta \in \mathcal{S}^*(\mathbf{Z})^c, \beta \in \widehat{\mathcal{S}}_n) \geq 1 - \sum_{\beta \in \mathcal{S}^*(\mathbf{Z})^c} \mathbb{P}(T_n(\beta) > 0) \rightarrow 1,$$

when $n \rightarrow \infty$. Under such an assumption the sets in (4.2) are all equal with probability converging to 1. However this assumption is quite strong and is not satisfied on numerical examples.

4.3 Asymptotic normality of the estimators

We now study the asymptotic distribution of the vector \mathbf{T}_n . To avoid degenerate cases we restrict ourselves to the clusters $\beta \in \mathcal{S}_k^\infty$ where \mathcal{S}_k^∞ is defined in (3.1). Consequently, the restricted vectors $\mathbf{p}_{\mathcal{S}_k^\infty}^*$, $\mathbf{p}_{n, \mathcal{S}_k^\infty}$, and $\mathbf{T}_{n, \mathcal{S}_k^\infty}$ of \mathbb{R}^{s^∞} are considered. We gather our results in the following theorem whose proof is given in Appendix B.

Theorem 1. *Assume that $u_n \rightarrow \infty$ and $k_n = n\mathbb{P}(|\mathbf{X}| > u_n) \rightarrow \infty$. Then the following convergence hold.*

1. On \mathcal{S}_k^∞ :

$$\sqrt{k_n} \text{Diag}(\mathbf{p}_{n, \mathcal{S}_k^\infty})^{-1/2} \left(\frac{\mathbf{T}_{n, \mathcal{S}_k^\infty}}{k_n} - \mathbf{p}_{n, \mathcal{S}_k^\infty} \right) \xrightarrow{d} \mathcal{N}(0, Id_{s^\infty}), \quad n \rightarrow \infty. \quad (4.3)$$

2. On $\mathcal{S}^*(\mathbf{Z})$:

$$\sqrt{k_n} \text{Diag}(\mathbf{p}_{\mathcal{S}^*(\mathbf{Z})})^{-1/2} \left(\frac{\mathbf{T}_{n, \mathcal{S}^*(\mathbf{Z})}}{k_n} - \mathbf{p}_{n, \mathcal{S}^*(\mathbf{Z})} \right) \xrightarrow{d} \mathcal{N}(0, Id_{s^*}), \quad n \rightarrow \infty. \quad (4.4)$$

3. Moreover, if we assume that

$$\forall \beta \in \mathcal{S}^*(\mathbf{Z}), \quad \sqrt{k_n}(p_n(\beta) - p(\beta)) \rightarrow 0, \quad n \rightarrow \infty, \quad (4.5)$$

then we have

$$\sqrt{k_n} \text{Diag}(\mathbf{p}_{\mathcal{S}^*(\mathbf{Z})})^{-1/2} \left(\frac{\mathbf{T}_{n, \mathcal{S}^*(\mathbf{Z})}}{k_n} - \mathbf{p}_{\mathcal{S}^*(\mathbf{Z})} \right) \xrightarrow{d} \mathcal{N}(0, Id_{s^*}), \quad n \rightarrow \infty. \quad (4.6)$$

Remark 3. Some comments on Theorem 1.

1. In all convergences the square root of the diagonal matrix is meant componentwise.
2. The convergence (4.4) is obtained by restricting the previous convergence to the clusters $\beta \in \mathcal{S}^*(\mathbf{Z})$ and writing down that

$$\text{Diag}(\mathbf{p}_{n, \mathcal{S}^*(\mathbf{Z})})^{1/2} \text{Diag}(\mathbf{p}_{\mathcal{S}^*(\mathbf{Z})})^{-1/2} \rightarrow Id_{s^*}, \quad n \rightarrow \infty.$$

This observation does not hold for $\mathcal{S}_k^\infty \setminus \mathcal{S}^*(\mathbf{Z})$ since for such β 's we have $p^*(\beta) = 0$.

3. The assumption (4.5) is a standard assumption to control the bias.

4. The left-hand side of the convergences (4.3), (4.4), and (4.6) are never degenerate even if the sets $\mathcal{S}^*(\mathbf{Z})$ or \mathcal{S}_k^∞ correspond to all clusters β . Indeed, in such a case the randomness comes from the extremes sampling via the condition $|\mathbf{X}_j| > u_n$.

From Equation (4.3) we obtain that the vector

$$\mathbf{U}_n := \sqrt{k_n} \text{Diag}(\mathbf{p}_{n, \mathcal{S}_k^\infty})^{-1/2} \left(\frac{\mathbf{T}_{n, \mathcal{S}_k^\infty}}{k_n} - \mathbf{p}_{n, \mathcal{S}_k^\infty} \right)$$

satisfies the convergence

$$\mathbf{U}_n^\top \cdot \mathbf{U}_n = k_n \left(\frac{\mathbf{T}_{n, \mathcal{S}_k^\infty}}{k_n} - \mathbf{p}_{n, \mathcal{S}_k^\infty} \right)^\top \text{Diag}(\mathbf{p}_{n, \mathcal{S}_k^\infty})^{-1} \left(\frac{\mathbf{T}_{n, \mathcal{S}_k^\infty}}{k_n} - \mathbf{p}_{n, \mathcal{S}_k^\infty} \right) \xrightarrow{d} \chi^2(s_\infty), \quad (4.7)$$

when $n \rightarrow \infty$ and where $\chi^2(s_\infty)$ denotes a chi-squared distribution with s_∞ degrees of freedom. This convergence can be rephrased as follows:

$$k_n \sum_{j=1}^{s_\infty} \frac{(T_{n,j}/k_n - p_{n,j})^2}{p_{n,j}} \xrightarrow{d} \chi^2(s_\infty), \quad n \rightarrow \infty. \quad (4.8)$$

Now, let us fix $s_1 < s_2$ and consider s_2 clusters $\beta_1, \dots, \beta_{s_2}$ in \mathcal{S}_k^∞ . Since the associated subsets C_β are disjoint, we have the relation

$$\sum_{j=s_1+1}^{s_2} T_n(\beta_j) = T_n(\cup_{j=s_1+1}^{s_2} \beta_j), \quad \text{and} \quad \sum_{j=s_1+1}^{s_2} p_n(\beta_j) = p_n(\cup_{j=s_1+1}^{s_2} \beta_j).$$

This implies that the vector

$$\mathbf{U}_n(s_1) = \sqrt{k_n} \left(\frac{T_n(\beta_1)/k_n - p_n(\beta_1)}{\sqrt{p_n(\beta_1)}}, \dots, \frac{T_n(\beta_{s_1})/k_n - p_n(\beta_{s_1})}{\sqrt{p_n(\beta_{s_1})}}, \frac{\sum_{j=s_1+1}^{s_2} (T_n(\beta_j)/k_n - p_n(\beta_j))}{\sqrt{\sum_{j=s_1+1}^{s_2} p_n(\beta_j)}} \right)^\top$$

converges in distribution to a random vector of \mathbb{R}^{s_1+1} with distribution $\mathcal{N}(0, Id_{s_1+1})$. Then, similarly to Equation (4.7), we have the convergence of $\mathbf{U}_n(s_1)^\top \cdot \mathbf{U}_n(s_1)$ to a chi-square distribution with s_1 degrees of freedom, i.e.

$$k_n \sum_{j=1}^{s_1} \frac{(T_n(\beta_j)/k_n - p_n(\beta_j))^2}{p_n(\beta_j)} + k_n \frac{(\sum_{j=s_1+1}^{s_2} (T_n(\beta_j)/k_n - p_n(\beta_j)))^2}{\sum_{j=s_1+1}^{s_2} p_n(\beta_j)} \xrightarrow{d} \chi^2(s_1 + 1), \quad n \rightarrow \infty. \quad (4.9)$$

This convergence is used to identify the parameter s in the bias selection, see Appendix C.1.

4.4 Asymptotic test for the biased clusters

In this section we advocate our approach by providing an asymptotic test similar to the standard score test. It motivates the use of the AIC approach for the bias selection. The idea is to test whether some parameters p_j of the models are equal or not. This is particularly relevant for the study of the biased clusters for which we assumed in every model \mathbf{M}_k that the associated parameters are the same.

We recall that we work under the event $\mathcal{S}_k^\infty = \widehat{\mathcal{S}}_n$, where \mathcal{S}_k^∞ is defined in (3.1). With such an assumption the convergence (4.3) holds true for the clusters in $\widehat{\mathcal{S}}_n$. We fix $s_1 < s_2$ and consider the model \mathbf{M}_k with s_2 parameters. We test the hypothesis $H_0: \mathbf{T}_n \sim \mathbf{M}_k$ with $p_j = \rho_n$ for all $j = s_1 + 1, \dots, s_2$. Our test statistic is given by

$$\mathcal{T}_n := k_n \sum_{j=s_1+1}^{s_2} \frac{(T_{n,j}/k_n - \hat{\rho}_n)^2}{\hat{\rho}_n},$$

where $\hat{\rho}_n = k_n^{-1}(s_2 - s_1)^{-1} \sum_{i=s_1+1}^{s_2} T_{n,i}$. For a given $\alpha > 0$, we reject H_0 if $\mathcal{T}_n > q_{1-\alpha}^{(s_2-s_1-1)}$, where $q_{1-\alpha}^{(s_2-s_1-1)}$ is the $(1-\alpha)$ -quantile of a chi-square distribution with $s_2 - s_1 - 1$ degrees of freedom. The following proposition, proved in Appendix D, states that such a test has asymptotically a significance level α .

Proposition 3 (Score test for the biased clusters). *With the previous notation we have the convergence*

$$\mathcal{T}_n \xrightarrow{d} \chi^2(s_2 - s_1 - 1), \quad n \rightarrow \infty.$$

Testing the equality of the probability of occurrence of some clusters is more accurate for the biased clusters. Indeed, for the extremal clusters these probabilities are not close to zero so that it is very unlikely that two of them are identical. On the other hand, since the probability of occurrence of the biased clusters are near zero, some of them may be equal. Therefore this test gives some insights on the bias selection of Section 3.1. As for the score test, we establish in Appendix C.1 a convergence to a chi-square distribution for the bias selection.

However, this test applied to the biased clusters is slightly different from the AIC-type approach we propose in Section 3.1. Indeed, for the test we only look at the bias between s_1 and s_2 and we assume that this bias is constant equal to ρ_n . In particular, we do not make any assumption on the clusters after s_2 . On the other hand, the bias selection as any AIC-type approach requires nested models. This is why for a given model \mathbf{M}_k with s extremal clusters associated to a vector $\mathbf{p} = (p_1, \dots, p_s)$ we assume that all clusters between s and r have the same probability of occurrence.

Remark 4. In both the test and the bias selection we assumed a bias uniformly spread over non-extremal clusters. However it might be more reasonable to consider an isotropic bias, i.e. invariant under a change of direction. Such alternative assumption is approximately consistent with the uniform one we assumed for convenience. But with an isotropic bias we might face some issues regarding the cluster $\{1, \dots, d\}$. Indeed, this cluster might be over-represented even if it is not an extremal cluster because of the topology of the subspace $C_{\{1, \dots, d\}}$. Therefore, if the cluster $\{1, \dots, d\}$ appears as an extremal cluster, the result must be considered with care.

