

Robust bounds in multivariate extremes



Sebastian Engelke and Jevgenijs Ivanovs

Robust bounds in multivariate extremes

Sebastian Engelke¹ and Jevgenijs Ivanovs²

¹Ecole Polytechnique Fédérale de Lausanne, email: sebastian.engelke@epfl.ch

²Department of Mathematics, Aarhus University, email: jevgenijs.ivanovs@math.au.dk

Abstract

Extreme value theory provides an asymptotically justified framework for estimation of exceedance probabilities in regions where few or no observations are available. For multivariate tail estimation, the strength of extremal dependence is crucial and it is typically modeled by a parametric family of spectral distributions. In this work we provide asymptotic bounds on exceedance probabilities that are robust against misspecification of the extremal dependence model. They arise from optimizing the statistic of interest over all dependence models within some neighborhood of the reference model. A certain relaxation of these bounds yields surprisingly simple and explicit expressions, which we propose to use in applications. We show the effectiveness of the robust approach compared to classical confidence bounds when the model is misspecified. The results are further applied to quantify the effect of model uncertainty on the Value-at-Risk of a financial portfolio.

Keywords: Extremal dependence, Pickands' function, model misspecification, stress test, robust bounds, convex optimization

1 Introduction

In parametric statistics there are several sorts of uncertainties that arise in the estimation of an unknown quantity of interest. The estimation uncertainty, for instance, refers to the error made by inferring the model parameters from only finitely many data points. Bootstrapping or results on asymptotic normality are typically applied to quantify this error and to derive confidence intervals. On the other hand, the parametric family used as a model for the data is a finitely dimensional subset of all distributions and is thus only an approximation of the true data generating distribution. The uncertainty due to this misspecification is usually called model uncertainty, and it is more difficult to quantify than the estimation uncertainty within a parametric model class. A popular way to provide confidence bounds, that are robust against wrong model assumptions, is to find the smallest and largest values of the statistic of interest with respect to all probability measures in some neighborhood of the estimated parametric distribution assuming that it contains the true data generating distribution. Moreover, one may view such a search for the worst case

as a systematic stress test within a set of plausible scenarios Breuer and Csiszár (2013b).

For a random vector $(X, \mathbf{Y}^\top) = (X, Y_1, \dots, Y_{d-1})$ with $d \geq 2$, in this paper we consider the optimization problem

$$V_\mu(\delta) = \sup_{\mathbb{P}'} \{\mathbb{E}' X : D_\mu(\mathbb{P}', \mathbb{P}) \leq \delta, \mathbb{E}' \mathbf{Y} = \mathbb{E} \mathbf{Y}\}, \quad \delta > 0, \quad (1.1)$$

where the supremum is taken over all probability measures in the δ -neighborhood of the reference model \mathbb{P} under the constraint that the expectation of \mathbf{Y} is preserved. Here and in the sequel \mathbb{E}' denotes the expectation under the model \mathbb{P}' , and all the measures are defined on a common measurable space (Ω, \mathcal{F}) . The proximity $D_\mu(\mathbb{P}', \mathbb{P})$ will be measured in terms of the L_μ^2 -distance between the densities of \mathbb{P}' and \mathbb{P} with respect to some dominating probability measure μ , which provides additional flexibility in selection of the neighborhood; it will be shown that the choice $\mu = \mathbb{P}$ essentially results in Rényi divergence of order 2. The random variable X is the statistic of interest and the constraint on the expectation of \mathbf{Y} allows to incorporate necessary model restrictions. They arise naturally in the application of the results to estimation of multivariate tail probabilities.

Importantly, the optimizing μ -density has an appealing form yielding the surprisingly simple, explicit expression

$$\mathbb{E} X + \sqrt{\delta \frac{\det\{\Sigma_\mu(X, \mathbf{Y})\}}{\det\{\Sigma_\mu(\mathbf{Y})\}}} \quad (1.2)$$

for the optimal value $V_\mu(\delta)$ when $\delta \in [0, \delta^*]$ is in a certain range, and otherwise this expression provides an upper bound on $V_\mu(\delta)$, where $\Sigma_\mu(\cdot)$ denotes the respective μ -covariance matrix. In this paper, we advocate using this simple square-root bound, and its analogue for the respective minimization problem, as robust bounds for $\mathbb{E} X$ under moment constraints; see Theorem 2. Interestingly, the above fraction of the determinants is a well-known expression in stochastic simulation theory where it arises as the minimal variance of $X + \mathbf{c}^\top \mathbf{Y}$, for arbitrary $\mathbf{c} \in \mathbb{R}^{d-1}$ (Asmussen and Glynn, 2007, Sec. V.2).

The general optimization problem (1.1) might be interesting in many different situations, see e.g. Dey and Juneja (2012); Glasserman and Xu (2014); Breuer and Csiszár (2013b) for applications of the robust approach to various problems in economics, risk and finance. Let us also note that a problem similar to (1.1) appears as the dual representation of a coherent risk measure Ahmadi-Javid (2012); Artzner et al. (1999). In this work we concentrate on the application to the risk of rare events and the estimation of their small tail probabilities, a field that has attracted strong attention in the last decade. Extreme value theory provides the theoretical foundation for statistical extrapolation into tail regions with few or no data; see Embrechts et al. (1997); de Haan and Ferreira (2006); Resnick (2008) for more details.

The univariate theory is well understood and is concerned with the quantification of tail probabilities $\mathbb{P}(Z > z)$ of a random variable Z , where $z > 0$ is a threshold close to the upper end point of its distribution function F . There are standard procedures to build confidence intervals for estimators of $\mathbb{P}(Z > z)$, but bounds that are robust against violation of the assumptions of the extremal types theorem

have only recently been studied in Blanchet and Murthy (2016). The authors of this paper solve the optimization problem

$$\overline{F}_\delta(z) = \sup_{\mathbb{P}'} \{ \mathbb{P}'(Z > z) : \widehat{D}(\mathbb{P}', \mathbb{P}) \leq \delta \},$$

where \overline{F}_δ is the worst case survival function over all probability measures \mathbb{P}' in some divergence neighborhood with radius $\delta > 0$ around the reference model \mathbb{P} . Here \widehat{D} is either the Kullback–Leibler divergence or the Rényi divergence of an arbitrary order; see also Section A.3. It is shown in Blanchet and Murthy (2016) that the worst case tail \overline{F}_δ is considerably heavier than the one of the reference distribution F .

For a d -dimensional random vector $\mathbf{Z} = (Z_1, \dots, Z_d)^\top$, multivariate extreme value theory studies probabilities $\mathbb{P}(\mathbf{Z} \in tB)$, where for $B \subset [0, \infty]^d$ bounded away from the origin and large $t > 0$ the dilated set tB is called a tail region. As in the univariate case, the idea is to extrapolate from regions with more data into the tails, but in the multivariate case the dependence between components Z_i at high quantiles is crucial. The mathematical concept of regular variation is needed in order to perform this extrapolation. Assuming that \mathbf{Z} is standardized to have unit Pareto marginal tails, multivariate extreme value theory justifies, in particular, the following approximation for any $z_i > 0$ and large $t > 0$:

$$\mathbb{P}(\exists i : Z_i > tz_i) \approx t^{-1} d \mathbb{E} \left(\max_{i=1}^d \frac{Y_i}{z_i} \right), \quad (1.3)$$

where $\mathbf{Y} = (Y_1, \dots, Y_d)$ takes values in the standard simplex and satisfies certain moment constraints, i.e.,

$$\mathbf{Y} \in \mathbb{S}^{d-1} = \{ \mathbf{y} \in [0, \infty]^d : \sum_{i=1}^d y_i = 1 \} \quad \text{and} \quad \forall i : \mathbb{E} Y_i = 1/d.$$

The distribution of \mathbf{Y} is called a *spectral distribution* and it encodes extremal dependence in the model. Many parametric models have been proposed for the spectral distribution (e.g., Hüsler and Reiss, 1989; Tawn, 1990; Boldi and Davison, 2007; Cooley et al., 2010). For a non-parametric approach to estimation of the spectral distribution we refer the reader to Einmahl and Segers (2009), where an optimization problem is used to enforce the moment constraint.

A natural problem is to find bounds for the asymptotic expression of the tail probability in (1.3) with fixed $\mathbf{z} = (z_1, \dots, z_d)^\top$ that are robust against model misspecification of the spectral distribution, i.e., the distribution of \mathbf{Y} . For the upper bound we are thus interested in the maximization problem

$$\sup_{\mathbb{P}'} \left\{ \mathbb{E}' \left(\max_{i=1}^d \frac{Y_i}{z_i} \right) : D_\mu(\mathbb{P}', \mathbb{P}) \leq \delta, \mathbb{E}' Y_i = 1/d \text{ for all } i \right\} \quad (1.4)$$

which is clearly a special case of (1.1) with $X = X(\mathbf{z}) = \max_{i=1}^d Y_i/z_i$. Importantly, we assume here that the dominating measure μ is supported by $\{ \mathbf{Y} \in \mathbb{S}^{d-1} \}$ and hence $\mathbf{Y} \in \mathbb{S}^{d-1}$ holds also \mathbb{P}' -a.s. In particular, $Y_d = 1 - \sum_{i=1}^{d-1} Y_i$ and so there are essentially $d - 1$ moment constraints which ensure that \mathbf{Y} has a valid spectral distribution also under the measure \mathbb{P}' .

Similarly, a lower bound can be defined as the optimal value of the corresponding minimization problem with \sup replaced by \inf in (1.4). The respective optimal values $\beta^*(\mathbf{z})$ and $\beta_*(\mathbf{z})$ of these optimization problems readily yield the robust asymptotic bounds

$$t^{-1}d\beta_*(\mathbf{z}) \lesssim \mathbb{P}(\exists i : Z_i > tz_i) \lesssim t^{-1}d\beta^*(\mathbf{z}), \quad \text{as } t \rightarrow \infty.$$

Note that according to (1.3) it is enough to consider $\mathbf{z} \in \mathbb{S}^{d-1}$. It should also be stressed that our bounds address misspecification of the extremal dependence model exclusively, and so they are guaranteed to hold for sufficiently large scaling factor t only. Furthermore, we essentially optimize over the class of max-stable distributions, which is different from the univariate case analysis in Blanchet and Murthy (2016).

In Section 2 we provide details on the divergence $D_\mu(\mathbb{P}', \mathbb{P})$, and recall necessary results on multivariate extreme value theory, regular variation and spectral measures. The convex optimization problem (1.1) is solved in Section 3 and the simple square-root bound for the optimal value $V_\mu(\delta)$ is derived in Section 4, where we also identify a necessary and sufficient condition for this upper bound to coincide with $V_\mu(\delta)$. Based on these general results, in Section 5 we investigate robust bounds for small probabilities of tail regions in the bivariate case arising from the optimization problem (1.4) for $d = 2$. Several examples are given in Section 5.1 to illustrate the results. In Section 5.2 we conduct an experiment that shows the effectiveness of the robust bounds compared to classical confidence bounds when the model is misspecified. As a further application of our theory, Section 6 discusses how worst case bounds on the Value-at-Risk of a financial portfolio under model uncertainty can be derived. The Appendix contains some parametric families of spectral distributions, further comments about the degenerate maximizer of the problem in (1.4), and results on optimization for other divergences.

2 Preliminaries and the setup

2.1 Distribution model risk

Distribution model risk refers to the error made when using a simplified model of reality that is only an approximation to the data generating process. From a probabilistic point of view, this amounts to computing the quantity of interest, say the probability $\mathbb{P}(A)$ of some event A , using a wrong probability measure \mathbb{P} , which nevertheless is close in some sense to the true measure \mathbb{P}_{true} . The robust approach to this problem is to consider all measures \mathbb{P}' in some neighborhood of \mathbb{P} that should contain \mathbb{P}_{true} as well, and to find the maximal and the minimal values among all $\mathbb{P}'(A)$. These numbers then provide robust bounds on the true value $\mathbb{P}_{\text{true}}(A)$. This approach has become quite popular in financial mathematics, see Hansen and Sargent (2001); Ahmadi-Javid (2012); Breuer and Csiszár (2013a); Glasserman and Xu (2014) and references therein, and Blanchet and Murthy (2016) for an application to univariate extreme value statistics.

