



HAL
open science

Beyond tail median and conditional tail expectation: extreme risk estimation using tail L_p –optimisation

Laurent Gardes, Stéphane Girard, Gilles Stupfler

► To cite this version:

Laurent Gardes, Stéphane Girard, Gilles Stupfler. Beyond tail median and conditional tail expectation: extreme risk estimation using tail L_p –optimisation. 2018. hal-01726328v1

HAL Id: hal-01726328

<https://inria.hal.science/hal-01726328v1>

Preprint submitted on 8 Mar 2018 (v1), last revised 7 Nov 2019 (v4)

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.

Beyond tail median and conditional tail expectation: extreme risk estimation using tail L^p -optimisation

Laurent Gardes⁽¹⁾, Stéphane Girard⁽²⁾ & Gilles Stupfler⁽³⁾

⁽¹⁾ Université de Strasbourg & CNRS, IRMA, UMR 7501, 7 rue René Descartes,
67084 Strasbourg Cedex, France

⁽²⁾ Université Grenoble Alpes, INRIA, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

⁽³⁾ School of Mathematical Sciences, University of Nottingham, University Park,
Nottingham NG7 2RD, United Kingdom

Abstract. The Conditional Tail Expectation is an indicator of tail behaviour that has recently gained traction in actuarial and financial applications. Contrary to the quantile or Value-at-Risk, it takes into account the frequency of a tail event together with the probabilistic behaviour of the variable of interest on this event. However, the asymptotic normality of the empirical Conditional Tail Expectation estimator requires that the underlying distribution possess a finite variance; this can be a strong restriction in heavy-tailed models which constitute the favoured class of models in actuarial and financial applications. One possible solution in very heavy-tailed models where this assumption fails could be to use the more robust Median Shortfall, but this quantity is actually just a quantile, which therefore only gives information about the frequency of a tail event and not about its typical magnitude. We construct a synthetic class of tail L^p -medians, which encompasses the Median Shortfall (for $p = 1$) and Conditional Tail Expectation (for $p = 2$). We show that, for $1 < p < 2$, a tail L^p -median always takes into account both the frequency and magnitude of tail events, and its empirical estimator is, within the range of the data, asymptotically normal under a condition weaker than a finite variance. We extrapolate this estimator, along with another technique, to proper extreme levels using the heavy-tailed framework. The estimators are showcased on a simulation study and on a set of real fire insurance data showing evidence of a very heavy right tail.

AMS Subject Classifications: 62G05, 62G30, 62G32.

Keywords: asymptotic normality, Conditional Tail Expectation, extreme value statistics, heavy-tailed distribution, L^p -optimisation, Median Shortfall, semiparametric extrapolation.

1 Introduction

A precise assessment of extreme risk is a crucial question in a number of fields of statistical applications. In actuarial science and finance, especially, a major question is to get a good understanding of potential extreme claims and losses that would constitute a threat to the survival of a company. This has historically been done by simply using a quantile, also called Value-at-Risk (VaR) in the actuarial and financial world. A quantile of a real-valued random variable X at level $\alpha \in (0, 1)$ is, by definition, the lowest value which is exceeded by X with probability not larger than $1 - \alpha$. As a consequence, quantiles only provide an information about the frequency of a tail event; in particular, the quantile $q(\alpha)$ gives no indication about how heavy the right tail of X is or, perhaps more concretely, as to what a typical value of X above $q(\alpha)$ would be. From a risk assessment perspective, this is clearly an issue, as quantiles allow to quantify what a risky situation is (*i.e.* the quantile level), but not what the consequences of a risky situation would be (*i.e.* the behaviour of X beyond the quantile level).

The notion of Conditional Tail Expectation (CTE) precisely addresses this point. This risk measure, defined as $\text{CTE}(\alpha) := \mathbb{E}[X | X > q(\alpha)]$, is exactly the average value of X given that it is larger than the quantile $q(\alpha)$. When X is continuous, $\text{CTE}(\alpha)$ coincides with the so-called Expected Shortfall $\text{ES}(\alpha)$, also known as Tail-Value-at-Risk or Conditional Value-at-Risk and discussed in [2, 48, 50], which is nothing but the average value of the quantile function $\tau \mapsto q(\tau)$ over the interval $[\alpha, 1)$. The potential of CTE for use in actuarial and financial risk management has been considered by a number of studies, such as [10, 25, 51, 56]. Outside academic contexts, the CTE risk measure is used in capital requirement calculations by the Canadian financial and actuarial sectors [36], as well as for guaranteeing the sustainability of life insurance annuities in the USA [45]. European regulators, *via* the Basel Committee on Banking Supervision, also recently recommended to use CTE rather than VaR in internal market risk models, see [5]. Whether the CTE or VaR should be used in a given financial application is still very much up for debate: we point to, among others, [4, 16, 24, 28, 42] who discuss the various intrinsic axiomatic properties of CTE and/or VaR and their practical interpretations.