5 Numerical results

5.1 Overview

We apply our algorithm on two numerical examples to demonstrate our method. For both examples our aim is to recover the clusters β on which \mathbf{Z} places mass, that is, the set $\mathcal{S}^*(\mathbf{Z})$. To this end we compare the estimated probability vector $\hat{\zeta}$ with the theoretical one \mathbf{p}^* via the Hellinger distance

$$h(\mathbf{p}^*, \hat{\zeta}) = \frac{1}{\sqrt{2}} \left[\sum_{\beta \in \mathcal{P}_d^*} (p^*(\beta)^{1/2} - \hat{\zeta}(\beta)^{1/2})^2 \right]^{1/2}. \quad (5.1)$$

The closer $h(\mathbf{p}^*, \hat{\zeta})$ is to 0, the better $\hat{\zeta}$ estimates \mathbf{p}^* . In order to compare our method with some existing ones, we also compute the Hellinger distance between the true probabilities $\mathbb{P}(\Theta \in C_\beta)$ and the estimated ones given by the algorithm called DAMEX of Goix et al. (2017) and the two methods of Simpson et al. (2019). We represent the mean Hellinger distance over $N = 100$ simulations. The parameters in the method of Goix et al. (2017) are chosen to be $\epsilon = 0.1$, $k = \sqrt{n}$, and $p = 0.1$, see the notation in their paper. Regarding the methods of Simpson et al. (2019) we use the parameters given by the authors in Section 4.2 of their paper, i.e. we set $\pi = 0.01$, and $p = 0.5$ and u_β to be the 0.75 quantile of observed Q values in region C_β for the first method, and $\delta = 0.5$ and u_β to be the 0.85 quantile of observed Q values in region C_β for the first method. We refer to Simpson et al. (2019) for more insights on these

parameters. Contrary to the aforementioned methods, we recall that MUSCLE does not require any hyperparameter. This is a main advantage from a statistical and computational point of view.

For the first example we consider a Gaussian copula with correlation parameter $\rho < 1$ and Pareto distributed marginals. With such a parameter ρ the random vectors \mathbf{Z} and Θ only places mass on the axes. These directions are all the more identifiable as the parameter ρ is close to 0 (no dependence). The second example is the one developed by [Simpson et al. \(2019\)](#) and deals with a max-mixture of Gaussian and extreme value logistic distributions. The latter distribution is characterized by a parameter $\alpha \in (0, 1)$ and leads to asymptotic dependence. We provide numerical examples for different values of ρ and α . The code can be found at <https://sites.google.com/view/nicolasmeyer/programs>.

5.2 Asymptotic independence

Let $\mathbf{X} = (X_1, \dots, X_d)^\top$ be a random vector with a Gaussian copula with a common correlation parameter $\rho < 1$ and marginal distributions satisfying $\mathbb{P}(X_j > x) = x^{-1}$. Then \mathbf{X} is regularly varying with tail index -1 and its marginals are asymptotically independent, see [Resnick \(1987\)](#), Corollary 5.28. The spectral measure only places mass on the subsets C_β such that $|\beta| = 1$, and so does the distribution of \mathbf{Z} , see [Example 1](#). The aim of our procedure is then to recover these d directions among the $2^d - 1$ clusters.

We first consider $d = 40$, a sample size $n = 30\,000$, and a correlation parameter $\rho = 0.5$. Following [Remark 1](#), we plot the evolution of the penalized log-likelihood [\(2.8\)](#) for a given sample $\mathbf{X}_1, \dots, \mathbf{X}_n$. [Figure 4](#) shows that this quantity first decreases for small values of k before it slightly increases with k . The minimum is reached for $\hat{k} = 1050$ which corresponds to a proportion of extreme values of $\hat{k}/n = 3.5\%$. Regarding the evolution of $\hat{s}(k)$, we observe that the value of $\hat{s}(k)$ remains constant for k close to \hat{k} . It stabilizes around an optimal value of $\hat{s}(\hat{k}) = 41$. Recall that in this example the true clusters corresponds to the $d = 40$ one-dimensional ones. It turns out that the algorithm identifies the 40 one-dimensional clusters. The extra cluster which appears is $\{1, \dots, d\}$ (see [Remark 4](#) for some insights on this cluster).

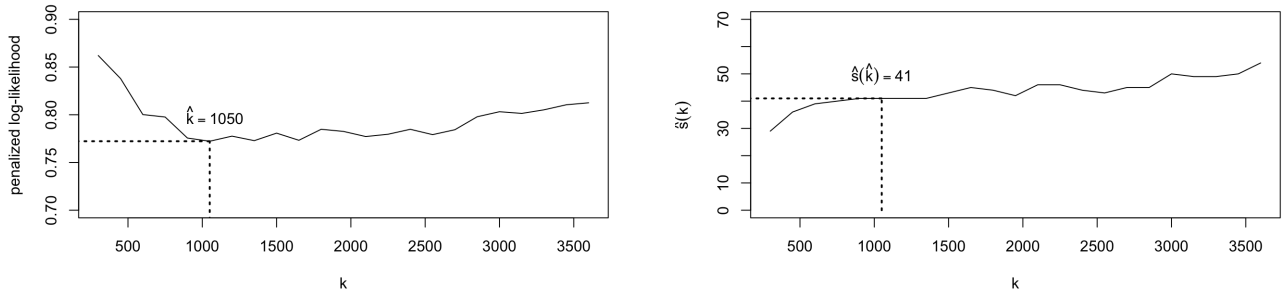


Figure 4: Evolution of the penalized log-likelihood given in [\(2.8\)](#) (left) and of $\hat{s}(k)$ (right) with respect to k . Here $\rho = 0.5$ and $n = 30\,000$.

We then compare the estimated probabilities $\hat{\zeta}$ given by our algorithm with the true ones $p^*(\beta) = 1/d$ for $|\beta| = 1$ and zero elsewhere. To do so, we still consider $d = 40$ and $n = 30\,000$. We apply our procedure for different correlation parameters $\rho \in \{0, 0.25, 0.5, 0.75\}$ and repeat our procedure over $N = 100$ simulations. Then we compute the Hellinger distance defined in [\(5.1\)](#) and we compare ourselves with the approach of [Goix et al. \(2017\)](#). For the two methods proposed by [Simpson et al. \(2019\)](#) it is necessary to compute the empirical mass on all $2^d - 1$ subsets C_β which can not be achieved for such a high dimension. This is why for this example we restrict our comparison with DAMEX.

[Figure 5](#) shows the mean Hellinger distance achieved by our method and the one of [Goix et al. \(2017\)](#) over 100 simulations. We observe that the performance of both methods deteriorates as the value of the

parameter ρ increases. For any $\rho \in \{0, 0.25, 0.5, 0.75\}$ our approach leads to better results than the one of [Goix et al. \(2017\)](#).

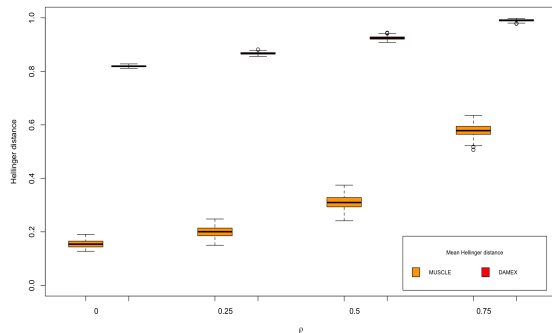


Figure 5: Mean Hellinger distance for $\rho \in \{0, 0.25, 0.5, 0.75\}$ over 100 simulations. In orange: MUSCLE. In red: [Goix et al. \(2017\)](#).

5.3 Max-mixture distribution

The second example is taken from [Simpson et al. \(2019\)](#) and consists of a mixture of Gaussian and extreme-value logistic distributions. For any $\beta \in \mathcal{P}_d^*$, let $\mathbf{A}_\beta \in \mathbb{R}_+^{|\beta|}$ be a random vector with standard Fréchet marginal distributions, and let $\{\mathbf{A}_\beta : \beta \in \mathcal{P}_d^*\}$ be independent random vectors. Then the vector $\mathbf{X} = (X_1, \dots, X_d)^\top$ whose components are defined via

$$X_j = \max_{\beta \in \mathcal{P}_d^* : i \in \beta} \lambda_{i, \beta} X_{j, \beta},$$

with $\lambda_{i, \beta} \in [0, 1]$ and $\sum_{\beta \in \mathcal{P}_d^* : i \in \beta} \lambda_{i, \beta} = 1$, has also standard Fréchet marginal distributions and is regularly varying.

For our simulations we consider the five-dimensional example introduced by [Simpson et al. \(2019\)](#) which we recall for completeness. We consider two bivariate Gaussian copulas with correlation parameter ρ and Fréchet marginals $\mathbf{A}_{\{1,2\}}$ and $\mathbf{A}_{\{4,5\}}$, and three extreme-value logistic copulas with dependence parameter α and Fréchet marginals $\mathbf{A}_{\{1,2,3\}}$, $\mathbf{A}_{\{3,4,5\}}$, and $\mathbf{A}_{\{1,2,3,4,5\}}$. For $\rho < 1$, the Gaussian copula is asymptotically independent and thus the spectral measure defined in (1.1) concentrates on the subsets $C_{\{1\}}$, $C_{\{2\}}$, $C_{\{4\}}$, and $C_{\{5\}}$. For $\alpha \in (0, 1)$ the logistic distribution is asymptotically dependent so that the spectral measure also places mass on the subsets $C_{\{1,2,3\}}$, $C_{\{3,4,5\}}$, and $C_{\{1,2,3,4,5\}}$. Following [Simpson et al. \(2019\)](#), we set

$$\begin{aligned} \lambda_{\{1,2\}} &= (5, 5)/7, & \lambda_{\{4,5\}} &= (5, 5)/7 \\ \lambda_{\{1,2,3\}} &= (1, 1, 3)/7, & \lambda_{\{3,4,5\}} &= (3, 1, 1)/7, & \lambda_{\{1,2,3,4,5\}} &= (1, 1, 1, 1, 1)/7, \end{aligned}$$

so that equal mass is assigned to each of the seven aforementioned subsets. In order to compute the mass the distribution of \mathbf{Z} assigns to every subset C_β we start from the distribution of Θ and use Monte-Carlo simulation. We then compare these probabilities with their estimated ones $\hat{\zeta}$ given by MUSCLE.

We run our algorithm for different values of $\rho \in \{0, 0.25, 0.5, 0.75\}$ and $\alpha \in \{0.1, 0.2, \dots, 0.9\}$. Figure 6 shows the average mean Hellinger distance for our method, the one of [Goix et al. \(2017\)](#), and the two of [Simpson et al. \(2019\)](#) over 100 simulations. Our method provides a mean Hellinger distance which stabilizes between 0.2 and 0.3 for all values of ρ and α . For $\alpha \leq 0.7$ the distance slightly decreases with alpha, while it increases for $\alpha \geq 0.8$. The standard deviation is quite small for $\alpha \leq 0.7$ and then increases

with α . Regarding the approach of [Goix et al. \(2017\)](#), the mean Hellinger distance tends to increase with α and with ρ . The smallest values is obtained for $\rho \in \{0, 0.25\}$ and for small α . The estimation particularly deteriorates for $\rho = 0.75$. Finally both methods proposed by [Simpson et al. \(2019\)](#) provide a mean Hellinger distance which increases with α and ρ . The second one seems to provide almost always better results than the first one.