A natural way to define a neighborhood of measures around \mathbb{P} is to consider some form of divergence. Fix a dominating probability measure μ , i.e., such that

$\mathbb{P} \ll \mu$, and suppose for now that $\mathbb{P}' \ll \mu$. Letting $L = d\mathbb{P}/d\mu$ and $L' = d\mathbb{P}'/d\mu$ be the corresponding Radon–Nikodym derivatives we consider the standard squared L_μ^2 -distance

$$D_\mu(\mathbb{P}', \mathbb{P}) = \mathbb{E}_\mu(L' - L)^2, \quad (2.1)$$

where \mathbb{E}_μ denotes the expectation under probability measure μ . We put $D_\mu(\mathbb{P}', \mathbb{P}) = \infty$ if \mathbb{P}' is not absolutely continuous with respect to μ . It is noted that (2.1) is a special case of the so-called Bregman divergence, see, e.g., Breuer and Csiszár (2013a). Furthermore, by choosing $\mu = \mathbb{P}$ we get

$$D_\mathbb{P}(\mathbb{P}', \mathbb{P}) = \mathbb{E}(L' - 1)^2 = \mathbb{E}L'^2 - 1 \quad (2.2)$$

for all $\mathbb{P}' \ll \mathbb{P}$ with $L' = d\mathbb{P}'/d\mathbb{P}$. Moreover,

$$D_\mathbb{P}(\mathbb{P}', \mathbb{P}) \leq \delta \quad \text{iff} \quad \log \mathbb{E}L'^2 \leq \log(1 + \delta) = \delta',$$

where $\log \mathbb{E}L'^2$ is the well-known Rényi (power) divergence of order 2 of \mathbb{P}' from \mathbb{P} . In other words, neighborhoods of measures defined by $D_\mathbb{P}(\cdot, \mathbb{P}) \leq \delta$ coincide with second order Rényi divergence neighborhoods with radius δ' .

It is clear that the choice of the dominating measure μ has an impact on the solution of the optimization problem (1.1). Suppose, for instance, that \mathbb{P}, \mathbb{P}' and μ are defined on $[0, 1]$ and that they are absolutely continuous with respect to Lebesgue measure with densities f, f' and g , respectively. Then it holds that

$$D_\mu(\mathbb{P}', \mathbb{P}) = \int_0^1 \left(\frac{f'(\omega)}{g(\omega)} - \frac{f(\omega)}{g(\omega)} \right)^2 g(\omega) d\omega = \int_0^1 (f'(\omega) - f(\omega))^2 \frac{1}{g(\omega)} d\omega, \quad (2.3)$$

and so μ provides a mechanism of weighing the squared distance between f' and f . A similar weight function appears in e.g. Bücher et al. (2011) in the context of estimating the Pickands' function. Thus the dominating measure μ may be chosen according to our uncertainty about the measure \mathbb{P} .

In this study we leave out a detailed analysis of the choice of μ . Our default choice in applications to multivariate extremes is $\mu = \mathbb{P}$, which corresponds to Rényi divergence of order 2. We also provide an example where this choice is inappropriate, in which case the uniform dominating measure is used. Finally, the remaining parameter $\delta > 0$, representing our trust into the measure \mathbb{P} , has to be chosen by hand or derived from data. In Section 5.2 we use a straightforward heuristic procedure to estimate it from data.

2.2 Regular variation and spectral distributions

A d -dimensional random vector \mathbf{Z} is *multivariate regularly varying* in the non-negative orthant if there exists a sequence $a_t \rightarrow \infty$, as $t \rightarrow \infty$, and a Radon measure ν on $E = [0, \infty]^d \setminus \{0\}$ equipped with its Borel σ -algebra such that

$$t\mathbb{P}(\mathbf{Z}/a_t \in \cdot) \xrightarrow{v} \nu, \quad t \rightarrow \infty \quad (2.4)$$

in the sense of vague convergence, see, e.g., (Resnick, 2007, Ch. 6). The so-called *exponent measure* ν then satisfies the scaling property $\nu(tB) = t^{-\alpha}\nu(B)$ for all $t > 0$

and all Borel sets $B \subset E$ bounded away from 0, where $\alpha > 0$ is called the *tail index* of regular variation. Moreover, by switching to polar coordinates $\mathbf{z} \mapsto (\|\mathbf{z}\|, \mathbf{z}/\|\mathbf{z}\|) = (r, \boldsymbol{\omega})$ for the L_1 -norm $\|\mathbf{z}\| = \sum_{i=1}^d |z_i|$ on \mathbb{R}^d , the measure ν factorizes into

$$c \alpha r^{-\alpha-1} dr \times H(d\boldsymbol{\omega}),$$

where $c > 0$ and H is a probability measure, called the *spectral measure*, on the simplex \mathbb{S}^{d-1} equipped with its Borel σ -algebra. Importantly, (2.4) implies the following weak convergence to H :

$$\mathbb{P} \left(\frac{\mathbf{Z}}{\|\mathbf{Z}\|} \in \cdot \mid \|\mathbf{Z}\| > t \right) \xrightarrow{w} H, \quad \text{as } t \rightarrow \infty. \quad (2.5)$$

Without loss of generality we assume that ν is non-degenerate in the sense that $\nu(\{\mathbf{z} : z_i > 1\}) \neq 0$ for all $i = 1, \dots, d$. Otherwise, we may simply remove the components of the vector \mathbf{Z} that decrease at a faster rate. This implies that all marginal survival functions $\bar{F}_i(z) = 1 - F_i(z)$ are regularly varying with the same index $-\alpha$, and, moreover, for some $m_i > 0$,

$$\frac{\bar{F}_i(z)}{\bar{F}_1(z)} \rightarrow m_i \quad \text{as } z \rightarrow \infty \quad (2.6)$$

with $m_1 = 1$. That is, $\bar{F}_i(z)$ are equivalent in the limit up to multiplicative constants.

It is common to split the problem of multivariate tail estimation into estimation of marginal tails and estimation of the spectral distribution. The theory for univariate tail estimation is well-studied and there are many established methods to estimate the survival functions de Haan and Ferreira (2006); Resnick (2007). We therefore assume that the marginal tail models are continuous and correctly specified, and that the Z_i have been transformed to unit Pareto tails. That is, we generally assume that $\alpha = 1$ and $\bar{F}_i(z) = 1/z$ for large z , apart from Section 6, where we return to the general setup and the issue of standardization.

With the above standardization in mind we may choose $a_t = t$ in (2.4) leading to the approximation

$$\mathbb{P}(\mathbf{Z} \in tB) \approx \nu(B)/t$$

for large t and B bounded away from the origin with $\nu(\partial B) = 0$. A natural choice of such a set is given by $B_{\mathbf{z}} = E \setminus [\mathbf{0}, \mathbf{z}]$, where we may assume that $\|\mathbf{z}\| = 1$ because of the scaling property of ν . That is, we are interested in approximating the probability that at least one marginal is relatively large, namely $Z_i > tz_i$ for some i . Letting $\mathbf{Y} \in \mathbb{S}^{d-1}$ have the spectral distribution H one finds that

$$\nu(B_{\mathbf{z}}) = c \mathbb{E} \int 1_{\{\exists i: rY_i > z_i\}} r^{-2} dr = c \mathbb{E} \left(\max_{i=1}^d \frac{Y_i}{z_i} \right). \quad (2.7)$$

Moreover, according to the above standardization the exponent measure must satisfy

$$t \mathbb{P}(Z_i > t) \rightarrow 1 = \nu(\{\mathbf{z} : z_i > 1\}) = c \mathbb{E} Y_i$$

and hence $\mathbb{E} Y_i = 1/c$ for all i . But since $\sum_{i=1}^d Y_i = 1$, it must be that $c = d$, which yields the approximation in (1.3), our starting point for the robust approach. Importantly, any \mathbf{Y} satisfying these moment constraints gives rise to a valid spectral measure.

Remark 1. *The sum norm used throughout this paper is special in the sense that the constant c does not depend on the spectral measure. This makes it possible to employ the optimization problem in (1.1). Other norms would lead to the objective $\mathbb{E}'X(\mathbf{z})/\mathbb{E}'Y_1$, which does not comply with (1.1).*

A common way of representing the dependence structure in the bivariate case Pickands (1981); Klüppelberg and May (2006); Bücher et al. (2011) is by means of the so-called *Pickands' function*

$$A(z) = 2\mathbb{E}\{(1-z)Y_1 \vee z(1-Y_1)\}, \quad z \in [0, 1]. \quad (2.8)$$

Indeed, an easy transformation of (2.7) yields

$$\nu(B_{\mathbf{z}}) = \left(\frac{1}{z_1} + \frac{1}{z_2} \right) A\left(\frac{z_1}{z_1 + z_2} \right).$$

Importantly, the Pickands' dependence function $A : [0, 1] \rightarrow [1/2, 1]$ is convex and satisfies $z \vee (1-z) \leq A(z) \leq 1$. Moreover, any such function defines a unique exponent measure ν , see (de Haan and Ferreira, 2006, p. 226).

3 Convex optimization

In this section we solve the optimization problem (1.1) which, according to (2.1), can be rewritten in the convenient form

$$V_\mu(\delta) = \sup_{L' \geq 0} \{ \mathbb{E}_\mu(L'X) : \mathbb{E}_\mu L' = 1, \mathbb{E}_\mu(L' - L)^2 \leq \delta, \mathbb{E}_\mu(L'Y) = \mathbb{E}_\mu(LY) \}, \quad (3.1)$$

where the supremum is taken over all measurable functions $L' : \Omega \rightarrow [0, \infty)$ satisfying the stated constraints. This is a convex optimization problem in an infinite dimensional space allowing for a rather explicit solution given in Theorem 1. For related results without moment constraints see Breuer and Csiszár (2013a); Blanchet and Murthy (2016); Dey and Juneja (2012); Glasserman and Xu (2014). The latter two works also provide short derivations based on the strong duality theorem. There is, however, no reference to the strong duality theorem for infinite dimensional spaces which does require verification of certain conditions. Moreover, the issue with a distribution of X with some mass at its right end is not addressed in the literature.

3.1 The underlying measurable space

Before solving the optimization problem (1.1) or its equivalent version (3.1), let us comment on the underlying measurable space (Ω, \mathcal{F}) . Letting $\mathcal{G} = \sigma(X, \mathbf{Y}) \subset \mathcal{F}$ we assume that $L = d\mathbb{P}/d\mu$ is \mathcal{G} -measurable. That is, the choice of the dominating measure μ does not introduce additional randomness in the model, which is trivially the case for our default choice $\mu = \mathbb{P}$. The optimization problem (1.1) formulated on the measurable space (Ω, \mathcal{F}) and its analogue formulated on the measurable space (Ω, \mathcal{G}) lead to the same optimal value $V_\mu(\delta)$. This follows from Jensen's inequality:

$$\mathbb{E}_\mu(L' - L)^2 = \mathbb{E}_\mu(\mathbb{E}_\mu[(L' - L)^2|\mathcal{G}]) \geq \mathbb{E}_\mu(\mathbb{E}[L'|\mathcal{G}] - L)^2,$$

where the latter is the respective divergence on (Ω, \mathcal{G}) . Therefore, we may always consider the induced distributions of (X, \mathbf{Y}^\top) without changing the robust bounds. In the setting of (1.4) we may thus work on the Borel σ -algebra of \mathbb{S}^{d-1} . In fact, this can be seen as the modeling choice requiring little justification.

3.2 The optimal Radon–Nikodym derivative

Let us immediately present the solution to the optimization problem (3.1). It is noted that the proof of this result provides good intuition on the form of the solution. Throughout the paper, we will denote a maximizer of (3.1), if it exists, by L^* , and for any random variable X we put $\mathbb{E}^*X = \mathbb{E}_\mu(L^*X)$.