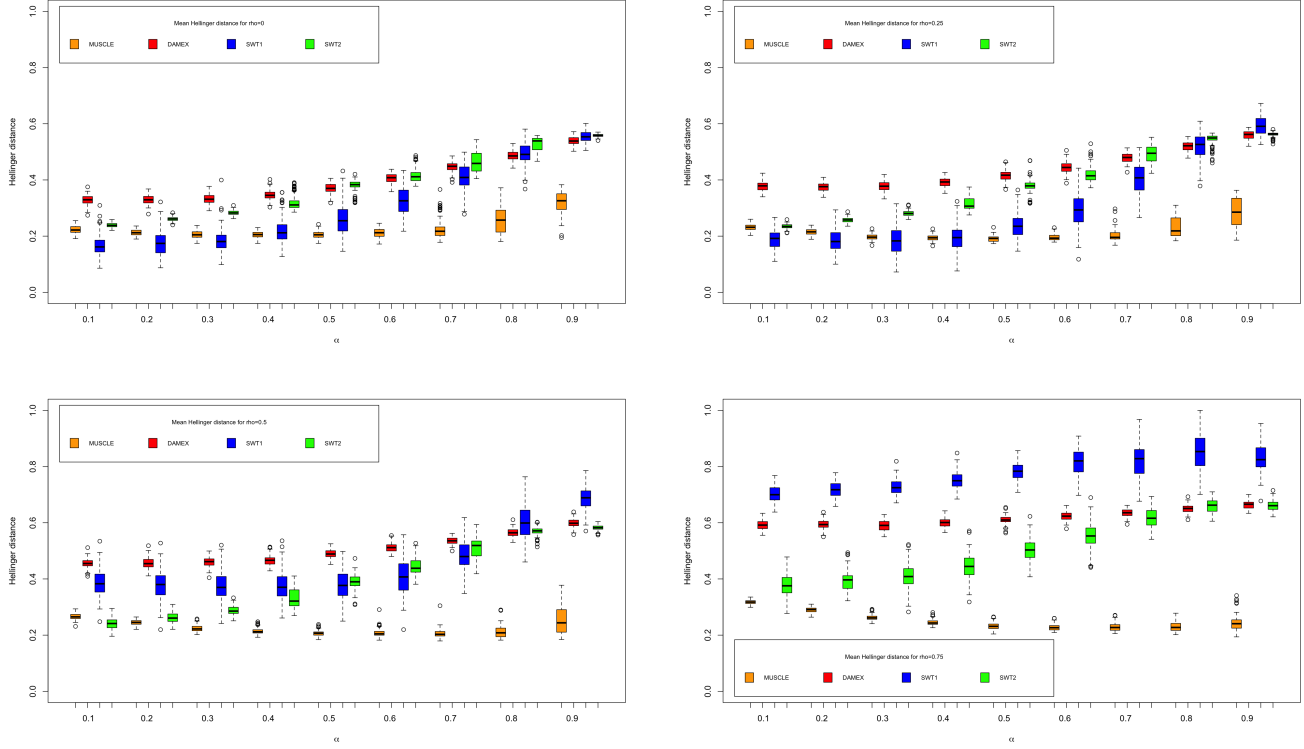


Figure 6: Mean Hellinger distance over 100 simulations for $\rho = 0$ (top left), $\rho = 0.25$ (top right), $\rho = 0.5$ (bottom left), $\rho = 0.75$ (bottom right). The abbreviation SWT1 (resp. SWT2) refers to the first (resp. second) method of [Simpson et al. \(2019\)](#)

While all methods provided by [Goix et al. \(2017\)](#) and [Simpson et al. \(2019\)](#) deteriorate when ρ or α increase, our procedure provides results which stabilize around a mean Hellinger distance of 0.2. This distance is the smallest one for $\rho = 0.5$ and $\rho = 0.75$ for all α compared to the one of the three other methods. For small ρ , MUSCLE better performs for large α . For small α the second method of [Simpson et al. \(2019\)](#) provides better results than our approach, while its standard deviation is larger. It turns out that except for small α with $\rho = 0$ and $\rho = 0.25$ our algorithm better detects the extremal clusters. These performances are all the more remarkable than they are achieved without any hyperparameter.

6 Application to real-world data

6.1 Preprocessing for real-world data

In Section 5 we considered examples with standard Pareto or standard Fréchet marginal distributions so that the tail index (see Equation (1.1)) of the considered vector is equal to 1. The influence of this index on the extremal clusters has been studied on some numerical results by [Meyer and Wintenberger \(2021\)](#). It turns out that a large tail index does not provide accurate results while a small one highlights

one-dimensional clusters, see Remark 11 in their article. A tail index of $\alpha = 1$ seems to provide the best results.

For real-world data the estimation of the tail index of a sample $\mathbf{x}_1, \dots, \mathbf{x}_n$ is achieved with a Hill plot (Hill (1975)). It consists in plotting

$$\hat{\alpha}(k) = \left(\frac{1}{k} \sum_{j=1}^k \log(|\mathbf{x}|_{(j)}) - \log(|\mathbf{x}|_{(k)}) \right)^{-1}, \quad k = 2, \dots, n,$$

where $|\mathbf{x}|_{(j)}$ denotes the order statistics of the norms $|\mathbf{x}_1|, \dots, |\mathbf{x}_n|$, i.e.

$$|\mathbf{x}|_{(1)} \geq \dots \geq |\mathbf{x}|_{(n)},$$

and to choose $\hat{\alpha}$ as the value around which the plot stabilizes.

Then, we consider the power transform $\mathbf{x}'_j = (\mathbf{x}_j)^{\hat{\alpha}}$. This transformation highlights the tail structure of the data without modifying the support of the spectral measure, see Meyer and Wintenberger (2021), Remark 8. It differs from the standardization discussed in the introduction for which the vectors are normalized via a rank transform.

6.2 Wind speed data

The data correspond to the daily-average wind speed for 1961-1978 at 12 synoptic meteorological stations in the Republic of Ireland ($n = 6574$, $d = 12$). They are available at <http://lib.stat.cmu.edu/datasets/> and have been analyzed in detail by Haslett and Raftery (1989). The stations are the following ones: Malin Head (Mal), Belmullet (Bel), Clones (Clo), Claremorris (Cla), Mullingar (Mul), Dublin (Dub), Shannon (Sha), Birr (Bir), Kilkenny (Kil), Valentia (Val), Roche's Pt. (Rpt), Rosslare (Ros). Seven of these stations are along the sea: Belmullet (west), Dublin (east), Malin Head (north), Roche's Pt. (south), Rosslare (east), Shannon (west), and Valentia (southwest). The five other stations are more than 50 kilometers away from the coast. We refer to Haslett and Raftery (1989) for a map of the stations.

The preprocessing of the data provides a Hill estimator of $\hat{\alpha} = 10.7$. Before applying MUSCLE and similarly to Section 5 we plot the evolution of the penalized log-likelihood in (2.8) as a function of k . The optimal value $\hat{k} = 460$ is clearly identified and corresponds to a proportion $\hat{k}/n = 7\%$. This choice of \hat{k} leads to a number of clusters $\hat{s}(\hat{k}) = 11$. Note that contrary to the numerical examples the value of $\hat{s}(k)$ does not stabilize when k is close to \hat{k} .

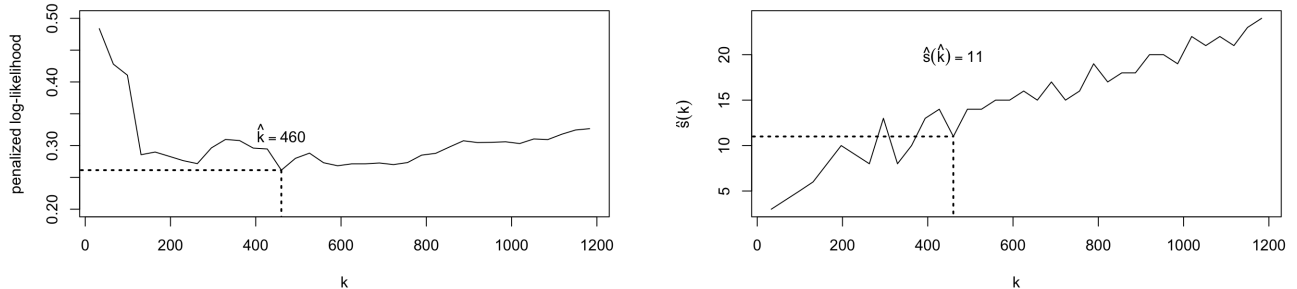


Figure 7: Evolution of the penalized log-likelihood given in (2.8) (left) and of $\hat{s}_n(k)$ (right) with respect to k for the wind speed data.

MUSCLE provides 11 extremal clusters which correspond to 6 stations: Belmullet, Malin Head, Roche's Pt., Rosslare, Dublin, and Shannon. All of these stations are located close to the sea, where

wind speed is likely to be higher than in inland cities. The only coastal city which does not appear in the extremal clusters is Valentia which is located more than 200 kilometers away from the other stations.

The 11 clusters and their inclusions are illustrated on Figure 8. The algorithm exhibits low-dimensional clusters as the largest ones are of dimension 3. The northern station Malin Head appears in all multivariate clusters. Most of the clusters are related to a specific localization: $\{\text{Sha}, \text{Bel}, \text{Mal}\}$ and $\{\text{Bel}, \text{Mal}\}$ correspond to stations in the north/west, $\{\text{Mal}, \text{Dub}\}$ and $\{\text{Mal}, \text{Ros}\}$ to stations in the north/east. Three extremal clusters gather the northern station Malin Head and the southern one Roche’s Pt.

We conclude that the aforementioned 11 clusters correspond to subsets C_β which gather the mass of the angular vector \mathbf{Z} . In particular, the subsets related to the clusters $\{\text{Sha}, \text{Bel}, \text{Mal}\}$, $\{\text{Rpt}, \text{Bel}, \text{Mal}\}$, $\{\text{Rpt}, \text{Ros}, \text{Mal}\}$, and $\{\text{Dub}, \text{Mal}\}$ gather some mass of \mathbf{Z} and are not included in larger subsets on which \mathbf{Z} places mass. Following Meyer and Wintenberger (2021), Theorem 2, these *maximal* subsets also concentrate the mass of the spectral measure. The remaining clusters, which correspond to *non-extremal* subsets, contain almost all the station Malin Head. We interpret this as follows: among the maximal subsets the wind speed in Malin Head is likely to be larger than in the other stations. We also refer to Meyer and Wintenberger (2021), Section 3.2, for a discussion on maximal and non-maximal subsets. A separate study can then be conducted on each group of stations for which standard methods for low-dimensional extremes can be applied, see Coles and Tawn (1991), Einmahl et al. (1993), Einmahl et al. (1997), Einmahl and Segers (2009).

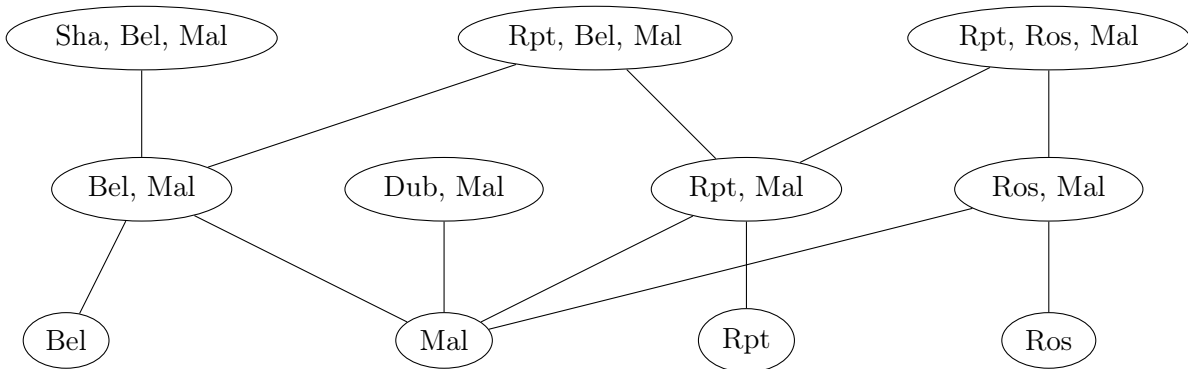


Figure 8: Representation of the 11 clusters and their inclusions.

In order to study the remaining stations, we remove the 6 extremal stations and reapply our procedure. MUSCLE then provides 16 clusters:

- Four one-dimensional clusters: Val, Clo, Cla, Mul.
- Five two-dimensional clusters: $\{\text{Val}, \text{Cla}\}$, $\{\text{Val}, \text{Clo}\}$, $\{\text{Val}, \text{Mul}\}$, $\{\text{Mul}, \text{Clo}\}$, $\{\text{Cla}, \text{Clo}\}$.
- Two three-dimensional clusters: $\{\text{Val}, \text{Cla}, \text{Clo}\}$, $\{\text{Val}, \text{Mul}, \text{Clo}\}$.
- and other clusters: $\{\text{Val}, \text{Cla}, \text{Mul}, \text{Clo}\}$, $\{\text{Val}, \text{Bir}, \text{Cla}, \text{Clo}\}$, $\{\text{Val}, \text{Bir}, \text{Cla}, \text{Mul}, \text{Clo}\}$, $\{\text{Val}, \text{Kil}, \text{Bir}, \text{Cla}, \text{Mul}, \text{Clo}\}$.

The station Valentia appears in almost all of these clusters. It is the only remaining coastal station, the other ones are inland ones. No particular tail dependence structure appears for these non-extremal stations. In particular, the largest clusters $\{\text{Val}, \text{Kil}, \text{Bir}, \text{Cla}, \text{Mul}, \text{Clo}\}$ indicates that it is likely that the wind speed in all of these six stations is simultaneously large.

6.3 Extreme variability for financial data

In the second example we deal with financial data. The data set we use corresponds to the value-average daily returns of 49 industry portfolios compiled and posted as part of the Kenneth French Data

Library. They are available at https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. A related study on a similar dataset has been conducted by Cooley and Thibaud (2019). We restrict our study to the period 1970 – 2019 which provides $n = 12\,613$ observations denoted by $\mathbf{x}_1^{\text{obs}}, \dots, \mathbf{x}_n^{\text{obs}} \in \mathbb{R}^{49}$. Our goal is to study the variability of these returns so that we take the componentwise absolute value $\mathbf{x}_j = |\mathbf{x}_j^{\text{obs}}|$ of the data. Thus, we study the non-negative vectors $\mathbf{x}_1, \dots, \mathbf{x}_n$ in \mathbb{R}_+^d with $n = 12\,613$ and $d = 49$. Following Section 6.1, we consider the vectors $\mathbf{x}_j' = (\mathbf{x}_j)^{\hat{\alpha}}$, where $\hat{\alpha} = 2.99$ is the Hill estimator of the sample $|\mathbf{x}_1|, \dots, |\mathbf{x}_n|$.

Following Remark 1, we plot the evolution of the estimator of the Kullback-Leibler divergence in (2.8) as a function of k . We see on Figure 9 that this estimator decreases until it reaches a minimal value for $\hat{k} = 441$, before increasing for $k \geq \hat{k}$. The level \hat{k} corresponds to a proportion $\hat{k}/n = 3\%$ and leads to a number of extremal clusters $\hat{s}(\hat{k}) = 14$. Contrary to the numerical results and as for the wind speed data, we do not observe a range of k for which the minimal value $\hat{s}(k)$ remains approximately constant.

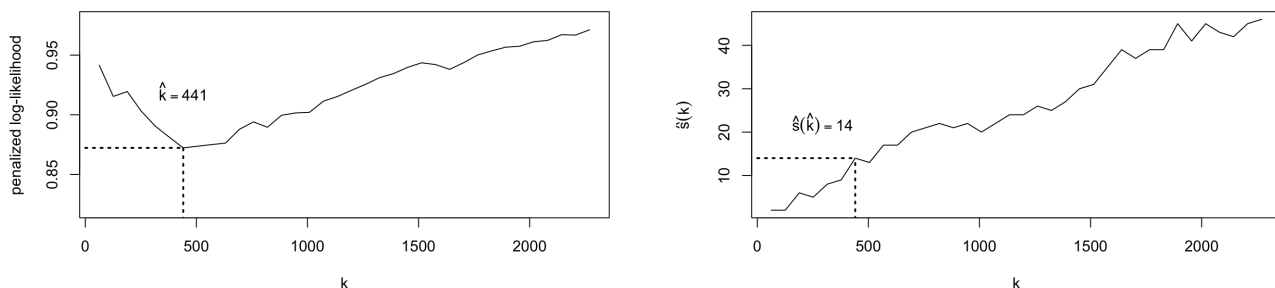


Figure 9: Evolution with respect to k of the penalized log-likelihood given in (2.8) (left) and of $\hat{s}_n(k)$ (right) for the financial data.

MUSCLE provides $\hat{s}(\hat{k}) = 14$ extremal clusters which gather 12 portfolios. These clusters and their inclusions are represented in Figure 10. The number of identified clusters is much smaller compared to the total number $2^{49} \approx 10^{15}$. Besides these clusters are at most three-dimensional so that our procedure drastically reduces the dimension of the study. Most of the extremal portfolios which appear in the clusters correspond to "office/executive" sectors, such as Health, Software, Hardware, Banks, Finance, Electronic Equipment (Chips), Real Estate. Some other clusters group portfolios related to heavy industries, such as Steel, Coal, and Gold. The only clusters gathering a heavy industry and service sectors are $\{\text{Coal, Banks}\}$ and $\{\text{Coal, Banks, Fin}\}$. The tail dependence of the variability of these different kinds of portfolios may result from the financing of the coal industry by several big banks, see Raval et al. (2020).

As for the wind speed data, we conclude that the 14 clusters given by MUSCLE correspond to subsets C_β which gather the mass of \mathbf{Z} . Among them, 8 are maximal: the spectral measure places mass on the associated C_β . Here again, standard approaches which hold for low-dimensional extremes can now be applied on these subsets.

After removing the 12 extremal components we reapply MUSCLE to obtain the dependence structure of the non-extremal portfolios. The algorithm provides a unique cluster with all 37 remaining portfolios. Hence these portfolios tend to have a dependent tail structure: their extreme variability is strongly correlated.

7 Conclusion

The statistical analysis introduced in this article provides a new approach to detect the extremal directions of a multivariate random vector \mathbf{X} . This method relies on the notion of sparse regular variation which

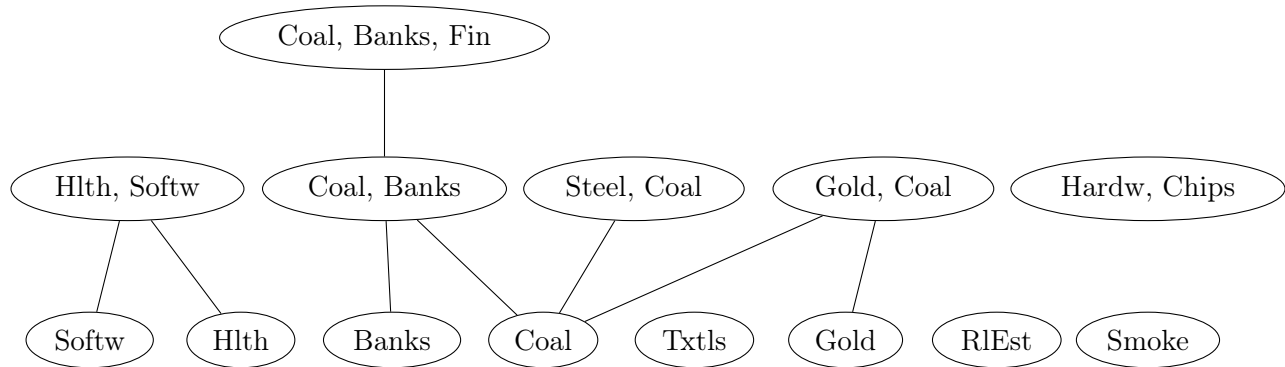


Figure 10: Representation of the 14 clusters and their inclusions. The abbreviations are the following ones: Softw = Computer Software, Txtls = Textiles, Hlth = Healthcare, REEst = Real Estate, Hardw = Hardware, Chips = Electronic Equipment, Fin = Finance.

better highlights the tail dependence of \mathbf{X} . Several convergence results are established in Section 4 and are used to build a rigorous statistical method based on model selection. This approach provides not only the clusters of directions on which the extremes of \mathbf{X} gather but also an optimal threshold above which the data are considered as extreme values. The latter issue has always been challenging and no theoretical-based procedure has been provided in a multivariate setting yet, even if it has been the subject of much attention in the literature. The model selection is done with an AIC-type minimization whose penalization allows to reduce the number of selected subsets. Including the choice of an appropriate level k then entails a multiplicative penalization. This approach leads to the parameter-free algorithm MUSCLE whose purpose is to recover the extremal clusters of a sample of iid sparsely regularly varying random vectors $\mathbf{X}_1, \dots, \mathbf{X}_n$.

The absence of any hyperparameter is a main difference with the existing methods (Goix et al. (2017), Simpson et al. (2019), Chiapino and Sabourin (2016), Chiapino et al. (2019)). In these articles the level k is always chosen as a function of n , usually $k = \sqrt{n}$. Moreover, the estimation of the tail dependence in such papers often relies on some hyperparameters which allow to reduce the number of extremal clusters β . Another main advantage of our procedure is that it is still efficient for large d . This follows from the expected linear-time algorithm introduced by Duchi et al. (2008) to compute the Euclidean projection, see Appendix A.

The numerical experiments provide promising results for asymptotic independent and dependent cases. Our algorithm is tested on samples of dimension $d = 40$ for which the number of possible clusters is $2^{40} - 1$. Despite this very high number of potential subsets C_β we succeed in capturing the extremal clusters and do not recover many extra subsets. In particular, the numerical results show that our approach is quite accurate for asymptotically independent data. Regarding the max-mixture example, our algorithm provides better results than the ones of Goix et al. (2017) and Simpson et al. (2019) for ρ close to 1, or small ρ and α close to 1. Moreover the results do not vary a lot with ρ and α . Finally the application of our algorithm on two types of real-world examples highlights sparse clusters and thus reduces the dimension of the study. We obtain a sparse tail dependence structure for extreme wind speed in different stations, as well as for the extreme variability of several industry portfolios. This reinforces the relevance of our approach for reducing the dimension in Extreme Value Theory.

A Algorithm

We introduce here the expected linear-time algorithm given in Duchi et al. (2008). It is based on a random selection of the coordinates.

Algorithm 2: Expected linear-time projection onto the positive sphere $\mathbb{S}_+^{d-1}(z)$.

Data: A vector $\mathbf{v} \in \mathbb{R}_+^d$ and a scalar $z > 0$

Result: The projected vector $\mathbf{w} = \pi(\mathbf{v})$

Initialize $U = \{1, \dots, d\}$, $s = 0$, $\rho = 0$;

while $U \neq \emptyset$ **do**

 Pick $k \in U$ at random;

 Partition U : $G = \{j \in U, v_j \geq v_k\}$ and $L = \{j \in U, v_j < v_k\}$;

 Calculate $\Delta\rho = |G|$, $\Delta s = \sum_{j \in G} v_j$;

if $(s + \Delta s) - (\rho + \Delta\rho)v_k < z$ **then**

$s = s + \Delta s$;

$\rho = \rho + \Delta\rho$;

$U \leftarrow L$;

else

$U \leftarrow G \setminus \{k\}$;

end

end

Set $\eta = (s - z)/\rho$;

Output: \mathbf{w} s.t. $w_i = v_i - \eta$.