Theorem 1. *Assume that $\mathbb{E}_\mu X^2, \mathbb{E}_\mu Y_i^2, \mathbb{E}_\mu L^2 < \infty$ and let $\mathbb{E}\mathbf{Y} = \mathbf{y}$. Then L^* is a maximizer of the optimization problem (3.1) if and only if $\mathbb{E}_\mu L^* = 1, \mathbb{E}_\mu(L^*\mathbf{Y}) = \mathbf{y}$ and at least one of the following holds:*

(i) *there exist $a > 0, b, c_i \in \mathbb{R}, i = 1, \dots, d-1$, such that*

$$L^* = (aX + b + \mathbf{c}^\top \mathbf{Y} + L)_+ \quad \mu\text{-a.s. and} \quad \mathbb{E}_\mu(L^* - L)^2 = \delta.$$

(ii) *there exist $c_i \in \mathbb{R}, i = 1, \dots, d-1$, such that the distribution of $X + \mathbf{c}^\top \mathbf{Y}$ under μ has a positive mass at its upper end, $L^* = 0$ everywhere else μ -a.s., and the constraint $\mathbb{E}_\mu(L^* - L)^2 \leq \delta$ holds.*

Proof. Note that $\mathbb{E}_\mu L'^2 \leq 2\{\mathbb{E}_\mu(L' - L)^2 + \mathbb{E}_\mu L^2\} < \infty$ if $\mathbb{E}_\mu(L' - L)^2 \leq \delta$. So we may consider a normed vector space of μ -square-integrable L' and its convex subset defined by the additional requirement of $L' \geq 0$. Note also that $\mathbb{E}_\mu(L'|X|), \mathbb{E}_\mu(L'|Y_i|) < \infty$. Next, for the convex optimization problem (3.1) we define the corresponding Lagrangian:

$$\mathcal{L}(L') = \mathbb{E}_\mu(L'X) - a(\mathbb{E}_\mu(L' - L)^2 - \delta) + b(\mathbb{E}_\mu L' - 1) + \mathbf{c}^\top(\mathbb{E}_\mu(L'\mathbf{Y}) - \mathbf{y}), \quad (3.2)$$

where $a \geq 0, b, c_i \in \mathbb{R}$. The strong duality theorem, see e.g. (Mitter, 2008, Thm. 4), asserts that L^* is a maximizer of the original problem if and only if L^* is a maximizer of $\sup_{L' \geq 0} \mathcal{L}(L')$ for some $a \geq 0, b, c_i \in \mathbb{R}$, such that the constraints hold as well as so-called complementary slackness:

$$\mathbb{E}_\mu L^* = 1, \quad \mathbb{E}_\mu(L^*\mathbf{Y}) = \mathbf{y}, \quad \mathbb{E}_\mu(L^* - L)^2 \leq \delta, \quad a(\mathbb{E}_\mu(L^* - L)^2 - \delta) = 0.$$

For this result to be true it is sufficient to verify Slater's condition: $\exists L' \geq 0$ such that $\mathbb{E}_\mu(L' - L)^2 < \delta$ and $\mathbb{E}_\mu L' = 1, \mathbb{E}_\mu(L'\mathbf{Y}) = \mathbf{y}$, but this is clearly satisfied by $L' = L$.

Hence it is left to solve the dual problem $\sup_{L' \geq 0} \mathcal{L}(L')$ for fixed $a \geq 0, b, c_i \in \mathbb{R}$. Since $\mathcal{L}(L')$ is concave in L' , a sufficient and necessary condition for a maximizer L^* of the dual problem is

$$g'_{L^*, L'}(0+) \leq 0 \quad \forall L' \geq 0, \quad \text{where } g_{L^*, L'}(t) = \mathcal{L}(L^*(1-t) + L't),$$

that is, one looks down from L^* . But $g_{L^*, L}(t)$ is given by

$$\mathbb{E}_\mu \left\{ (L^*(1-t) + L't)(X + b + \mathbf{c}^\top \mathbf{Y}) - a(L^*(1-t) + L't - L)^2 + a\delta - b - \mathbf{c}^\top \mathbf{y} \right\},$$

which can be differentiated under the expectation sign, see e.g. (Williams, 1991, A16), yielding

$$\mathbb{E}_\mu(L' - L^*)(X + b + \mathbf{c}^\top \mathbf{Y} - 2a(L^* - L)) \leq 0, \quad \forall L' \geq 0. \quad (3.3)$$

This implies that $X + b + \mathbf{c}^\top \mathbf{Y} - 2a(L^* - L) \leq 0$ and $L^* = 0$ when the inequality is strict μ -a.s., because otherwise we may choose $L' \geq 0$ to invalidate (3.3). But the latter clearly implies (3.3) and so we have the equivalence. Thus for $a > 0$ we get

$$L^* = \left(\frac{X + b + \mathbf{c}^\top \mathbf{Y}}{2a} + L \right)_+ \quad \mu\text{-a.s.},$$

which is equivalent to (i). If $a = 0$ then $X + \mathbf{c}^\top \mathbf{Y} \leq -b$ and $L^* = 0$ when the inequality is strict μ -a.s. Hence $\mu(X + \mathbf{c}^\top \mathbf{Y} = -b) > 0$, because otherwise $\mathbb{E}_\mu L^* = 0$. This yields (ii). \square

Remark 2. Suppose that for some $\delta' > 0$ there is L^* as in (ii) of Theorem 1 which also satisfies the equality constraints. Then such L^* must be a maximizer of (3.1) for any $\delta \geq \delta'$. This implies that the corresponding optimal value $V_\mu(\delta')$ is the maximal possible for any $\delta > 0$, and in particular it does not increase with further increase of δ . In the following we let δ^{**} be the minimal such δ' , and $\delta^{**} = \infty$ if no such δ' exists. The optimizer L^* then has the form given in (ii) of Theorem 1 if and only if $\delta \geq \delta^{**}$.

We believe that some further clarification of Theorem 1 is necessary. Normally, we only need to look at (i), whereas (ii) corresponds to a rather pathological case explained in Remark 2. A necessary condition for the latter is that δ is sufficiently large, $\delta \geq \delta^{**}$, and also that (X, \mathbf{Y}^\top) satisfies the condition mentioned in (ii) for some \mathbf{c} , because otherwise $\delta^{**} = \infty$. The following two simple examples will provide some further intuition.

Example 1. Consider the optimization problem without moment constraints when $\mu = \mathbb{P}$ and the distribution of X has a positive mass $p > 0$ at its upper end \bar{x} . Then the optimizer L^* in (ii) puts all the mass on $\{X = \bar{x}\}$ achieving $\mathbb{E}^* X = \bar{x}$ which is the maximal possible value for any $\delta > 0$. But we must have

$$\delta + 1 \geq \mathbb{E} L^{*2} = p \mathbb{E}(L^{*2} | X = \bar{x}) \geq p \mathbb{E}^2(L^* | X = \bar{x}) = 1/p,$$

because $\mathbb{E}(L^* | X = \bar{x}) = 1/p$. So if $\delta \geq 1/p - 1$ then we can choose $L^* = 1/p \mathbf{1}_{\{X = \bar{x}\}}$ yielding the maximal possible optimal value \bar{x} , but otherwise we must consider L^* from (i). In particular, we have $\delta^{**} = 1/p - 1$.

Example 2. This example shows that in general (ii) does not require the distribution of X to have a mass at its upper end. Take $\mu = \mathbb{P}$ and consider the case of one constraint where $X = Y \mathbf{1}_A$ for some event A and $Y > 0$ on A^c . Hence $X - Y = -Y \mathbf{1}_{A^c} \leq 0$ a.s. meaning that $X - Y$ has mass $\mathbb{P}(A)$ at 0. Clearly, $\mathbb{E}' X \leq \mathbb{E}' Y = y$ for any $\mathbb{P}' \ll \mathbb{P}$, whereas $\mathbb{E}^* X = \mathbb{E}^* Y = y$ if L^* puts all the mass on A . It is only left to ensure that there is such L^* preserving the expectation of Y .

The following observation will later lead to the square-root bound (1.2), an upper bound on $V_\mu(\delta)$.

Remark 3. *If the non-negativity constraint on L' is removed in the optimization problem (3.1) then*

$$\hat{L} = aX + b + \mathbf{c}^\top \mathbf{Y} + L, \quad a > 0, b, c_i \in \mathbb{R}$$

satisfying $\mathbb{E}_\mu \hat{L} = 1, \mathbb{E}_\mu (\hat{L} - L)^2 = \delta, \mathbb{E}_\mu (\hat{L} \mathbf{Y}) = \mathbf{y}$ is a maximizer.

Finally, we note that the optimization problem (1.1) can be solved for some other popular divergences such as Rényi and Kullback–Leibler divergences; see Appendix A.3 for details. In both cases it is assumed that the dominating measure μ coincides with \mathbb{P} . In this paper, however, we aim at providing simple explicit bounds while giving flexibility in defining the neighborhoods of measures by choosing an appropriate dominating measure μ , and therefore we exclusively focus on the divergence $D_\mu(\mathbb{P}', \mathbb{P})$ defined in (2.1).

3.3 Computing the optimal value

Let us consider the main case (i) of Theorem 1 where $L^* = (aX + b + \mathbf{c}^\top \mathbf{Y} + L)_+$. In order to find $a > 0, b, c_i \in \mathbb{R}, i = 1, \dots, d-1$, we need to solve a system of $d+1$ corresponding equations:

$$\mathbb{E}_\mu(L^* - L) = 0, \quad \mathbb{E}_\mu(L^* - L)^2 = \delta, \quad \mathbb{E}_\mu\{(L^* - L)\mathbf{Y}\} = \mathbf{0}. \quad (3.4)$$

If a solution is found then the optimal value is given by $V_\mu(\delta) = \mathbb{E}_\mu(L^* X)$. Note, however, that the maximizer L^* may be of a different form given in case (ii) of Theorem 1, which corresponds to $\delta \geq \delta^{**}$ and the maximal possible optimal value. In some cases δ^{**} has an explicit formula, whereas in some other cases identification of δ^{**} requires solving yet another convex optimization problem. This latter problem usually can be avoided in practice when plotting $V_\mu(\delta)$ as a function of δ , because solving (3.4) becomes problematic only when $V_\mu(\delta)$ is close to its maximal value. Some further details concerning δ^{**} in the particular case of problem (1.4) are given in Section A.2.

In the important case of $\mu = \mathbb{P}$ we have $L \equiv 1$. Incorporating the latter into the constant b reduces the number of equations by one. Indeed, we then may consider $L^* = aU_+$ with $U = U(b, \mathbf{c}) = X + b + \mathbf{c}^\top \mathbf{Y}$ and the equations

$$\mathbb{E}(U_+) = 1/a, \quad \text{var}(U_+) = \delta/a^2, \quad \text{cov}(U_+, \mathbf{Y}) = \mathbf{0}^\top.$$

Hence we need to find constants $b, c_i \in \mathbb{R}, i = 1, \dots, d-1$, such that $U > 0$ with positive probability and

$$\text{cov}(U_+, \mathbf{Y}) = \mathbf{0}^\top, \quad \delta = \text{var}(U_+)/\mathbb{E}^2(U_+).$$

If a solution is found then $V_\mathbb{P}(\delta) = \mathbb{E}(U_+ X)/\mathbb{E}(U_+)$ and the corresponding Radon–Nikodym derivative is $L^* = U_+/\mathbb{E}(U_+)$. Moreover, this suggests a parametric approach to plot $V_\mathbb{P}(\delta), \delta > 0$: (1) fix b in some range and try to find $c_i, i = 1, \dots, d-1$,

such that $\text{cov}(U_+, \mathbf{Y}) = \mathbf{0}^\top$ and $U > 0$ with positive probability, (2) plot $(\delta, V_\mathbb{P}(\delta))$ for various values of b .

Solving the above systems of non-linear equations may not be trivial, but it can be done numerically for a moderate dimension d . Note that evaluation of the left hand sides in (3.4) for a given choice of constants a, b, c_i requires integration with respect to the joint distribution of (X, \mathbf{Y}^\top) . Thus each evaluation is costly even for small d . In the following section we provide an upper bound for the optimal value of a simple explicit form that does not require solving any equation. Moreover, in our applications we observed that this bound often coincides with the optimal value or is very close to it.

4 Robust bounds of a simple form

Throughout this section we assume that $\mathbb{E}_\mu X^2, \mathbb{E}_\mu Y_i^2, \mathbb{E}_\mu L^2 < \infty$, and recall that $\Sigma_\mu(X, \mathbf{Y})$ denotes the μ -covariance matrix of the vector (X, \mathbf{Y}^\top) . The following result provides a robust bound of a simple form on $\mathbb{E}'X$ under the moment constraints $\mathbb{E}'\mathbf{Y} = \mathbb{E}\mathbf{Y}$ and in the neighborhood defined by $D_\mu(\mathbb{P}', \mathbb{P}) \leq \delta$. In the following we call it a *robust square-root bound*.

Theorem 2. *If $\Sigma_\mu(X, \mathbf{Y})$ is invertible then the optimal value of (3.1) for $\delta > 0$ satisfies*

$$V_\mu(\delta) \leq \mathbb{E}X + \sqrt{\delta \frac{\det\{\Sigma_\mu(X, \mathbf{Y})\}}{\det\{\Sigma_\mu(\mathbf{Y})\}}},$$

which holds with equality if and only if

$$X - \mathbb{E}_\mu X - \text{cov}_\mu(X, \mathbf{Y})\Sigma_\mu(\mathbf{Y})^{-1}(\mathbf{Y} - \mathbb{E}_\mu \mathbf{Y}) + \sqrt{\frac{\det\{\Sigma_\mu(X, \mathbf{Y})\}}{\delta \det\{\Sigma_\mu(\mathbf{Y})\}}}L \geq 0 \quad \mu\text{-a.s.} \quad (4.1)$$

Proof. The covariance matrix $\Sigma_\mu(X, \mathbf{Y})$ is positive definite and so must be $\Sigma_\mu(\mathbf{Y})$, showing that $\det\{\Sigma_\mu(X, \mathbf{Y})\}, \det\{\Sigma_\mu(\mathbf{Y})\} > 0$. Consider case (i) of Theorem 1 and note that we may rescale the constants b, \mathbf{c} so that $L^* = (aU + L)_+$, where $U = X + b + \mathbf{c}^\top \mathbf{Y}$.

First, we assume that $aU + L \geq 0$ μ -a.s. Then according to (3.4) we have

$$\mathbb{E}_\mu U = 0, \quad \text{var}_\mu(U) = \delta/a^2, \quad \text{cov}_\mu(\mathbf{Y}, U) = \mathbf{0}.$$

Denoting $\boldsymbol{\sigma} = \text{cov}_\mu(\mathbf{Y}, X)$ the latter reads as

$$\mathbf{0} = \text{cov}_\mu(\mathbf{Y}, U) = \text{cov}_\mu(\mathbf{Y}, X + b + \mathbf{c}^\top \mathbf{Y}) = \boldsymbol{\sigma} + \Sigma_\mu(\mathbf{Y})\mathbf{c}$$

showing that $\mathbf{c}^\top = -\boldsymbol{\sigma}^\top \Sigma_\mu(\mathbf{Y})^{-1}$. Similarly,

$$\begin{aligned} \text{var}_\mu(U) &= \text{cov}_\mu(U, X + b + \mathbf{c}^\top \mathbf{Y}) = \text{cov}_\mu(U, X) = \text{var}_\mu(X) + \mathbf{c}^\top \boldsymbol{\sigma} \\ &= \text{var}_\mu(X) - \boldsymbol{\sigma}^\top \Sigma_\mu(\mathbf{Y})^{-1} \boldsymbol{\sigma} \end{aligned} \quad (4.2)$$

which is $\det\{\Sigma_\mu(X, \mathbf{Y})\}/\det\{\Sigma_\mu(\mathbf{Y})\}$ according to the well-known formula for the determinant of a block matrix, see Schur (1917) or (Ouellette, 1981, Eq. (1.3));

the expression in (4.2) is called the Schur complement of $\Sigma_\mu(\mathbf{Y})$ with respect to $\Sigma_\mu(X, \mathbf{Y})$. Hence we find that

$$a = \sqrt{\delta \det\{\Sigma_\mu(\mathbf{Y})\} / \det\{\Sigma_\mu(X, \mathbf{Y})\}}, \quad b = -\mathbb{E}_\mu X + \boldsymbol{\sigma}^\top \Sigma_\mu(\mathbf{Y})^{-1} \mathbb{E}_\mu \mathbf{Y}.$$

According to (4.2) we finally get

$$\begin{aligned} V_\mu(\delta) &= \mathbb{E}^* X = \mathbb{E}_\mu\{(aU + L)X\} = \mathbb{E}X + a \operatorname{cov}_\mu(U, X) \\ &= \mathbb{E}X + a \frac{\det\{\Sigma_\mu(X, \mathbf{Y})\}}{\det\{\Sigma_\mu(\mathbf{Y})\}}, \end{aligned}$$

which readily yields $V_\mu(\delta)$ under the assumption of non-negativity of $aU + L$. But in any case we have an upper bound according to Remark 3. Finally, we have an exact expression for $V_\mu(\delta)$ if

$$\begin{aligned} 0 &\leq U + L/a \\ &= X - \mathbb{E}_\mu X + \boldsymbol{\sigma}^\top \Sigma_\mu(\mathbf{Y})^{-1} \mathbb{E}_\mu \mathbf{Y} - \boldsymbol{\sigma}^\top \Sigma_\mu(\mathbf{Y})^{-1} \mathbf{Y} + \sqrt{\frac{\det\{\Sigma_\mu(X, \mathbf{Y})\}}{\delta \det\{\Sigma_\mu(\mathbf{Y})\}}} L \end{aligned}$$

holds μ -a.s., which completes the proof. \square

The assumption that $\Sigma_\mu(X, \mathbf{Y})$ is invertible is not a restriction, because otherwise either some moment constraints are redundant and so can be removed, or X can be expressed as a linear combination of Y_i and so $\mathbb{E}X$ is determined by the moment constraints. Moreover, there is a link to the control variates method for variance reduction, where $\det\{\Sigma_\mu(X, \mathbf{Y})\} / \det\{\Sigma_\mu(\mathbf{Y})\}$ corresponds to the minimal possible variance of $X + \mathbf{c}^\top \mathbf{Y}$ for an arbitrary vector \mathbf{c} , see (Asmussen and Glynn, 2007, Sec. V.2). Furthermore, if (X, \mathbf{Y}^\top) is jointly normal then the above fraction of the determinants is the variance of X conditional on $\mathbf{Y} = \mathbf{y}$. Finally, in some applications it may be important to understand when the optimizing measure is equivalent to the dominating measure μ , i.e., L^* is strictly positive μ -a.s. It is easy to see that this happens if and only if (4.1) holds with strict inequality. As before here we assume that $\det\{\Sigma_\mu(X, \mathbf{Y})\} \neq 0$, which additionally ensures that the case (ii) of Theorem 1 cannot be used to construct a strictly positive L^* .

The condition (4.1) implies that there exists $\delta^* \in [0, \infty]$ such that the robust square-root bound is exact for all $\delta \in [0, \delta^*]$, but otherwise it is conservative. Because of the form of the square-root bound we observe that necessarily $\delta^* \leq \delta^{**}$, where the latter is defined in Remark 2. Figure 1 illustrates these quantities, the optimal value and the square-root bound. Note also that the optimal value $V_\mu(\delta)$ must be a concave function of δ , which is easily seen from (3.1).

The corresponding minimization problem is solved by considering $-X$ instead of X in (3.1), in which case we define L_*, δ_* and δ_{**} analogously to L^*, δ^* and δ^{**} . In particular, we have the following lower square-root bound.

Corollary 1. *Under the assumptions of Theorem 2 it holds for $\delta > 0$ that*

$$\inf_{\mathbb{P}'} \{\mathbb{E}' X : D_\mu(\mathbb{P}', \mathbb{P}) \leq \delta, \mathbb{E}' \mathbf{Y} = \mathbb{E} \mathbf{Y}\} \geq \mathbb{E} X - \sqrt{\delta \frac{\det\{\Sigma_\mu(X, \mathbf{Y})\}}{\det\{\Sigma_\mu(\mathbf{Y})\}}}$$




Figure 1: The optimal value $V_\mu(\delta)$ (lower curve) and the square-root bound (upper curve) as functions of δ .

with equality if and only if

$$X - \mathbb{E}_\mu X - \text{cov}_\mu(X, \mathbf{Y}) \Sigma_\mu(\mathbf{Y})^{-1} (\mathbf{Y} - \mathbb{E}_\mu \mathbf{Y}) - \sqrt{\frac{\det\{\Sigma_\mu(X, \mathbf{Y})\}}{\delta \det\{\Sigma_\mu(\mathbf{Y})\}}} L \leq 0 \quad \mu\text{-a.s.} \quad (4.3)$$

Proof. Consider $-X$ in place of X in Theorem 2 and note that $\det\{\Sigma_\mu(X, \mathbf{Y})\}$ stays the same. The condition for equality immediately follows from (4.1). \square

Remark 4. When $\mu = \mathbb{P}$, with $\Sigma = \Sigma_{\mathbb{P}}$ it holds that

$$\begin{aligned} \delta^* &= \frac{\det\{\Sigma(X, \mathbf{Y})\}}{b_*^2 \det\{\Sigma(\mathbf{Y})\}} 1_{\{b^* > 0\}} + \infty 1_{\{b^* \leq 0\}}, \\ \delta_* &= \frac{\det\{\Sigma(X, \mathbf{Y})\}}{b_*^2 \det\{\Sigma(\mathbf{Y})\}} 1_{\{b_* < 0\}} + \infty 1_{\{b_* \geq 0\}}, \end{aligned} \quad (4.4)$$

where b^* and b_* are the essential supremum and the essential infimum of

$$\text{cov}(X, \mathbf{Y}) \Sigma(\mathbf{Y})^{-1} (\mathbf{Y} - \mathbb{E} \mathbf{Y}) - (X - \mathbb{E} X),$$

respectively. In particular, if $|X|$ and all $|Y_i|$ are bounded a.s. then necessarily $\delta^*, \delta_* > 0$.

In the case of no moment constraints, $d = 1$, the square-root bounds on $\mathbb{E}' X$ in the respective ball of measures are given by

$$\mathbb{E} X \pm \sqrt{\delta \text{var}_\mu(X)}.$$

In the case of one constraint, $d = 2$, we obtain the square-root bounds

$$\mathbb{E} X \pm \sqrt{\delta \text{var}_\mu(X) (1 - \text{corr}_\mu^2(X, Y))}, \quad (4.5)$$

and the corresponding δ^*, δ_* can be computed from (4.1) and (4.3). Notice that the bounds become tighter in presence of a constraint when X and Y are correlated.

It is important to note that the exact robust bounds become tighter or stay the same when an extra moment constraint is added, which follows immediately from (3.1). The same is true for the square-root bounds. This either follows from the proof of Theorem 2 and Remark 3, or from the following analysis based on block matrix algebra. Letting $\mathbf{Y}^+ = (Y_1, \dots, Y_d)^\top$, we need to show that

$$\det\{\Sigma_\mu(X, \mathbf{Y}^+)\} / \det\{\Sigma_\mu(\mathbf{Y}^+)\} \leq \det\{\Sigma_\mu(X, \mathbf{Y})\} / \det\{\Sigma_\mu(\mathbf{Y})\}.$$

This inequality follows by rewriting it using the Schur complements as in (4.2) and applying the block matrix inversion formula (Ouellette, 1981, Thm. 2.7). By doing so we find that this inequality is strict unless

$$\text{cov}_\mu(X, Y_d) = \text{cov}_\mu(X, \mathbf{Y})\Sigma_\mu(\mathbf{Y})^{-1}\text{cov}_\mu(\mathbf{Y}, Y_d).$$

In other words, this condition corresponds to the case when the extra moment constraint on Y_d does not improve the square-root bounds, assuming that the enlarged covariance matrix $\Sigma_\mu(X, \mathbf{Y}^+)$ is invertible.