B Proof of Theorem 1

We consider the vector $\mathbf{V}_{n, \mathcal{S}_k^\infty} \in \mathbb{R}^{s_\infty}$ whose components are

$$V_{n, \beta} = \frac{1}{\sqrt{k_n p_n(\beta)}} \left(\mathbf{1}\{\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n\} - \frac{k_n}{n} p_n(\beta) \right), \quad \beta \in \mathcal{S}_k^\infty.$$

This vector has null expectation. We denote by $\Sigma_n \in \mathcal{M}_{s_\infty}(\mathbb{R})$ its covariance matrix. The diagonal entries correspond to the variance of a Bernoulli distribution:

$$\Sigma_n(\beta, \beta) = \frac{1}{k_n p_n(\beta)} \mathbb{P}(\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n) [1 - \mathbb{P}(\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n)] = \frac{1}{n} - \frac{k_n}{n^2} p_n(\beta).$$

Regarding the non-diagonal entries they can be computed as follows:

$$\begin{aligned} & \Sigma_n(\beta, \beta') \\ &= \mathbb{E}[V_{n, \beta} V_{n, \beta'}] \\ &= \frac{1}{k_n \sqrt{p_n(\beta) p_n(\beta')}} \left(\mathbb{E}[\mathbf{1}\{\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n\} \mathbf{1}\{\pi(\mathbf{X}/u_n) \in C_{\beta'}, |\mathbf{X}| > u_n\}] \right. \\ &\quad \left. - \frac{k_n}{n} p_n(\beta) \mathbb{E}[\mathbf{1}\{\pi(\mathbf{X}/u_n) \in C_{\beta'}, |\mathbf{X}| > u_n\}] - \frac{k_n}{n} p_n(\beta') \mathbb{E}[\mathbf{1}\{\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n\}] + \frac{k_n^2}{n^2} p_n(\beta) p_n(\beta') \right) \\ &= -\frac{1}{k_n \sqrt{p_n(\beta) p_n(\beta')}} \frac{k_n^2}{n^2} p_n(\beta) p_n(\beta') \\ &= -\frac{k_n}{n^2} \sqrt{p_n(\beta) p_n(\beta')}. \end{aligned}$$

This implies that we can rewrite the covariance matrix Σ_n as

$$\Sigma_n = \frac{1}{n} Id_{s_\infty} - \frac{k_n}{n^2} \sqrt{\mathbf{P}_{n, \mathcal{S}_k^\infty}} \cdot \sqrt{\mathbf{P}_{n, \mathcal{S}_k^\infty}}^\top,$$

where the square root is meant componentwise. In particular we have $n\Sigma_n \rightarrow Id_{s_\infty}$ when $n \rightarrow \infty$.

Consider now a triangular array $\mathbf{V}_n^{(1)}, \dots, \mathbf{V}_n^{(n)}$ with the same distribution as $\mathbf{V}_{n, \mathcal{S}_k^\infty}$. We prove that this triangular array satisfies Lindeberg's condition:

$$\sum_{j=1}^n \mathbb{E} \left[\frac{1}{k_n} \max_{\beta} \frac{1}{p_n(\beta)} \left| \mathbb{1}\{\pi(\mathbf{X}_j/u_n) \in C_\beta, |\mathbf{X}_j| > u_n\} - \frac{k_n}{n} p_n(\beta) \right|^2 \mathbb{1}_{\{\max_{\beta} |V_{n,\beta}^{(j)}| > \epsilon\}} \right] \rightarrow 0, \quad n \rightarrow \infty,$$

for all $\epsilon > 0$, or equivalently that

$$\mathbb{E} \left[\frac{n}{k_n} \max_{\beta} \frac{1}{p_n(\beta)} \left| \mathbb{1}\{\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n\} - \frac{k_n}{n} p_n(\beta) \right|^2 \mathbb{1}_{\{\max_{\beta} |V_{n,\beta}^{(1)}| > \epsilon\}} \right] \rightarrow 0, \quad n \rightarrow \infty. \quad (\text{B.1})$$

Fix $\epsilon > 0$. Recall that \mathcal{S}_k^∞ gathers all clusters β such that $k_n p_n(\beta) \rightarrow \infty$. Thus, there exists n_0 such that for all $n \geq n_0$,

$$\max_{\beta \in \mathcal{S}_k^\infty} \left| \mathbb{1}\{\pi(\mathbf{X}/u_n) \in C_\beta, |\mathbf{X}| > u_n\} - \frac{k_n}{n} p_n(\beta) \right| \leq \epsilon k_n \min_{\beta \in \mathcal{S}_k^\infty} p_n(\beta),$$

since the term on the left-hand side is always bounded by 1. This implies that for n large enough, the inequality $\max_{\beta} |V_{n,\beta}| > \epsilon$ is never satisfied. Hence, Lindeberg's condition in (B.1) holds and yields to the following convergence

$$\sum_{j=1}^n \mathbf{V}_n^{(j)} \xrightarrow{d} \mathcal{N}(0, Id_{s_\infty}), \quad n \rightarrow \infty.$$

This convergence can be rephrased as

$$\sqrt{k_n} \text{Diag}(\mathbf{p}_{n, \mathcal{S}_k^\infty})^{-1/2} \left(\frac{\mathbf{T}_{n, \mathcal{S}_k^\infty}}{k_n} - \mathbf{p}_{n, \mathcal{S}_k^\infty} \right) \xrightarrow{d} \mathcal{N}(0, Id_{s_\infty}), \quad n \rightarrow \infty,$$

which proves (4.3).

To obtain the convergence (4.4), it suffices to restrict the previous convergence to the clusters $\beta \in \mathcal{S}^*(\mathbf{Z})$ and to notice that

$$\text{Diag}(\mathbf{p}_{n, \mathcal{S}^*(\mathbf{Z})})^{1/2} \text{Diag}(\mathbf{p}_{\mathcal{S}^*(\mathbf{Z})})^{-1/2} \rightarrow Id_{s^*}, \quad n \rightarrow \infty.$$

Finally, to prove (4.6) it suffices to show that $\sqrt{k_n} \text{Diag}(\mathbf{p}_{\mathcal{S}^*(\mathbf{Z})})^{-1/2} (\mathbf{p}_{n, \mathcal{S}^*(\mathbf{Z})} - \mathbf{p}_{\mathcal{S}^*(\mathbf{Z})}) \rightarrow 0$ which holds true under assumption (4.5).

C Model selection

C.1 Bias selection

We recall the expression of the log-likelihood $\log L_{\mathbf{M}_k}$ as a function of $\mathbf{p} := (p_1, \dots, p_s)^\top \in \mathbb{R}^s$:

$$\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n) = \log(k!) - \sum_{i=1}^{2^d-1} \log(T_{n,i}!) + \sum_{i=1}^s T_{n,i} \log(p_i) + \left(\sum_{i=s+1}^r T_{n,i} \right) \log \left(\frac{1 - \sum_{j=1}^s p_j}{r - s} \right).$$

The expectation of this log-likelihood is then given by

$$\mathbb{E}[\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n)] = k! - \sum_{i=1}^{2^d-1} \mathbb{E}[\log(T_{n,i}!)] + \sum_{i=1}^s p_{n,i} \log p_i + \sum_{i=s+1}^r p_{n,i} \log \left(\frac{1 - \sum_{j=1}^s p_j}{r - s} \right).$$

A similar computation as for the maximum likelihood estimator in Section 3.1 entails that this expectation is maximized for

$$\mathbf{p} = \tilde{\mathbf{p}} := \left(\frac{p_{n,1}}{\sum_{j=1}^r p_{n,j}}, \dots, \frac{p_{n,s}}{\sum_{j=1}^r p_{n,j}} \right)^\top \in \mathbb{R}^s.$$

note that since we work under the event $\mathcal{S}_k^\infty = \hat{\mathcal{S}}_n$ we have $r = \hat{s}_n \geq s^*$ and thus $\sum_{j=1}^{\hat{s}_n} p_{n,j} \geq \sum_{j=1}^{s^*} p_{n,j} \rightarrow 1$.

The first step of the bias selection consists in establishing a Taylor expansion for the estimator in Equation (3.5).

Lemma 2. *There exists $c_1 \in (0, 1)$ such that*

$$KL(\mathbf{P}_k \| \mathbf{M}_k) |_{\mathbf{p}=\hat{\mathbf{p}}} = KL(\mathbf{P}_k \| \mathbf{M}_k) |_{\mathbf{p}=\tilde{\mathbf{p}}} + \frac{1}{2} (\hat{\mathbf{p}} - \tilde{\mathbf{p}})^\top \frac{\partial^2}{\partial \mathbf{p}^2} \mathbb{E}[-\log L_{\mathbf{M}_k}(\mathbf{p}, \mathbf{T}_n)] |_{c_1 \hat{\mathbf{p}} + (1-c_1)\tilde{\mathbf{p}}} (\hat{\mathbf{p}} - \tilde{\mathbf{p}}). \quad (\text{C.1})$$

Since the quantity $\tilde{\mathbf{p}}$ is deterministic, the first term of the right-hand side in (C.1) can be written as

$$KL(\mathbf{P}_k \| \mathbf{M}_k) |_{\mathbf{p}=\tilde{\mathbf{p}}} = \mathbb{E}[\log L_{\mathbf{P}}(\mathbf{T}_n)] - \mathbb{E}[\log L_{\mathbf{M}_k}(\tilde{\mathbf{p}}; \mathbf{T}_n)].$$

The idea is then to provide a Taylor expansion of $\log L_{\mathbf{M}_k}(\tilde{\mathbf{p}}; \mathbf{T}_n)$ around the vector $\hat{\mathbf{p}}$. This is the purpose of the following lemma.

Lemma 3. *There exists $c_2 \in (0, 1)$ such that*

$$\log L_{\mathbf{M}_k}(\tilde{\mathbf{p}}; \mathbf{T}_n) = \log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n) + \frac{1}{2} (\tilde{\mathbf{p}} - \hat{\mathbf{p}})^\top \frac{\partial^2}{\partial \mathbf{p}^2} \log L_{\mathbf{M}_k}(c_2 \tilde{\mathbf{p}} + (1 - c_2)\hat{\mathbf{p}}; \mathbf{T}_n) (\tilde{\mathbf{p}} - \hat{\mathbf{p}}). \quad (\text{C.2})$$

Lemmas 2 and 3 are a consequence of the following result known as "Cauchy's Mean-Value Theorem" (see Hille (1964) for a proof).

Lemma 4. *Let f and g be two continuous functions on the closed interval $[a, b]$, $a < b$, and differentiable on the open interval (a, b) . Then there exists some $c \in (a, b)$ such that*

$$(f(b) - f(a))g'(c) = (g(b) - g(a))f'(c).$$

Proof of Lemma 2. Let f be the function defined as $f(t) = h(t\hat{\mathbf{p}} + (1-t)\tilde{\mathbf{p}})$ for $t \in [0, 1]$, where h is defined as

$$h(\mathbf{p}) = KL(\mathbf{P}_k \| \mathbf{M}_k) + \frac{\partial}{\partial \mathbf{p}} KL(\mathbf{P}_k \| \mathbf{M}_k)(\hat{\mathbf{p}} - \mathbf{p}).$$

Some short calculations give the following relations:

$$\begin{aligned} f(1) &= h(\hat{\mathbf{p}}) = KL(\mathbf{P}_k \| \mathbf{M}_k) |_{\mathbf{p}=\hat{\mathbf{p}}}, \\ f(0) &= h(\tilde{\mathbf{p}}) = KL(\mathbf{P}_k \| \mathbf{M}_k) |_{\mathbf{p}=\tilde{\mathbf{p}}} + \frac{\partial}{\partial \mathbf{p}} KL(\mathbf{P}_k \| \mathbf{M}_k) |_{\mathbf{p}=\tilde{\mathbf{p}}} (\hat{\mathbf{p}} - \tilde{\mathbf{p}}) \\ &= KL(\mathbf{P}_k \| \mathbf{M}_k) |_{\mathbf{p}=\tilde{\mathbf{p}}} - \underbrace{\frac{\partial}{\partial \mathbf{p}} \mathbb{E}[\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n)] |_{\mathbf{p}=\tilde{\mathbf{p}}} (\hat{\mathbf{p}} - \tilde{\mathbf{p}})}_{=0 \text{ by definition of } \tilde{\mathbf{p}}} = KL(\mathbf{P}_k \| \mathbf{M}_k) |_{\mathbf{p}=\tilde{\mathbf{p}}}, \\ f'(t) &= \frac{\partial h}{\partial \mathbf{p}}(t\hat{\mathbf{p}} + (1-t)\tilde{\mathbf{p}})(\hat{\mathbf{p}} - \tilde{\mathbf{p}}) \\ &= (\hat{\mathbf{p}} - [t\hat{\mathbf{p}} + (1-t)\tilde{\mathbf{p}}])^\top \frac{\partial^2}{\partial \mathbf{p}^2} KL(\mathbf{P}_k \| \mathbf{M}_k) |_{t\hat{\mathbf{p}}+(1-t)\tilde{\mathbf{p}}} (\hat{\mathbf{p}} - \tilde{\mathbf{p}}) \\ &= (1-t)(\hat{\mathbf{p}} - \tilde{\mathbf{p}})^\top \frac{\partial^2}{\partial \mathbf{p}^2} \mathbb{E}[-\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n)] |_{t\hat{\mathbf{p}}+(1-t)\tilde{\mathbf{p}}} (\hat{\mathbf{p}} - \tilde{\mathbf{p}}). \end{aligned}$$

We apply Lemma 4 to the functions f and $g : t \mapsto (t - 1)^2$. There exists $c_1 \in (0, 1)$ such that $(f(1) - f(0))g'(c_1) = (g(1) - g(0))f'(c_1)$, i.e.

$$KL(\mathbf{P}_k \parallel \mathbf{M}_k)|_{\mathbf{p}=\hat{\mathbf{p}}} - KL(\mathbf{P}_k \parallel \mathbf{M}_k)|_{\mathbf{p}=\tilde{\mathbf{p}}} = -\frac{1}{2}(\hat{\mathbf{p}} - \tilde{\mathbf{p}})^\top \frac{\partial^2}{\partial \mathbf{p}^2} \mathbb{E}[\log L_{\mathbf{M}_k}(\mathbf{T}_n)]|_{c_1 \hat{\mathbf{p}} + (1-c_1)\tilde{\mathbf{p}}} (\hat{\mathbf{p}} - \tilde{\mathbf{p}}).$$

This concludes the proof. \square

We now prove Lemma 3.