5 Bounds on Pickands' dependence function

In this section we apply the bounds from the previous sections to assess the model misspecification error in multivariate extremes with the focus on tail probabilities in (1.3). For illustration, we consider the bivariate case and note that the extension to higher dimensions is analogous. Possible computational challenges will be discussed in Section 6 presenting another application of the robust approach to multivariate extremes.

Recall that we exclusively address misspecification of the spectral distribution. In the bivariate case, this distribution is defined on the simplex \mathbb{S}^1 , and it is thus effectively one-dimensional. In the following we assume that the corresponding random variable $Y = Y_1 \in [0, 1]$ μ -a.s. Alternatively to (1.4), we may directly consider the robust bound on the Pickands' dependence function defined in (2.8):

$$\sup_{\mathbb{P}'} \{2\mathbb{E}'\{(1-z)Y \vee z(1-Y)\} : D_\mu(\mathbb{P}', \mathbb{P}) \leq \delta, \mathbb{E}Y = 1/2\}, \quad (5.1)$$

i.e., we take $X = X(z) = 2\{(1-z)Y \vee z(1-Y)\}$ for a fixed $z \in (0, 1)$. The respective optimal value $A^*(z; \delta)$ provides the asymptotic robust upper bound on the tail probability

$$\mathbb{P}(Z_1 > tz_1 \text{ or } Z_2 > tz_2) \lesssim t^{-1} \left(\frac{1}{z_1} + \frac{1}{z_2} \right) A^* \left(\frac{z_1}{z_1 + z_2}; \delta \right), \quad t \rightarrow \infty$$

The lower bound is obtained similarly using the optimal value $A_*(z; \delta)$ of the corresponding minimization problem.

In addition to the exact robust bounds we will use the corresponding square-root bounds, and also bounds in the model class for comparison. More precisely, we consider the following:

- (a) Robust bounds $A^*(z; \delta)$ and $A_*(z; \delta)$ given by Theorem 1, which can be computed as explained in Section 3.3. Details on the identification of δ^{**} and δ_{**} are postponed to the Appendix A.2, because this is not essential for applications.
- (b) Robust square-root bounds $\hat{A}^*(z; \delta) \geq A^*(z; \delta)$ and $\hat{A}_*(z; \delta) \leq A_*(z; \delta)$ given by (4.5). These are conservative bounds that are easy to compute. Moreover, they are exact when $\delta \leq \delta^*$ and $\delta \leq \delta_*$ for upper and lower bounds, respectively.

- (c) Exact bounds in the model class that are not robust under model misspecification. That is, we impose the restriction that Y under \mathbb{P}' belongs to the chosen model class. These bounds are easy to compute for, e.g., one-parameter families, but otherwise it can be a hard problem. This paper addresses model misspecification issues and so the bounds in the model class will be given only for comparison.

Remark 5. *The bounds in (a) and (b) on Pickands' function directly provide robust bounds on the extremal coefficient $\theta = 2A(1/2) \in [1, 2]$, a commonly used summary statistic for dependence in multivariate and spatial extreme value statistics (Schlather and Tawn, 2003).*

Regarding the optimization problem (5.1) it is convenient to switch to the induced distributions of Y , see also Section 3.1. Thus we assume that $\Omega = [0, 1]$ and \mathcal{F} is the respective Borel σ -algebra, and that $Y(\omega) = \omega$. In Section 2.1 we claimed that the choice of the dominating measure μ reflects our uncertainty about the measure \mathbb{P} . Throughout this paper, however, our main choice is $\mu = \mathbb{P}$ leading to the Rényi divergence of order 2. For the purpose of illustration we also use $\mu = \text{Leb}[0, 1]$ assigning uniform weights; see (2.3). In the following we discuss some further simplifications of the general theory in these two particular cases.

If $\mu = \mathbb{P}$, then computing the robust bounds requires solving a system of two non-linear equations, see Section 3.3 for details. On the contrary, the square-root bounds (4.5) are always explicit, and we only need to compute $\mathbb{E}X$, $\text{var}(X)$ and $\text{corr}(X, Y)$. Moreover, Remark 4 provides simple expressions for δ^* and δ_* in terms of b^* and b_* . Importantly, the latter two can be given explicitly under a minor assumption that $0, 1, z$ are in the support of the distribution of Y :

$$\begin{aligned} b^* &= \mathbb{E}X - \rho/2 + (-2z) \vee (\rho z - 2z(1 - z)) \vee (\rho - 2(1 - z)), \\ b_* &= \mathbb{E}X - \rho/2 + (-2z) \wedge (\rho - 2(1 - z)), \end{aligned}$$

where $\rho = \text{cov}(X, Y)/\text{var}(Y)$. Indeed,

$$b^* = \mathbb{E}X - \rho/2 + \text{ess sup}(\rho Y - X), \quad b_* = \mathbb{E}X - \rho/2 + \text{ess inf}(\rho Y - X)$$

according to Remark 4, and $\rho Y - X$ achieves its maximum in one of the points $Y = 0, Y = z$ or $Y = 1$, and its minimum in $Y = 0$ or $Y = 1$. Further comments concerning δ^{**} and δ_{**} will be given in Section A.2.

If $\mu = \text{Leb}$, then for the exact robust bounds we need to solve a system of three non-linear equations. Concerning the square-root bounds, we observe from (4.5) that their width is determined by the dominating measure μ , whereas L affects the center $\mathbb{E}X$ and the values of δ^*, δ_* only; see (4.1) and (4.3). A simple calculation based on (4.5) yields the following square-root bounds for an arbitrary density L :

$$\mathbb{E}X \pm \sqrt{\delta \frac{4}{3} z^3 (1 - z)^3}. \quad (5.2)$$

Furthermore, for any δ one can provide a lower bound on the density $L(w)$ so that the square-root upper and lower bounds are exact. In particular, one can show that if L approaches 0 at one of the ends of the interval $[0, 1]$ then $\delta_* = 0$, that is, the lower square-root bound is never exact. Similarly, if L approaches 0 at z then $\delta^* = 0$.

5.1 Illustration of the bounds

In the beginning of this section we provided a list of three bounds on the Pickands' dependence function $A(z)$: (a) exact robust bounds, (b) conservative square-root bounds and (c) bounds in the model class. Let us illustrate these bounds for different divergence levels δ with an example of a Hüsler–Reiss spectral distribution; see Appendix A.1 for several common parametric families of spectral distributions. The Hüsler–Reiss distribution has a single parameter $\lambda \in (0, \infty)$ and we fix it to $\lambda = 0.6$. Furthermore, we consider $z = 0.4$ and use two dominating measures: $\mu = \mathbb{P}$ and $\mu = \text{Leb}$. The upper panels of Figure 2 show the bounds as functions of divergence δ . The middle and lower panels depict the Hüsler–Reiss density $h_{0.6}(\omega)$ as well as the optimizing densities corresponding to the upper and lower bounds for a particular choice of $\delta = 0.4$. In order to make comparisons easier, the densities with respect to the Lebesgue measure are depicted in both cases, and so for $\mu = \mathbb{P}$ we plot $L^*(\omega)h_{0.6}(\omega)$ and $L_*(\omega)h_{0.6}(\omega)$ rather than $L^*(\omega)$ and $L_*(\omega)$. Finally, there is no density corresponding to the square-root bound when δ is larger than δ^* or δ_* for the upper and lower bound, respectively. Nevertheless, there always exists the corresponding pseudo-density which is not necessarily non-negative; see Remark 3. These pseudo-densities are also included in Figure 2.

Additionally, we find $\delta^* = 0.36, \delta_* = 0.14$ when $\mu = \mathbb{P}$, and $\delta^* = 0.43, \delta_* = 0$ when $\mu = \text{Leb}$, respectively. Thus in the case of $\mu = \text{Leb}$ the upper square-root bound is exact for the chosen level $\delta = 0.4$ and so L^* and the corresponding pseudo-density coincide. In the other cases, exact and square-root bounds do not coincide, but it can be seen that the square-root bound is still a very good approximation of the exact robust bound even when δ is much larger than δ^* or δ_* . Furthermore, the upper bounds in the model class are obtained for $\lambda = 0.737$ and $\lambda = 0.844$ according to $\mu = \mathbb{P}$ and $\mu = \text{Leb}$, and the lower bounds for $\lambda = 0.367$ and $\lambda = 0.366$.

Let us make some final observations concerning the optimizing densities. Firstly, when maximizing $A(z)$ the probability mass is shifted from around z towards 0 and 1. Conversely, when minimizing $A(z)$ the mass is shifted towards z . Secondly, when $\mu = \mathbb{P}$ the density corresponding to the exact upper bound approaches 0 at both ends, and it does not do so when $\mu = \text{Leb}$. The reason is that the chosen Hüsler–Reiss spectral density decays faster than any power at 0 and at 1, and so Rényi divergence as defined in (2.2) is finite only if the density of \mathbb{P}' decays fast at 0 and at 1. This issue will arise again in Section 5.2 describing our experiments.

5.2 Experiments

In this section we show how the robust bounds are able to capture correctly the uncertainty due to model misspecification in a statistical estimation problem. They remain reliable in situations where classical confidence bounds would underestimate the statistical error.

As an application of our results we consider the estimation of tail probabilities of the bivariate, regularly varying random vector \mathbf{Z} ; see Section 2.2. Throughout this section we assume that \mathbf{Z} follows an asymmetric logistic distribution with dependence parameter $a \in (0, 1)$ and asymmetry parameters $b_1, b_2 \in [0, 1]$ as defined in Appendix A.1.2, so that the limiting spectral measure $H = \text{AL}(a, b_1, b_2) = \mathbb{P}_{\text{true}}$




Figure 2: Upper panels contain the value of $A(z)$ for \mathbb{P} being Hüsler-Reiss with $\lambda = 0.6$ and $z = 0.4$ (blue), exact robust bounds (black), the square-root bounds (green) and bounds in the model class (dashed purple) as functions of δ . Middle and lower panels show the Hüsler-Reiss density and the optimizing (Lebesgue-)densities for $\delta = 0.4$ corresponding to upper and lower bounds, respectively. The pseudo-densities corresponding to square-root bounds are given in green.

in (2.5) has the density (A.4). We conduct several experiments that illustrate the use of the robust bounds in practical applications. All our experiments are carried out according to the following scheme.

- (a) Simulate n data $\mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(n)}$ from a bivariate asymmetric logistic distribution using the *R*-package Stephenson (2002).
- (b) Transform the samples to polar coordinates as in Section 2.2, and choose $r > 0$ such that there are $k < n$ of the radii exceeding the threshold r . According to (2.5), the corresponding angles, say $Y^{(1)}, \dots, Y^{(k)}$, are approximate realizations of the spectral distribution. The choice of the threshold r is a trade-off between the sample size k and the approximation error.
- (c) Choose a parametric family for the spectral distribution and fit it to the observations $Y^{(1)}, \dots, Y^{(k)}$, using maximum likelihood estimation. The parametric family can either be the correct asymmetric logistic model, or a misspecified model such as the Hüsler–Reiss or the extremal- t described in Appendices A.1.1 and A.1.3, respectively. This model of the spectral distribution defines our probability measure \mathbb{P} .
- (d) Plot the Pickands’ dependence functions A_{true} and A corresponding to the true asymmetric logistic model \mathbb{P}_{true} and the estimate \mathbb{P} from (c), respectively; see (2.8).
- (e) Estimate the divergence of the data from the fitted model $\delta = D_{\mu}(\mathbb{P}_{\text{data}}, \mathbb{P})$ using a naive approach: estimate the density (and point masses) from the given k observations and plug it into (2.1) together with the model density from (c). Alternative methods for divergence estimation can be found in, e.g., Póczos and Schneider (2011). For comparison, we also compute the true divergence $\delta_{\text{true}} = D_{\mu}(\mathbb{P}_{\text{true}}, \mathbb{P})$ of the true underlying asymmetric logistic distribution from the fitted model.
- (f) Compute the robust square-root bounds $\hat{A}^*(z; \delta)$ and $\hat{A}_*(z; \delta)$ for Pickands’ function A using δ computed in (e). The exact robust bounds, which are considerably harder to compute, are very close to the square-root bounds and we omit them for clarity of the plots.
- (g) Compute the classical 95%-confidence bounds for the Pickands’ function by non-parametric bootstrap. This is based on 300 estimates of the model parameters as in (c), each for a resample of the data with replacement. We plot these bounds around A .