Proof of Lemma 3. Consider $f(t) = h(t\tilde{\mathbf{p}} + (1-t)\hat{\mathbf{p}})$, for $t \in [0, 1]$ where h is defined as

$$h(\mathbf{p}) = \log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n) + \frac{\partial}{\partial \mathbf{p}} \log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n)(\tilde{\mathbf{p}} - \mathbf{p}).$$

After some calculations we obtain

$$\begin{aligned} f(1) &= h(\tilde{\mathbf{p}}) = \log L_{\mathbf{M}_k}(\tilde{\mathbf{p}}; \mathbf{T}_n), \\ f(0) &= h(\hat{\mathbf{p}}) = \log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n) + \underbrace{\frac{\partial}{\partial \mathbf{p}} \log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n)(\tilde{\mathbf{p}} - \hat{\mathbf{p}})}_{=0 \text{ by definition of } \hat{\mathbf{p}}}, \\ f'(t) &= \frac{\partial h}{\partial \mathbf{p}}(t\tilde{\mathbf{p}} + (1-t)\hat{\mathbf{p}})(\tilde{\mathbf{p}} - \hat{\mathbf{p}}) \\ &= (\tilde{\mathbf{p}} - [t\tilde{\mathbf{p}} + (1-t)\hat{\mathbf{p}}])^\top \frac{\partial^2}{\partial \mathbf{p}^2} \log L_{\mathbf{M}_k}(t\tilde{\mathbf{p}} + (1-t)\hat{\mathbf{p}}; \mathbf{T}_n)(\tilde{\mathbf{p}} - \hat{\mathbf{p}}) \\ &= (1-t)(\tilde{\mathbf{p}} - \hat{\mathbf{p}})^\top \frac{\partial^2}{\partial \mathbf{p}^2} \log L_{\mathbf{M}_k}(t\tilde{\mathbf{p}} + (1-t)\hat{\mathbf{p}}; \mathbf{T}_n)(\tilde{\mathbf{p}} - \hat{\mathbf{p}}). \end{aligned}$$

We apply Lemma 4 to the functions f and $g : t \mapsto (t - 1)^2$. There exists $c_2 \in (0, 1)$ such that $(f(1) - f(0))g'(c_2) = (g(1) - g(0))f'(c_2)$, i.e.

$$\log L_{\mathbf{M}_k}(\tilde{\mathbf{p}}; \mathbf{T}_n) - \log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n) = \frac{1}{2}(\tilde{\mathbf{p}} - \hat{\mathbf{p}})^\top \frac{\partial^2}{\partial \mathbf{p}^2} \log L_{\mathbf{M}_k}(c_2\tilde{\mathbf{p}} + (1-c_2)\hat{\mathbf{p}}; \mathbf{T}_n)(\tilde{\mathbf{p}} - \hat{\mathbf{p}}).$$

This concludes the proof. \square

After taking the expectation with respect to $\hat{\mathbf{p}}$ in Equations (C.1) and (C.2), and after combining these equations, we obtain the following expression for the expectation of the estimator in (3.5):

$$\begin{aligned} \mathbb{E}[KL(\mathbf{P}_k \parallel \mathbf{M}_k)|_{\mathbf{p}=\hat{\mathbf{p}}}] &= \mathbb{E}[\log L_{\mathbf{P}_k}(\mathbf{T}_n)] - \mathbb{E}[\log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n)] \\ &= \underbrace{\mathbb{E}\left[-\frac{1}{2}(\hat{\mathbf{p}} - \tilde{\mathbf{p}})^\top \left(\frac{\partial^2}{\partial \mathbf{p}^2} \log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n) \Big|_{c_2 \hat{\mathbf{p}} + (1-c_2)\tilde{\mathbf{p}}} + \frac{\partial^2}{\partial \mathbf{p}^2} \mathbb{E}[\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n)]|_{c_1 \hat{\mathbf{p}} + (1-c_1)\tilde{\mathbf{p}}} \right) (\hat{\mathbf{p}} - \tilde{\mathbf{p}}) \right]}_{\psi_n}. \end{aligned} \tag{C.3}$$

In the following lemma we establish the convergence of ψ_n to a chi-square distribution with $s + 1$ degrees of freedom.

Lemma 5. *For a sequence (k_n) such that $k_n \rightarrow \infty$ and $k_n/n \rightarrow 0$, the following weak convergence holds:*

$$\psi_n \xrightarrow{d} \chi^2(s + 1), \quad n \rightarrow \infty.$$

Proof of Lemma 5. Starting from the second order derivative of the log-likelihood $\log L_{\mathbf{M}_k}$,

$$\frac{\partial^2}{\partial \mathbf{p}^2} \log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n) = -\text{Diag} \left(\frac{T_{n,1}}{p_1^2}, \dots, \frac{T_{n,s}}{p_s^2} \right) - \frac{\sum_{i=s+1}^r T_{n,i}}{(1 - \sum_{i=1}^s p_i)^2} \mathbf{1} \cdot \mathbf{1}^\top,$$

we obtain

$$\begin{aligned} \psi_n &= \sum_{j=1}^s \left(\frac{T_{n,j}}{k} - \tilde{p}_j \right)^2 \left(\frac{T_{n,j}}{2(c_2 T_{n,j}/k + (1-c_2)\tilde{p}_j)^2} + \frac{p_{n,j}}{2(c_1 T_{n,j}/k + (1-c_1)\tilde{p}_j)^2} \right) \\ &+ \left(\sum_{j=1}^s \left(\frac{T_{n,j}}{k} - \tilde{p}_j \right) \right)^2 \left(\frac{\sum_{j=s+1}^r T_{n,j}}{2(c_2 \sum_{j=1}^s T_{n,j}/k + (1-c_2) \sum_{j=1}^s \tilde{p}_j)^2} + \frac{\sum_{j=s+1}^r p_{n,j}}{2(c_1 \sum_{j=1}^s T_{n,j}/k + (1-c_1) \sum_{j=1}^s \tilde{p}_j)^2} \right). \end{aligned}$$

We rewrite ψ_n as

$$\psi_n = k \sum_{j=1}^s \frac{(T_{n,j}/k - \tilde{p}_j)^2}{p_{n,j}} A_{n,j} + \left(\frac{\sum_{j=s+1}^r (T_{n,j}/k - \tilde{p}_j)}{\sum_{j=s+1}^r p_{n,j}} \right)^2 B_n \quad (\text{C.4})$$

where

$$A_{n,j} := \frac{p_{n,j} T_{n,j}/k}{2(c_2 T_{n,j}/k + (1-c_2)\tilde{p}_j)^2} + \frac{p_{n,j}^2}{2(c_1 T_{n,j}/k + (1-c_1)\tilde{p}_j)^2},$$

and

$$B_n := \frac{\sum_{j=s+1}^r p_{n,j} \sum_{j=s+1}^r T_{n,j}}{2(c_2 \sum_{j=s+1}^r T_{n,j}/k + (1-c_2) \sum_{j=s+1}^r \tilde{p}_j)^2} + \frac{(\sum_{j=s+1}^r p_{n,j})^2}{2(c_1 \sum_{j=s+1}^r T_{n,j}/k + (1-c_1) \sum_{j=s+1}^r \tilde{p}_j)^2}.$$

Since $\sum_{i=1}^{\hat{s}_n} p_{n,i} \rightarrow 1$ as $n \rightarrow \infty$, we obtain $\tilde{p}_j = p_{n,j}/\sum_{i=1}^{\hat{s}_n} p_{n,i} \sim p_{n,j}$. Therefore, by Slutsky's lemma and convergence (4.9) it suffices to prove that $A_{n,j}$ and B_n converges to 1 in probability.

Theorem 1 entails that for any $j \in \mathcal{S}_k^\infty$ we have the convergence in probability

$$\frac{T_{n,j}}{k p_{n,j}} = \frac{1}{\sqrt{k p_{n,j}}} \sqrt{k} \frac{T_{n,j}/k - p_{n,j}}{\sqrt{p_{n,j}}} + 1 \rightarrow 1, \quad n \rightarrow \infty. \quad (\text{C.5})$$

Therefore, for any $c \in (0, 1)$ we have

$$\frac{c T_{n,j}/k + (1-c)\tilde{p}_j}{p_{n,j}} = c \frac{T_{n,j}}{k p_{n,j}} + \frac{1-c}{\sum_{j=1}^{\hat{s}_n} p_{n,j}} \rightarrow c + (1-c) = 1, \quad n \rightarrow \infty.$$

This proves that for all $j = 1, \dots, s$ we have the convergence in probability

$$\frac{p_{n,j} T_{n,j}/k}{2(c_2 T_{n,j}/k + (1-c_2)\tilde{p}_j)^2} + \frac{p_{n,j}^2}{2(c_1 T_{n,j}/k + (1-c_1)\tilde{p}_j)^2} \rightarrow 1, \quad n \rightarrow \infty, \quad j = 1, \dots, s,$$

The same arguments apply to B_n since we have $\sum_{j=s+1}^r T_{n,j} = T_n(\cup_{j=s+1}^r \beta_j)$ and $\sum_{j=s+1}^r p_{n,j} = p_n(\cup_{j=s+1}^r \beta_j)$.

This proves that the convergence of ψ_n to a chi-square distribution with $s+1$ degrees of freedom. \square

Lemma 5 entail the following approximation of (C.3) when n is large:

$$\mathbb{E}[KL(\mathbf{P}_k || \mathbf{M}_k) |_{\mathbf{p}=\hat{\mathbf{p}}}] \approx \mathbb{E}[\log L_{\mathbf{P}_k}(\mathbf{T}_n)] - \mathbb{E}[\log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n)] + \mathbb{E}[\chi^2(s+1)].$$

The first expectation is constant with respect to the parameter \mathbf{p} . Therefore the variations of the divergence $KL(\mathbf{P}_k || \mathbf{M}_k) |_{\mathbf{p}=\hat{\mathbf{p}}}$ can be estimated by the variations of $\log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n) + (s+1)$.