Remark 6. *Let us remark that instead of simulating data from the asymmetric logistic distribution we could have used any bivariate distribution from its max-domain of attraction, because we rely on a limiting result in (b) to approximate realizations of Y . Importantly, it is the limiting asymmetric logistic distribution and the corresponding spectral distribution of Y which are of main interest since they provide a way to extrapolate tail probabilities out of the sample.*

The basic information on the four experiments is given in Table 1, and the corresponding plots are given in Figure 3. In all experiments we use $k = 500$ exceedances. In experiment #1 we fit a symmetric extremal- t model to an asymmetric logistic

Table 1: Details on the four experiments. AL: asymmetric logistic, HR: Hüsler–Reiss, ET: extremal- t .

#	\mathbb{P}_{true}	n	Fitted model \mathbb{P}	μ	δ	δ_{true}
1	AL(0.4, 0.7, 1)	20000	ET(0.65, 1.21)	\mathbb{P}	0.34	0.35
2	AL(0.5, 1, 1)	20000	HR(0.61)	Leb	0.05	0.06
3	AL(0.5, 1, 1)	2000	ET(0.88, 3.37)	\mathbb{P}	0.05	0.05
4	AL(0.5, 0.9, 0.5)	20000	AL(0.49, 0.85, 0.54)	\mathbb{P}	0.02	0.02

model, where both allow for point masses at 0 and 1. The bootstrap confidence bounds are quite tight in this example, but they do not contain the true model on most of the domain. This shows that these classical bounds are overly confident if the fitted model is misspecified. The robust bounds, on the other hand, are wider and they do contain the true model everywhere. We note that the square-root bounds may go outside the triangle of admissible Pickands’ functions, see first row in Figure 3. This, however, can be easily fixed by simply restricting the bounds to stay inside the triangle.

In experiment #2 we fit a symmetric Hüsler–Reiss model to a symmetric logistic model ($b_1 = b_2 = 1$) with no point masses. Here the dominating measure is $\mu = \text{Leb}$ for the reasons that we discuss below. Even though both models are symmetric, the Hüsler–Reiss family is not flexible enough to well-approximate the generating logistic distribution. This is underlined by the fact that the Pickands’ function does not stay inside the bootstrap bounds, but only inside the wider robust bounds.

In the first two experiments we simulate $n = 20\,000$ data points, corresponding roughly to 55 years of daily observations. We choose r to be the 97.5% quantile of all radii, and use for fitting the $k = 500$ observations whose radii exceed r . In experiment #3 we only have $n = 2000$ data points and still use $k = 500$, corresponding to the 75% quantile for r . Comparing the histograms #2 and #3 in Figure 3, we note that in the latter case there are less observations close to 0 and 1. This illustrates that the data used for fitting comes from a pre-limit distribution. The δ we are estimating therefore represents the divergence of the data, that is, the pre-limit distribution, from the fitted model. This number can be considerably larger or smaller than δ_{true} , the divergence of the generating logistic distribution from the fitted model. In this case the robust theory still works well, but the estimated δ becomes unreliable. In experiment #3 we chose a run with similar divergences δ and δ_{true} .

Experiment #4 shows the case of fitting the well-specified asymmetric logistic family to the data. As expected, both the bootstrap and the robust bounds contain the true model. It is interesting to observe that both types of bounds almost coincide, meaning that the robust version is not overly conservative in the well-specified case.

Let us briefly discuss the choice of a dominating measure μ . We use $\mu = \mathbb{P}$, i.e., the classical Rényi divergence as defined in (2.2), whenever possible, that is, whenever δ_{true} is finite. This is the case when fitting extremal- t in experiments #1 and #3 and asymmetric logistic in #4, but not in experiment #2. Even though the true symmetric logistic density with no point masses is absolutely continuous with respect to the fitted Hüsler–Reiss density, the Rényi divergence of the former from the latter is infinite, because the Hüsler–Reiss density decays faster than any power

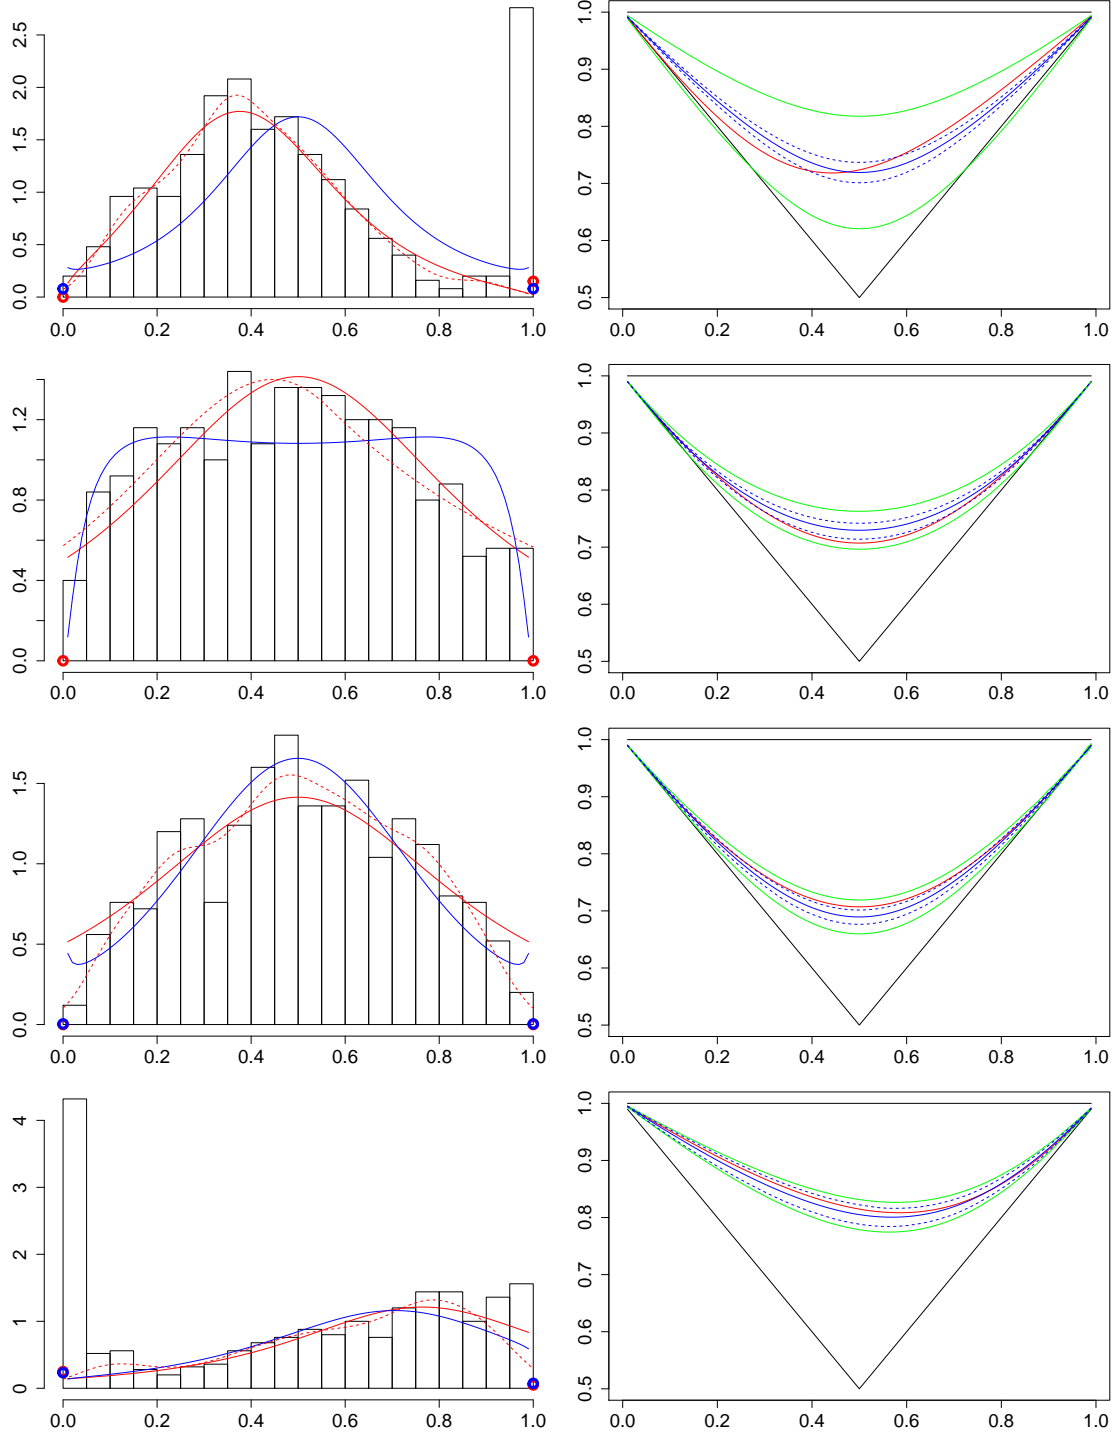


Figure 3: Left column: histogram of approximate Y , the true spectral distribution (solid red), the fitted distribution (blue), non-parametric density estimate used to compute divergence (dashed red). Right column: $A_{\text{true}}(z)$ (red), fitted $A(z)$ (solid blue) with its bootstrap bounds (dashed blue), and the robust square-root bounds (green). The rows correspond to experiments #1–4 described in Table 1, respectively.

at 0. Therefore, we take $\mu = \text{Leb}$ as a dominating measure in this case. Notice however, that the Lebesgue dominating measure does not allow for point masses which is desired in the other experiments.

Our experiments show that the easily computable robust square-root bounds can be applied effectively to measure uncertainty related to misspecified dependence structures in multivariate extremes. These readily available bounds are often exact, or very close to the exact robust bounds, see also Section 5.1. Thus the more challenging computation of the exact bounds is usually not required. Let us note that estimation of δ can be subtle, but it can be improved by an adequate choice of the dominating measure μ . Another important problem concerns reliable estimation of δ when data is coming from a pre-limit distribution. We leave these statistical questions for future research.

6 Robust bounds on the Value-at-Risk of a portfolio

In recent years diversification effects in heavy-tailed portfolios received considerable attention; see Mainik and Embrechts (2013); Mainik and Rüschendorf (2010); Zhou (2010). Suppose that \mathbf{Z} is a d -dimensional vector of dependent risk factors, and consider the portfolio $P = \sum_{i=1}^d w_i Z_i$, where $w_1, \dots, w_d \geq 0$ are some non-negative weights, not all being 0; one may assume that $\sum_i w_i = 1$ but this is not required. In order to have comparable risks one assumes that \mathbf{Z} is multivariate regularly varying with some index $\alpha > 0$ and non-degenerate exponent measure ν , so that (2.6) holds true.

Let $\text{VaR}_i(p)$ be the *Value-at-Risk* of the i th component, i.e., it satisfies

$$\bar{F}_i(\text{VaR}_i(p)) = p,$$

where $p > 0$ is a number close to 0. It follows that $\text{VaR}_i(p)$ is regularly varying at 0 with index $-1/\alpha$, and moreover from (2.6) we find

$$\frac{\text{VaR}_i(p)}{\text{VaR}_1(p)} \rightarrow m_i^{1/\alpha} \quad \text{as } p \downarrow 0. \quad (6.1)$$

In the following we consider the Value-at-Risk $\text{VaR}_P(p)$ of the portfolio P and provide the corresponding asymptotic robust bounds.