C.2 Threshold selection

We start from Equation (3.8) which we recall here:

$$\mathbb{E}[KL(\mathbf{P}'_n \parallel \mathbf{M}'_n) | \hat{\mathbf{p}}] = \mathbb{E}[\log L_{\mathbf{P}_n}(\mathbf{T}'_n)] + \mathbb{E} \left[\mathbb{E} \left[-\log L_{\mathbf{M}_{n-T'_{n,2^d}}}(\mathbf{p}; \mathbf{T}'_{n,\{1,\dots,2^d-1\}}) \mid T'_{n,2^d} \right] \Big| \hat{\mathbf{p}} \right] - \mathbb{E} \left[\mathbb{E} \left[\phi(n, q', T'_{n,2^d}) \right] \Big| \hat{\mathbf{p}} \right]. \quad (\text{C.6})$$

The first term on the right-hand side is a constant. For the second term we write

$$\log L_{\mathbf{M}_{n-T'_{n,2^d}}}(\mathbf{p}; \mathbf{T}'_{n,\{1,\dots,2^d-1\}}) = \log((n-T'_{n,2^d})!) - \sum_{i=1}^{2^d-1} \log(T'_{n,i}!) + \sum_{i=1}^{s'} T'_{n,i} \log(p'_i) + \left(\sum_{i=s'+1}^{\hat{s}_n} T'_{n,i} \right) \log(p')$$

The two first terms satisfy

$$\mathbb{E} \left[-\log((n-T'_{n,2^d})!) + \sum_{i=1}^{2^d-1} \log(T'_{n,i}!) \mid T'_{n,2^d} \right] = -\mathbb{E} \left[\log((n-T'_{n,2^d})!) - \sum_{j=1}^{2^d-1} \log(T'_{n,j}!) \right].$$

The idea is then to condition the last two terms of the log-likelihood with respect to $T'_{n,2^d}$ in order to apply the results of the bias selection. Since \mathbf{T}'_n follows a multinomial distribution with parameter \mathbf{p}'_n given in (3.6), the conditional distribution of $(T'_{n,1}, \dots, T'_{n,2^d-1})^\top \mid T'_{n,2^d}$ is multinomial with parameters $n - T'_{n,2^d}$ and \mathbf{p}_n . This entails in particular that $\mathbb{E}[T'_{n,j} \mid T'_{n,2^d}] = (n - T'_{n,2^d})p_{n,j}$. Hence, we obtain

$$\begin{aligned} & \mathbb{E} \left[-\log L_{\mathbf{M}_{n-T'_{n,2^d}}}(\mathbf{p}; \mathbf{T}'_{n,\{1,\dots,2^d-1\}}) + \log((n-T'_{n,2^d})!) - \sum_{j=1}^{2^d-1} \log(T'_{n,j}!) \mid T'_{n,2^d} \right] \\ &= \sum_{j=1}^s \mathbb{E}[T'_{n,j} \mid T'_{n,2^d}] \log(p_j) + \log(p) \sum_{j=s+1}^{2^d-1} \mathbb{E}[T'_{n,j} \mid T'_{n,2^d}] \\ &= (n - T'_{n,2^d}) \left(\sum_{j=1}^s p_{n,j} \log(p_j) + \log(p) \sum_{j=s+1}^{2^d-1} p_{n,j} \right) \\ &= \frac{n - T'_{n,2^d}}{k} \left(k \sum_{j=1}^s p_{n,j} \log(p_j) + k \log(p) \sum_{j=s+1}^{2^d-1} p_{n,j} \right) \\ &= \frac{n - T'_{n,2^d}}{k} \left(\mathbb{E}[-\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n(k))] + \log(k!) - \mathbb{E} \left[\sum_{j=1}^{2^d-1} \log(T_{n,j}(k)!) \right] \right), \end{aligned}$$

where $\mathbf{T}_n(k) \in \mathbb{R}^{2^d-1}$ corresponds to the vector of Section C.1 and for which we emphasize the dependence in k . The last relation results from (3.4) since $\mathbf{T}_n(k)$ is given $T'_{n,2^d} = n - k$ implicitly in the model \mathbf{M}_k . Then, we evaluate the previous quantity in $\hat{\mathbf{p}}$ and take its expectation:

$$\begin{aligned} & \mathbb{E} \left[\mathbb{E} \left[-\log L_{\mathbf{M}_{n-T'_{n,2^d}}}(\mathbf{p}; \mathbf{T}'_{n,\{1,\dots,2^d-1\}}) + \log((n-T'_{n,2^d})!) - \sum_{j=1}^{2^d-1} \log(T'_{n,j}!) \mid T'_{n,2^d} \right] \Big| \hat{\mathbf{p}} \right] \\ &= \frac{n(1 - q_n)}{k} \left(\mathbb{E} \left[\mathbb{E} \left[-\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n(k)) \right] \Big| \hat{\mathbf{p}} \right] + \log(k!) - \mathbb{E} \left[\sum_{j=1}^{2^d-1} \log(T_{n,j}(k)!) \right] \right). \end{aligned}$$

For the third term in Equation (3.8), we have

$$\begin{aligned}\mathbb{E}[\phi(n, q', T'_{n,2^d})]_{|\hat{\mathbf{p}}'} &= \mathbb{E}\left[\log\left(\frac{n!}{(n - T'_{n,2^d})!}\right) - \log(T'_{n,2^d}!)\right] + \mathbb{E}\left[(n - T'_{n,2^d}) \log(q') + T'_{n,2^d} \log(1 - q')\right]_{|\hat{q}'} \\ &= \mathbb{E}\left[\log\left(\frac{n!}{(n - T'_{n,2^d})!}\right) - \log(T'_{n,2^d}!)\right] + nq_n \log(k/n) + n(1 - q_n) \log(1 - k/n),\end{aligned}$$

as we use the estimator $\hat{q}' = k/n$.

We consider a large n and we use the results of the previous section. Then, we estimate the Kullback-Leibler divergence $KL(\mathbf{P}'_n \|\mathbf{M}'_n)$ in (3.7) by the unbiased estimator $\mathbb{E}[KL(\mathbf{P}'_n \|\mathbf{M}')_{|\hat{\mathbf{p}}'}]$ in (3.8) which can be approximated by the following quantity

$$\mathbb{E}[\log L_{\mathbf{P}'_n}(\mathbf{T}'_n)] + \frac{n(1 - q_n)}{k} \left(\mathbb{E}[-\log L_{\mathbf{M}_k}(\hat{\mathbf{p}}; \mathbf{T}_n(k))] + (s + 1) \right) + R_{n,k}, \quad (\text{C.7})$$

where $R_{n,k}$ is defined as

$$\begin{aligned}R_{n,k} &= -\mathbb{E}\left[\log\left(\frac{n!}{(n - T'_{n,2^d})!}\right) - \log(T'_{n,2^d}!)\right] - nq_n \log(k/n) - n(1 - q_n) \log(1 - k/n) \\ &\quad + \frac{n(1 - q_n)}{k} \left(\log(k!) - \mathbb{E}\left[\sum_{j=1}^{2^d-1} \log(T_{n,j}(k)!) \right] \right) + \mathbb{E}\left[-\log((n - T'_{n,2^d})!) + \sum_{j=1}^{2^d-1} \log(T'_{n,j}!)\right].\end{aligned}$$

After withdrawing the terms which are constant with respect to k , it remains

$$-nq_n \log(k/n) - n(1 - q_n) \log(1 - k/n) + \frac{n(1 - q_n)}{k} \left(\log(k!) - \mathbb{E}\left[\sum_{j=1}^{2^d-1} \log(T_{n,j}(k)!) \right] \right). \quad (\text{C.8})$$

Regarding the last term, Stirling's approximation $\log(m!) = m \log(m) - m + o(m)$ yields

$$\begin{aligned}\log(k!) - \mathbb{E}\left[\sum_{j=1}^{2^d-1} \log(T_{n,j}(k)!) \right] &\approx k \log(k) - k - \sum_{j=1}^{2^d-1} \left(kp_{n,j} \log(kp_{n,j}) - kp_{n,j} \right) \\ &\approx k \log(k) - \sum_{j=1}^{2^d-1} kp_{n,j} \log(k) - \sum_{j=1}^{2^d-1} kp_{n,j} \log(p_{n,j}) \\ &\approx -k \sum_{j=1}^{2^d-1} p_{n,j} \log(p_{n,j}),\end{aligned}$$

where we use that the $p_{n,j}$ add up to 1. This implies that the last term in (C.8) is approximately constant. Regarding the first one, Assumption 1 implies that

$$|-nq_n \log(k/n)| \leq nq_n \log(n) \rightarrow 0, \quad n \rightarrow \infty.$$

The only term remaining which has not been neglected in $R_{n,k}$ is then $-n(1 - q_n) \log(1 - k/n)$.

D Proof of Proposition 3

We recall the expression of the log-likelihood $\log L_{\mathbf{M}_k}$, for $\mathbf{p} \in \mathbb{R}^{s_2}$:

$$\log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n) = \log(k!) - \sum_{i=1}^{2^d-1} \log(T_{n,i}!) + \sum_{i=1}^{s_2} T_{n,i} \log(p_i) + \left(\sum_{i=s_2+1}^r T_{n,i} \right) \log\left(\frac{1 - \sum_{j=1}^{s_2} p_j}{r - s_2}\right). \quad (\text{D.1})$$

The gradient and the Hessian matrix of this log-likelihood are given by

$$\frac{\partial}{\partial \mathbf{p}} \log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n) = \left(\frac{T_{n,j}}{p_j} - \frac{\sum_{i=s_2+1}^r T_{n,i}}{1 - \sum_{i=1}^{s_2} p_i} \right)_{1 \leq j \leq s_2}, \quad (\text{D.2})$$

and

$$\frac{\partial^2}{\partial \mathbf{p}^2} \log L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n) = -\text{Diag} \left(\frac{T_{n,1}}{p_1^2}, \dots, \frac{T_{n,s}}{p_{s_2}^2} \right) - \frac{\sum_{i=s_2+1}^r T_{n,i}}{(1 - \sum_{i=1}^{s_2} p_i)^2} \mathbf{1} \cdot \mathbf{1}^\top, \quad (\text{D.3})$$

where $\mathbf{1} \in \mathbb{R}^s$. In particular the information matrix satisfies

$$I_{\mathbf{M}_k}(\mathbf{p}) := -\mathbb{E} \left[\frac{\partial^2}{\partial \mathbf{p}^2} L_{\mathbf{M}_k}(\mathbf{p}; \mathbf{T}_n) \right] = k_n \text{Diag} \left(\frac{p_{n,1}}{p_1^2}, \dots, \frac{p_{n,s_2}}{p_{s_2}^2} \right) + k_n \frac{\sum_{i=s_2+1}^r p_{n,i}}{(1 - \sum_{i=1}^{s_2} p_i)^2} \mathbf{1} \cdot \mathbf{1}^\top.$$

Under H_0 the log-likelihood in (D.1) varies like

$$\sum_{i=1}^{s_1} T_{n,i} \log(p_i) + \sum_{i=s_1+1}^{s_2} T_{n,i} \log(\rho_n) + \left(\sum_{i=s_2+1}^r T_{n,i} \right) \log \left(\frac{1 - \sum_{j=1}^{s_1} p_j - (s_2 - s_1)\rho_n}{r - s} \right).$$

The parameters $\hat{p}_1, \dots, \hat{p}_{s_1}, \hat{\rho}_n$ which maximize this quantity are

$$\hat{p}_j = \frac{T_{n,j}}{k_n}, \quad j = 1, \dots, s_1, \quad \text{and} \quad \hat{\rho}_n = \frac{\sum_{i=s_1+1}^{s_2} T_{n,i}}{k_n(s_2 - s_1)}.$$