As discussed in Section 2.2, it is a common procedure to first estimate the marginal tails and then to address tail dependence, comparable to the copula concept in multivariate modeling. In this work we focus on the latter, more difficult task, and so we assume that the marginal survival functions \bar{F}_i are correctly specified. Transforming the marginals to unit Pareto

$$\hat{Z}_i = \frac{1}{\bar{F}_i(Z_i)},$$

we obtain normalized multivariate regularly varying $\hat{\mathbf{Z}}$, to which we associate $\hat{\mathbf{Y}} \in \mathbb{S}^{d-1}$ having the corresponding spectral distribution.

According to (Zhou, 2010, Thm. 3.1) the Value-at-Risk $\text{VaR}_P(p)$ of the portfolio P satisfies

$$\left(\frac{\text{VaR}_P(p)}{\text{VaR}_1(p)}\right)^\alpha \rightarrow d\mathbb{E}\left(\sum_{i=1}^d w_i(m_i\hat{Y}_i)^{1/\alpha}\right)^\alpha \quad \text{as } p \downarrow 0. \quad (6.2)$$

That is, the Value-at-Risk of the portfolio is asymptotically equivalent to the Value-of-Risk of every individual risk factor up to a diversification constant, which is easily identified from (6.2) and (6.1). In particular, letting

$$X = \left(\sum_{i=1}^d w_i(m_i\hat{Y}_i)^{1/\alpha}\right)^\alpha \quad (6.3)$$

we have the approximation for small $p > 0$

$$\text{VaR}_P(p) \approx \text{VaR}_1(p)(d\mathbb{E}X)^{1/\alpha}.$$

Suppose now that our model \mathbb{P} for the extremal dependence between the risk factors, i.e., for the spectral distribution $\hat{\mathbf{Y}}$, is subject to model uncertainty. A prominent problem in the financial literature on model uncertainty is to obtain worst case bounds on the Value-at-Risk of a portfolio (cf., Embrechts et al., 2013, 2015). In this regard we may directly apply our results from Sections 3 and 4 by considering the optimization problem (3.1) with X given in (6.3) and the moment constraints $\mathbb{E}\hat{Y}_i = 1/d$. Let us note again that there are essentially $d - 1$ constraints since $\hat{Y}_d = 1 - \sum_{i=1}^{d-1} \hat{Y}_i$. For fixed uncertainty radius $\delta > 0$, Theorem 1 yields the desired exact worst case bounds on $\text{VaR}_P(p)$ which are found numerically by solving a system of $d + 1$ non-linear equations. In higher dimension, solving these equations might be computationally challenging. On the other hand, the upper and lower square-root bounds in Theorem 2 and Corollary 1 coincide with the exact bounds for $\delta < \delta^*$ and $\delta < \delta_*$, respectively, and are otherwise very good approximations. Moreover, they can be easily computed even in higher dimensions. Indeed we only need to evaluate the covariance matrix $\Sigma_\mu(X, \hat{\mathbf{Y}})$ with respect to the chosen dominating measure μ . In the default case $\mu = \mathbb{P}$ this can be done, for instance, by Monte Carlo methods based on independent samples from $\hat{\mathbf{Y}}$. An algorithm for exact and efficient simulation of $\hat{\mathbf{Y}}$ can be found in Dombry et al. (2016).

Differently to Zhou (2010) and the above discussion, in Mainik and Embrechts (2013) the asymptotic relation between $\text{VaR}_P(p)$ and $\text{VaR}_1(p)$ is expressed using the spectral distribution of the original non-standardized \mathbf{Z} . This approach avoids separating the problem into marginal tail estimation and estimation of the tail dependence structure, which may be beneficial in some situations. Nevertheless, it does neither allow to use standard models for the spectral measure, nor to use the robust bounds presented in this paper. To be more explicit, let $\mathbf{Y} \in \mathbb{S}^{d-1}$ have the spectral distribution corresponding to \mathbf{Z} . Then according to (Mainik and Embrechts, 2013, Cor. 2.3) it holds that

$$\left(\frac{\text{VaR}_P(p)}{\text{VaR}_1(p)}\right)^\alpha \rightarrow \mathbb{E}\left(\sum_{i=1}^d w_i Y_i\right)^\alpha / \mathbb{E}Y_1^\alpha \quad \text{as } p \downarrow 0,$$

where \mathbf{Y} is not required to satisfy any moment constraints. It is now clear that optimization of the right hand side does not fit into our framework.

Acknowledgments

We are grateful to Enkelejd Hashorva for useful remarks, and to two anonymous referees for in-depth reading of this manuscript and for various comments that substantially improved the paper. Financial support by the Swiss National Science Foundation Project 161297 (first author) and the T.N. Thiele Center at Aarhus University (second author) is gratefully acknowledged.

Appendix

A.1 Some parametric families of spectral distributions

In the following we provide several commonly used parametric models for the spectral distribution H of $(Y_1, Y_2) \in \mathbb{S}^1$ in the bivariate case under L_1 -norm. Without loss of generality, we restrict our attention to the first component $Y = Y_1$, so that H is a probability measure on $[0, 1]$ equipped with its Borel σ -algebra. The following formulas are known but not readily available in the literature, and so we present them here for completeness.

A.1.1 Hüsler–Reiss

If the max-stable distribution is a bivariate Hüsler–Reiss distribution with dependence parameter $\lambda \in (0, \infty)$, then the density of the corresponding spectral distribution is

$$h_\lambda(\omega) = \frac{1}{\omega^2(1-\omega)4\lambda\sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \left(\lambda + \frac{\log \frac{1-\omega}{\omega}}{2\lambda} \right)^2 \right\}, \quad \omega \in [0, 1].$$

It is easy to check that this distribution is symmetric around $1/2$ and that $\mathbb{E}Y = 1/2$. Moreover, $A(1/2) = \Phi(\lambda)$, where Φ is the standard normal distribution function.

A.1.2 Asymmetric logistic

If the max-stable distribution is an asymmetric logistic distribution with dependence parameter $a \in (0, 1)$ and asymmetry parameters $b_1, b_2 \in [0, 1]$, then the corresponding spectral distribution has point masses $\mathbb{P}(Y = 0) = (1 - b_2)/2$ and $\mathbb{P}(Y = 1) = (1 - b_1)/2$, and the density is

$$h_{a,b_1,b_2}(\omega) = \frac{1-a}{2a} \frac{(b_1 b_2)^{1/a}}{(\omega(1-\omega))^{1+1/a}} \left\{ \left(\frac{b_1}{\omega} \right)^{1/a} + \left(\frac{b_2}{1-\omega} \right)^{1/a} \right\}^{a-2}, \quad \omega \in (0, 1). \quad (\text{A.4})$$

A.1.3 Extremal- t

If the max-stable distribution is an extremal- t distribution with parameters $\rho \in [-1, 1]$ and $a > 0$, then the distribution of the corresponding spectral distribution

has point masses

$$\mathbb{P}(Y = 0) = \mathbb{P}(Y = 1) = 1 - F_{a+1} \left\{ \rho \sqrt{\frac{a+1}{1-\rho^2}} \right\},$$

and the density for $\omega \in (0, 1)$ is

$$h_{\rho,a}(\omega) = \frac{(1-\rho^2)^{\frac{a+1}{2}} \Gamma(\frac{a+2}{2})}{2a\sqrt{\pi}\Gamma(\frac{a+1}{2})} (\omega(1-\omega))^{1/a-1} \\ \times \left\{ \omega^{2/a} - 2\rho(\omega(1-\omega))^{1/a} + (1-\omega)^{2/a} \right\}^{-\frac{a+2}{2}}.$$

Here, F_a is the t -distribution function with $a > 0$ degrees of freedom, i.e.,

$$F_a(x) = \frac{\Gamma(\frac{a+1}{2})}{\sqrt{a\pi}\Gamma(\frac{a}{2})} \int_{-\infty}^x \left(1 + \frac{t^2}{a} \right)^{-\frac{a+1}{2}} dt, \quad x \in \mathbb{R}.$$

A.2 Degenerate optimizers for the Pickands' function

In this section we study the degenerate case (ii) of Theorem 1 for the Pickands' dependence function $A(z)$ defined in (2.8), see also Section 5. That is, we assume that

$$X = X(z) = 2\{(1-z)Y \vee z(1-Y)\}, \quad Y \in [0, 1] \quad (\text{A.5})$$

for some fixed $z \in [0, 1]$, and that the constraint is $\mathbb{E}Y = 1/2$. Our aim is to identify the corresponding optimal values and the thresholds δ^{**} and δ_{**} , see Remark 2. The following result shows that the degenerate case corresponds to the trivial bounds $z \vee (1-z) \leq A(z) \leq 1$ as expected, but a certain assumption on the μ -support of Y is necessary.

Lemma 1. *Consider (A.5) and assume that μ -support of Y contains $0, z, 1$. Then the case (ii) of Theorem 1 occurs if and only if there exists a Radon–Nikodym derivative L^* (with respect to μ) such that*

$$\mathbb{E}_\mu(L^* - L)^2 \leq \delta, \quad \mathbb{P}^*(Y = 0) = \mathbb{P}^*(Y = 1) = 1/2,$$

in which case $\mathbb{E}^*X = 1$.

In case of the minimization problem the corresponding requirement on a Radon–Nikodym derivative L_ is*

$$\mathbb{E}_\mu(L_* - L)^2 \leq \delta, \quad \begin{cases} \mathbb{P}_*(Y \geq z) = 1, & z < 1/2, \\ \mathbb{P}_*(Y = 1/2) = 1, & z = 1/2, \\ \mathbb{P}_*(Y \leq z) = 1, & z > 1/2, \end{cases} \quad \mathbb{E}_*Y = 1/2,$$

in which case $\mathbb{E}_*X = z \vee (1-z)$.

Proof. Observe that the maximum of $X + cY$ is obtained for $Y = 1$ or $Y = 0$ or both (draw a picture). Since $\mathbb{E}^*Y = 1/2$ we must have $\mathbb{P}^*(Y = 0) = \mathbb{P}^*(Y = 1) = 1/2$, which yields the result.

The minimum of $X + cY$ is obtained either at $Y \leq z$ or at $Y \geq z$ or at the single points $0, z, 1$ (z is the bending point). Again the constraint $\mathbb{E}_*Y = 1/2$ leads to the result. The corresponding optimal value is $2(1-z)\mathbb{E}_*Y = 1-z$ when $z \leq 1/2$, and it is z when $z > 1/2$. \square

For the maximization problem, in the case of $\mu = \text{Leb}$, it is impossible to have $\mathbb{P}^*(Y = 0), \mathbb{P}^*(Y = 1) > 0$ and so according to Lemma 1 there cannot be a degenerate maximizer for any δ , i.e. $\delta^{**} = \infty$. In the case of $\mu = \mathbb{P}$ we have the following:

$$\delta^{**} = \frac{1}{4p_0} + \frac{1}{4p_1} - 1, \quad (\text{A.6})$$

where $p_i = \mathbb{P}(Y = i)$, and in particular p_0 and p_1 must be positive to have $\delta^{**} < \infty$. This follows from Lemma 1 and the following arguments. Note that

$$\mathbb{E}L^{*2} = \mathbb{E}(L^{*2}|Y = 0)p_0 + \mathbb{E}(L^{*2}|Y = 1)p_1 \geq l_0^2 p_0 + l_1^2 p_1,$$

where $l_i = \mathbb{E}(L^*|Y = i)$ and so $l_0 p_0 = l_1 p_1 = 1/2$. Hence, $L^* = l_0 \mathbf{1}_{\{Y=0\}} + l_1 \mathbf{1}_{\{Y=1\}}$ guarantees the minimal value for $\mathbb{E}L^{*2}$ among the allowed ones for any fixed δ . Therefore, a sufficient and necessary condition for existence of a degenerate maximizer is $l_0^2 p_0 + l_1^2 p_1 - 1 \leq \delta$, which readily yields (A.6).