Thus we obtain that the gradient (D.2) and the information matrix (D.3) evaluated in $\hat{\mathbf{p}}_{H_0} := (\hat{p}_1, \dots, \hat{p}_{s_1}, \hat{\rho}_n, \dots, \hat{\rho}_n)^\top \in \mathbb{R}^{s_2}$ correspond to

$$\frac{\partial}{\partial \mathbf{p}} \log L_{\mathbf{M}_k}(\hat{\mathbf{p}}_{H_0}; \mathbf{T}_n) = k_n \left(0, \dots, 0, \frac{(s_2 - s_1)T_{n,s_1+1}}{\sum_{i=s_1+1}^{s_2} T_{n,i}} - 1, \dots, \frac{(s_2 - s_1)T_{n,s_2}}{\sum_{i=s_1+1}^{s_2} T_{n,i}} - 1 \right)^\top,$$

and

$$I_{\mathbf{M}_k}(\hat{\mathbf{p}}_{H_0}) = \text{Diag}(\mathbf{W}_n) + k_n^3 \frac{\sum_{i=s_2+1}^r p_{n,i}}{(\sum_{i=s_2+1}^r T_{n,i})^2} \mathbf{1} \cdot \mathbf{1}^\top,$$

where

$$\mathbf{W}_n = \left(\frac{k_n^3 p_{n,1}}{T_{n,1}^2}, \dots, \frac{k_n^3 p_{n,s_1}}{T_{n,s_1}^2}, \frac{k_n^3 \rho_n (s_2 - s_1)^2}{(\sum_{i=s_1+1}^{s_2} T_{n,i})^2}, \dots, \frac{k_n^3 \rho_n (s_2 - s_1)^2}{(\sum_{i=s_1+1}^{s_2} T_{n,i})^2} \right)^\top \in \mathbb{R}^{s_2}.$$

For the information matrix we used the fact that $p_{n,j} = \rho_n$ for all $j = s_1 + 1, \dots, s_2$ under H_0 . The inverse of this matrix is given by the Sherman–Morrison formula,

$$I_{\mathbf{M}_k}(\hat{\mathbf{p}}_{H_0})^{-1} = \text{Diag}(\mathbf{W}_n^{-1}) - k_n^3 \gamma_n^{-1} \mathbf{W}_n^{-1} (\mathbf{W}_n^{-1})^\top,$$

where \mathbf{W}_n^{-1} is meant componentwise and where

$$\gamma_n := \sum_{j=1}^{s_1} \frac{T_{n,j}^2}{p_{n,j}} + \frac{(\sum_{j=s_1+1}^{s_2} T_{n,j})^2}{\rho_n (s_2 - s_1)} + \frac{(\sum_{j=s_2+1}^r T_{n,j})^2}{\sum_{j=s_2+1}^r p_{n,j}}.$$

Similarly to the score test we compute the quantity

$$\begin{aligned} & \left(\frac{\partial}{\partial \mathbf{p}} \log L_{\mathbf{M}_k(s_2)}(\hat{\mathbf{p}}_{H_0}; \mathbf{T}_n) \right)^\top I_{\mathbf{M}_k(s_2)}(\hat{\mathbf{p}}_{H_0})^{-1} \left(\frac{\partial}{\partial \mathbf{p}} \log L_{\mathbf{M}_k(s_2)}(\hat{\mathbf{p}}_{H_0}; \mathbf{T}_n) \right) \\ &= \sum_{j=s_1+1}^{s_2} k_n^2 \left(\frac{(s_2 - s_1)T_{n,j}}{\sum_{i=s_1+1}^{s_2} T_{n,i}} - 1 \right)^2 \frac{(\sum_{j=s_1+1}^{s_2} T_{n,j})^2}{k_n^3 \rho_n (s_2 - s_1)^2} - k_n^3 \gamma_n^{-1} \left[\sum_{j=s_1+1}^{s_2} k_n \left(\frac{(s_2 - s_1)T_{n,j}}{\sum_{i=s_1+1}^{s_2} T_{n,i}} - 1 \right) \frac{(\sum_{j=s_1+1}^{s_2} T_{n,j})^2}{k_n^3 \rho_n (s_2 - s_1)^2} \right]^2 \\ &= k_n \sum_{j=s_1+1}^{s_2} \left(\frac{T_{n,j}}{k_n} - \hat{\rho}_n \right)^2 \frac{1}{\rho_n}. \end{aligned} \quad (\text{D.4})$$

The second term vanishes since

$$\sum_{j=s_1+1}^{s_2} \left(\frac{(s_2 - s_1)T_{n,j}}{\sum_{i=s_1+1}^{s_2} T_{n,i}} - 1 \right) = \frac{(s_2 - s_1) \sum_{j=s_1+1}^{s_2} T_{n,j}}{\sum_{i=s_1+1}^{s_2} T_{n,i}} - (s_2 - s_1) = 0.$$

For the remaining term, we write $(T_{n,j}/k_n - \rho_n)^2 = (T_{n,j}/k_n - \hat{\rho}_n + \hat{\rho}_n - \rho_n)^2$ and develop this square. We obtain

$$\sum_{j=s_1+1}^{s_2} \frac{(T_{n,j}/k_n - \hat{\rho}_n)^2}{\rho_n} = \sum_{j=s_1+1}^{s_2} \frac{(T_{n,j}/k_n - \rho_n)^2}{\rho_n} - (s_2 - s_1) \frac{(\hat{\rho}_n - \rho_n)^2}{\rho_n}.$$

We consider the vector $\bar{\mathbf{p}}_n \in \mathbb{R}^{s_2-s_1}$ whose components are $\sqrt{k_n}(T_{n,j}/k_n - \rho_n)/\sqrt{\rho_n}$ for $j = s_1+1, \dots, s_2$. Then the quadratic form (D.4) can then be written as $\bar{\mathbf{p}}_n^\top (\text{Diag}(\mathbf{1}) - (s_2 - s_1)^{-1} \mathbf{1} \cdot \mathbf{1}^\top) \bar{\mathbf{p}}_n$. We notice that the matrix $\text{Diag}(\mathbf{1}) - (s_2 - s_1)^{-1} \mathbf{1} \cdot \mathbf{1}^\top$ is an orthogonal projection on the subspace $\{\mathbf{x} \in \mathbb{R}^{s_2-s_1} : x_1 + \dots + x_{s_2-s_1} = 0\}$ so that

$$k_n \sum_{j=s_1+1}^{s_2} \frac{(T_{n,j}/k_n - \hat{\rho}_n)^2}{\rho_n} = |(\text{Diag}(\mathbf{1}) - (s_2 - s_1)^{-1} \mathbf{1} \cdot \mathbf{1}^\top) \bar{\mathbf{p}}_n|_2.$$

Theorem 1 ensures that this quantity converges to a chi-square distributions with $s_2 - s_1 - 1$ degrees of freedom. Finally, the convergence in probability $\rho_n/\hat{\rho}_n \rightarrow 1$ when $n \rightarrow \infty$ and Slutsky's lemma ensure that

$$\mathcal{T}_n = k_n \sum_{j=s_1+1}^{s_2} \frac{(T_{n,j}/k_n - \hat{\rho}_n)^2}{\hat{\rho}_n} \xrightarrow{d} \chi^2(s_2 - s_1 - 1), \quad n \rightarrow \infty,$$

which concludes the proof.

References

- Abdous, B. and Ghoudi, K. (2005). Non-parametric estimators of multivariate extreme dependence functions. *Nonparametric Statistics*, 17(8):915–935.
- Akaike, H. (1973). Information theory and an extension of the maximum likelihood principle. In *2nd International Symposium on Information Theory*, pages 267–281, Budapest. Akademia Kiado.
- Beirlant, J., Goegebeur, Y., Segers, J., Teugels, J., De Waal, D., and Ferro, C. (2006). *Statistics of Extremes: Theory and Applications*. Wiley.
- Caeiro, F. and Gomes, M. I. (2015). Threshold selection in extreme value analysis. *Extreme Value Modeling and Risk Analysis: Methods and Applications*, pages 69–86.
- Chautru, E. (2015). Dimension reduction in multivariate extreme value analysis. *Electronic Journal of Statistics*, 9(1):383–418.
- Chiapino, M. and Sabourin, A. (2016). Feature clustering for extreme events analysis, with application to extreme stream-flow data. In *International Workshop on New Frontiers in Mining Complex Patterns*, pages 132–147. Springer.
- Chiapino, M., Sabourin, A., and Segers, J. (2019). Identifying groups of variables with the potential of being large simultaneously. *Extremes*, 22(2):193–222.

- Coles, S. G. and Tawn, J. A. (1991). Modelling extreme multivariate events. *Journal of the Royal Statistical Society: Series B (Methodological)*, 53(2):377–392.
- Condat, L. (2016). Fast projection onto the simplex and the ℓ_1 ball. *Mathematical Programming*, 158(1-2):575–585.
- Cooley, D. and Thibaud, E. (2019). Decompositions of dependence for high-dimensional extremes. *Biometrika*, 106(3):587–604.
- de Haan, L. and Ferreira, A. (2006). *Extreme Value Theory: An Introduction*. Springer Series in Operations Research and Financial Engineering. Springer, New York.
- Duchi, J., Shalev-Shwartz, S., Singer, Y., and Chandra, T. (2008). Efficient projections onto the ℓ_1 -ball for learning in high dimensions. In *Proceedings of the 25th international conference on Machine learning*, pages 272–279. ACM.
- Einmahl, J., de Haan, L., and Huang, X. (1993). Estimating a multidimensional extreme-value distribution. *Journal of Multivariate Analysis*, 47(1):35–47.
- Einmahl, J., de Haan, L., and Sinha, A. (1997). Estimating the spectral measure of an extreme value distribution. *Stochastic Processes and their Applications*, 70(2):143–171.
- 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.
- Goix, N., Sabourin, A., and Cléménçon, S. (2017). Sparse representation of multivariate extremes with applications to anomaly detection. *Journal of Multivariate Analysis*, 161:12–31.
- Haslett, J. and Raftery, A. E. (1989). Space-time modelling with long-memory dependence: Assessing Ireland’s wind power resource. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 38(1):1–21.
- Heffernan, J. E. and Tawn, J. A. (2004). A conditional approach for multivariate extreme values (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3):497–546.
- Hill, B. M. (1975). A simple general approach to inference about the tail of a distribution. *The Annals of Statistics*, 3(5):1163–1174.
- Hille, E. (1964). *Analysis*, volume 1. Blaisdell, New-York.
- Hult, H. and Lindskog, F. (2006). Regular variation for measures on metric spaces. *Publications de l’Institut Mathématique*, 80(94):121–140.
- Kiriliouk, A., Rootzén, H., Segers, J., and Wadsworth, J. L. (2019). Peaks over thresholds modeling with multivariate generalized Pareto distributions. *Technometrics*, 61(1):123–135.
- Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86.
- Kyrillidis, A., Becker, S., Cevher, V., and Koch, C. (2013). Sparse projections onto the simplex. In *International Conference on Machine Learning*, volume 28, pages 235–243.
- Ledford, A. W. and Tawn, J. A. (1996). Statistics for near independence in multivariate extreme values. *Biometrika*, 83(1):169–187.

- Lee, J., Fan, Y., and Sisson, S. A. (2015). Bayesian threshold selection for extremal models using measures of surprise. *Computational Statistics & Data Analysis*, 85:84–99.
- Massart, P. (2007). *Concentration inequalities and model selection*. Springer, Berlin.
- Meyer, N. and Wintenberger, O. (2021). Sparse regular variation. *to appear in Advances in Applied Probability*, 4(53).
- Raval, A., Owen, W., Hume, N., and Stephen, M. (3 Sept. 2020). Biggest banks sustain coal financing despite defunding drive. *Financial Times*.
- Resnick, S. I. (1987). *Extreme Values, Regular Variation and Point Processes*. Springer, New-York.
- Resnick, S. I. (2007). *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*. Springer, New-York.
- Rootzén, H., Tajvidi, N., et al. (2006). Multivariate generalized Pareto distributions. *Bernoulli*, 12(5):917–930.
- Segers, J. (2012). Max-stable models for multivariate extremes. *Revstat Statistical Journal*, 10:61–82.
- Simpson, E., Wadsworth, J., and Tawn, J. (2019). Determining the dependence structure of multivariate extremes. *arXiv:1809.01606*.
- Stărică, C. (1999). Multivariate extremes for models with constant conditional correlations. *Journal of Empirical Finance*, 6(5):515–553.
- Wan, P. and Davis, R. A. (2019). Threshold selection for multivariate heavy-tailed data. *Extremes*, 22(1):131–166.