For the minimization problem, the case of $z = 1/2$ is easy, $\delta_{**} = \infty$ for $\mu = \text{Leb}$, and $\delta_{**} = 1/\mathbb{P}(Y = 1/2) - 1$ for $\mu = \mathbb{P}$. For $z \neq 1/2$ the value of δ_{**} depends on the distribution of Y on $Y \geq z$ if $z < 1/2$ (on $Y \leq z$ if $z > 1/2$). More precisely, we need to identify a Radon–Nikodym derivative L_{**} which assigns no mass to $Y < z$, satisfies the constraints $\mathbb{E}_\mu(L_{**}) = 1, \mathbb{E}_\mu(L_{**}Y) = 1/2$ and minimizes $\mathbb{E}_\mu(L_{**} - L)^2$. This optimization problem is solved by $L_{**} = (b + cY + L)_+ \mathbf{1}_{\{Y \geq z\}}$ for $b, c \in \mathbb{R}$ such that the constraints hold. Finally, the minimal value $\mathbb{E}_\mu(L_{**} - L)^2$ is our δ_{**} .

A.3 Optimization for Rényi and Kullback–Leibler divergences

For completeness, we consider our optimization problem (3.1) for some other popular divergences: Rényi divergence of order $\eta > 1$ given by

$$\widehat{D}_\eta(\mathbb{P}', \mathbb{P}) = \frac{1}{\eta - 1} \log \mathbb{E}L'^\eta,$$

and Kullback–Leibler divergence given by

$$\widehat{D}_1(\mathbb{P}', \mathbb{P}) = \mathbb{E}(L' \log L'),$$

where it is assumed that $\mathbb{P}' \ll \mathbb{P}$ with $L' = d\mathbb{P}'/d\mathbb{P}$ and that the dominating measure μ coincides with \mathbb{P} . An easy adaptation of the proof of Theorem 1 shows that a maximizer \mathbb{P}^* of

$$\sup_{\mathbb{P}'} \{ \mathbb{E}' X : \widehat{D}_\eta(\mathbb{P}', \mathbb{P}) \leq \delta, \mathbb{E}' \mathbf{Y} = \mathbb{E} \mathbf{Y} \}$$

must have a Radon–Nikodym derivative $L^* \geq 0$ which satisfies $\mathbb{E}L^* = 1, \mathbb{E}(L^* \mathbf{Y}) = \mathbb{E} \mathbf{Y}$ and one of the following:

(i) $\widehat{D}_\eta(\mathbb{P}^*, \mathbb{P}) = \delta$ and there exist $a > 0, b, c_i \in \mathbb{R}$ such that

$$L^* = (aX + b + \mathbf{c}^\top \mathbf{Y})_+^{1/(\eta-1)} \quad \text{a.s.} \quad \text{when } \eta > 1, \quad (\text{A.7})$$

$$L^* = \exp(aX + b + \mathbf{c}^\top \mathbf{Y}) \quad \text{a.s.} \quad \text{when } \eta = 1, \quad (\text{A.8})$$

- (ii) there exist c_i such that the distribution of $X + \mathbf{c}^\top \mathbf{Y}$ has a positive mass at its right end, $L^* = 0$ everywhere else a.s., and the constraint $\hat{D}_\eta(\mathbb{P}^*, \mathbb{P}) \leq \delta$ holds.

Conversely, any such L^* corresponds to a maximizer \mathbb{P}^* . In the case $\eta > 1$ we assume that $\mathbb{E}|X|^{\eta/(\eta-1)}, \mathbb{E}|Y|^{\eta/(\eta-1)} < \infty$, and in the case $\eta = 1$ we assume that $\mathbb{E}(|X|L'), \mathbb{E}(|Y_i|L') < \infty$ for all L' satisfy $\mathbb{E}(L' \log L') \leq \delta$. Furthermore, regardless of these assumptions, if there exists L^* as above and such that $\mathbb{E}(|X|L^*), \mathbb{E}(|Y_i|L^*) < \infty$ then it must be a maximizer, which can be seen by considering an appropriate convex subset of L' in the proof of Theorem 1.

Note that taking $\eta = 2$ we retrieve the result of Theorem 1 for $\mu = \mathbb{P}$. In the case $d = 1$ (no moment constraints) the expression for L^* in (i) appears in e.g. Blanchet and Murthy (2016). Furthermore, Breuer and Csiszár (2013a) considers more general divergences but the results are less explicit. Finally, we elaborate on the case of Kullback–Leibler divergence extending the result of Ahmadi-Javid (2012) by introducing moment constraints.

Proposition 1. *Assume that $X, Y_i \geq 0$ are positive random variables with finite expectation, and $G(a, \mathbf{c}) = \mathbb{E}e^{aX + \sum_i c_i Y_i}$ is finite on some domain $\mathcal{D} \subset (0, \infty) \times \mathbb{R}^d$ with non-empty interior. Suppose there exist $(a, \mathbf{c}) \in \mathcal{D}$ such that*

$$\frac{G_i(a, \mathbf{c})}{G(a, \mathbf{c})} = \mathbb{E}Y_i, \quad a \frac{G_0(a, \mathbf{c})}{G(a, \mathbf{c})} + \sum_{i=1}^d c_i \mathbb{E}Y_i - \log G(a, \mathbf{c}) = \delta,$$

where $G_i(\cdot)$ is a derivative with respect to the i th variable (pointing inside the domain if on the boundary). Then

$$V^{KL}(\delta) = \sup_{\mathbb{P}'} \{ \mathbb{E}' X : \hat{D}_1(\mathbb{P}', \mathbb{P}) \leq \delta, \mathbb{E}' \mathbf{Y} = \mathbb{E} \mathbf{Y} \} = \frac{G_0(a, \mathbf{c})}{G(a, \mathbf{c})},$$

which corresponds to the exponential change of measure $L^* = e^{aX + \sum_i c_i Y_i} / G(a, \mathbf{c})$.

Proof. According to (A.8) we consider

$$L^* = \exp(aX + b + \sum_i c_i Y_i) =: \exp(U), \quad a > 0, b, c_i \in \mathbb{R}$$

together with the constraints: $\mathbb{E}e^U = 1, \mathbb{E}(Ue^U) = \delta, \mathbb{E}(Y_i e^U) = \mathbb{E}Y_i$. We may rewrite these using the moment generating function G :

$$\begin{aligned} G(a, \mathbf{c}) &= e^{-b}, \\ aG_0(a, \mathbf{c}) + bG(a, \mathbf{c}) + \sum_i c_i G_i(a, \mathbf{c}) &= \delta e^{-b}, \\ e^b G_i(a, \mathbf{c}) &= \mathbb{E}Y_i, \\ V^{KL}(\delta) &= e^b G_0(a, \mathbf{c}). \end{aligned}$$

The equations in the statement are now immediate. Finally, we note that $\mathbb{E}(XL^*)$ and $\mathbb{E}(Y_i L^*)$ are finite which completes the proof. \square

References

- A. Ahmadi-Javid. Entropic value-at-risk: A new coherent risk measure. *Journal of Optimization Theory and Applications*, 155(3):1105–1123, 2012.
- P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.
- S. Asmussen and P. W. Glynn. *Stochastic simulation: algorithms and analysis*, volume 57 of *Stochastic Modelling and Applied Probability*. Springer, New York, 2007.
- J. Blanchet and K. R. Murthy. On distributionally robust extreme value analysis. *arXiv preprint arXiv:1601.06858*, 2016.
- M.-O. Boldi and A. C. Davison. A mixture model for multivariate extremes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 69:217–229, 2007.
- T. Breuer and I. Csiszár. Measuring distribution model risk. *Mathematical Finance*, 2013a.
- T. Breuer and I. Csiszár. Systematic stress tests with entropic plausibility constraints. *Journal of Banking & Finance*, 37(5):1552–1559, 2013b.
- A. Bücher, H. Dette, and S. Volgushev. New estimators of the Pickands dependence function and a test for extreme-value dependence. *Ann. Statist.*, 39(4):1963–2006, 2011. ISSN 0090-5364.
- D. Cooley, R. A. Davis, and P. Naveau. The pairwise beta distribution: a flexible parametric multivariate model for extremes. *J. Multivariate Anal.*, 101:2103–2117, 2010.
- L. de Haan and A. Ferreira. *Extreme Value Theory*. Springer, New York, 2006.
- S. Dey and S. Juneja. Incorporating fat tails in financial models using entropic divergence measures. *arXiv preprint arXiv:1203.0643*, 2012.
- C. Dombry, S. Engelke, and M. Oesting. Exact simulation of max-stable processes. *Biometrika*, 103:303–317, 2016.
- J. H. J. Einmahl and J. Segers. Maximum empirical likelihood estimation of the spectral measure of an extreme-value distribution. *Ann. Statist.*, 37(5B):2953–2989, 2009.
- P. Embrechts, C. Klüppelberg, and T. Mikosch. *Modelling Extremal Events: for Insurance and Finance*. Springer, London, 1997.
- P. Embrechts, G. Puccetti, and L. Rüschendorf. Model uncertainty and var aggregation. *Journal of Banking & Finance*, 37(8):2750–2764, 2013.
- P. Embrechts, B. Wang, and R. Wang. Aggregation-robustness and model uncertainty of regulatory risk measures. *Finance and Stochastics*, 19(4):763–790, 2015.
- P. Glasserman and X. Xu. Robust risk measurement and model risk. *Quantitative Finance*, 14(1):29–58, 2014.
- L. P. Hansen and T. J. Sargent. Robust control and model uncertainty. *The American Economic Review*, 91(2):60–66, 2001.

- J. Hüsler and R.-D. Reiss. Maxima of normal random vectors: between independence and complete dependence. *Statist. Probab. Lett.*, 7:283–286, 1989.
- C. Klüppelberg and A. May. Bivariate extreme value distributions based on polynomial dependence functions. *Mathematical Methods in the Applied Sciences*, 29:1467–1480, 2006.
- G. Mainik and P. Embrechts. Diversification in heavy-tailed portfolios: properties and pitfalls. *Annals of Actuarial Science*, 7(01):26–45, 2013.
- G. Mainik and L. Rüschendorf. On optimal portfolio diversification with respect to extreme risks. *Finance Stoch.*, 14(4):593–623, 2010.
- S. K. Mitter. Convex optimization in infinite dimensional spaces. In *Recent advances in learning and control*, pages 161–179. Springer, 2008.
- D. V. Ouellette. Schur complements and statistics. *Linear Algebra Appl.*, 36:187–295, 1981. ISSN 0024-3795.
- J. Pickands, III. Multivariate extreme value distributions. In *Proceedings of the 43rd session of the International Statistical Institute, Vol. 2 (Buenos Aires, 1981)*, volume 49, pages 859–878, 894–902, 1981. With a discussion.
- B. Póczos and J. G. Schneider. On the estimation of α -divergences. In *AISTATS*, pages 609–617, 2011.
- S. I. Resnick. *Heavy-tail phenomena*. Springer Series in Operations Research and Financial Engineering. Springer, New York, 2007. Probabilistic and statistical modeling.
- S. I. Resnick. *Extreme Values, Regular Variation and Point Processes*. Springer, New York, 2008.
- M. Schlather and J. A. Tawn. A dependence measure for multivariate and spatial extreme values: Properties and inference. *Biometrika*, 90:139–156, 2003.
- J. Schur. Über potenzreihen, die im innern des einheitskreises beschränkt sind. *Journal für die reine und angewandte Mathematik*, 147:205–232, 1917. URL <http://eudml.org/doc/149467>.
- A. G. Stephenson. evd: Extreme value distributions. *R News*, 2(2):0, June 2002. URL <http://CRAN.R-project.org/doc/Rnews/>.
- J. A. Tawn. Modelling multivariate extreme value distributions. *Biometrika*, 77:245–253, 1990.
- D. Williams. *Probability with martingales*. Cambridge Mathematical Textbooks. Cambridge University Press, Cambridge, 1991.
- C. Zhou. Dependence structure of risk factors and diversification effects. *Insurance Math. Econom.*, 46(3):531–540, 2010.