Aside from axiomatic considerations, the computation of the Conditional Tail Expectation of X requires of course that the first moment $\mathbb{E}|X| < \infty$ be finite. Moreover, its estimation, by the empirical conditional tail moment, requires that $\mathbb{E}(X^2) < \infty$ if this estimator is to be asymptotically normal. This should of course be expected, since a condition for the sample average to be an asymptotically normal estimator of the sample mean is precisely the finiteness of the variance. In heavy-tailed models, where (right) tail heaviness of X is controlled by the so-called tail index γ , this translates into assuming that $\gamma < 1/2$ (see Theorem 1 in [20] and Theorem 2 in [22]). At the finite-sample level, it is often desirable to have a finite fourth moment, or $\gamma < 1/4$; Table 1 in the simulation study of [22] shows a quite strong deterioration in the finite-sample performance

of the estimator as γ increases to $1/4$. The finite variance condition, and even more so the finite fourth moment condition, can be restrictive in the heavy-tailed framework which is of interest in insurance/finance and that will be the focus of this paper. Several actuarial and financial data sets have been found to violate such integrability assumptions: we refer to, for instance, the Danish fire insurance data set considered in Example 4.2 of [47], emerging market stock returns data [32, 43], exchange rates data [33], as well as the R package “CASdatasets” which includes examples of actuarial and financial data sets where observed values of γ are found to be higher than $1/4$. The deterioration in finite-sample performance as the number of finite moments decreases is due in no small part to increased variability in the right tail of X as γ increases, see Section 3.1 in [23]. To counteract this variability, one could of course move away from the CTE and use the Median Shortfall MS (see [41, 42]), just as one can use the median instead of the mean for added robustness. However, $MS(\alpha)$ is nothing but $q([1 + \alpha]/2)$, and as such suffers from the same problem as the VaR does, in terms of not containing any information on the behaviour of X beyond $MS(\alpha)$.

Our goal here is to design a class of tail indicators that realise a compromise between the sensitivity of the Conditional Tail Expectation and the robustness of the Median Shortfall. We will start by showing that these two quantities can be obtained in a new unified framework, which we call the class of *tail L^p -medians*. A tail L^p -median at level $\alpha \in (0, 1)$ is obtained, for $p \geq 1$, by minimising an L^p -moment criterion which only considers the probabilistic behaviour of X above a quantile $q(\alpha)$. As such, it satisfies a number of interesting properties we will investigate. Most importantly, the class of tail L^p -medians is linked to the notion of L^p -quantiles introduced by [13] and recently studied in heavy-tailed frameworks by [14], a tail L^p -median being the L^p -quantile of order $1/2$ (or L^p -median) of the variable $X | X > q(\alpha)$. In particular, the classical median of X is the L^1 -median and the mean of X is the L^2 -median (or expectile of order $1/2$, in the terminology of [44]), and we will show that similarly, the Median Shortfall is the tail L^1 -median, while the Conditional Tail Expectation is the tail L^2 -median. Like L^p -quantiles, tail L^p -medians with $p > 1$ depend on both the frequency of the event $\{X > q(\alpha)\}$ and the actual behaviour of X beyond $q(\alpha)$. At the technical level, a condition for a tail L^p -median to exist is that $\gamma < 1/(p - 1)$, and it can be empirically estimated at high levels by an asymptotically Gaussian estimator if $\gamma < 1/[2(p - 1)]$. When $p \in (1, 2)$, the tail L^p -median and the L^p -quantile therefore simultaneously exist and are accurately estimable in a wider class of models than the CTE is.

However, there are a number of differences between tail L^p -medians and L^p -quantiles. For example, the theoretical analysis of the asymptotic behaviour of tail L^p -medians, as $\alpha \rightarrow 1$, is technically more complex than that of L^p -quantiles. The asymptotic results that arise show that the tail L^p -median at level α is asymptotically proportional to the quantile $q(\alpha)$, as $\alpha \rightarrow 1$, through a non-explicit but very accurately approximable constant, and the remainder term in the asymptotic relationship is exclusively controlled by extreme value parameters of X . This stands

in contrast with L^p -quantiles, which also are asymptotically proportional to the quantile $q(\alpha)$ through a simpler constant, although the remainder term crucially features the expectation and left-tail behaviour of X . The remainder term plays an important role in the estimation, as it determines the bias term in our eventual estimators, and we will then argue that the extreme value behaviour of tail L^p -medians is easier to understand and more natural than that of L^p -quantiles. We will also explain why, for heavy-tailed models, extreme tail L^p -medians are able to interpolate monotonically between extreme MS and extreme CTE, as p varies in $(1, 2)$ and for $\gamma < 1$. By contrast, L^p -quantiles are known not to interpolate monotonically between quantiles and expectiles (see *e.g.* Figure 1 in [14]). The interpolation property also makes it possible to interpret an extreme tail L^p -median as a weighted average between extreme MS and extreme CTE, and as such, give an interpretation of extreme tail L^p -medians within the framework of Wang distortion risk measures, introduced originally by [52]. The class of Wang distortion risk measures is a well-established and intuitive set of risk indicators that has been used in a number of financial and actuarial contexts, including the pricing of insurance premiums and bonds, as well as tackling capital allocation problems (see *e.g.* [7, 53, 54]). The interpretation of extreme tail L^p -medians in terms of Wang distortion risk measures is thus likely to be helpful as far as practical applicability of tail L^p -medians is concerned.

We shall then examine how to estimate an extreme tail L^p -median. In extreme value terminology, this amounts to estimating a tail L^p -median at a level $\alpha = \alpha_n$ such that $\alpha_n \rightarrow 1$ as the size n of the sample of data tends to infinity. We shall start by suggesting two estimation methods in the so-called intermediate case, when $n(1 - \alpha_n) \rightarrow \infty$. Although the design stage of our tail L^p -median estimators has similarities with that of the L^p -quantile estimators in [14], the investigation of the asymptotic properties of the tail L^p -median estimators is more challenging and technically involved. These methods will provide us with basic estimators that we will extrapolate at a proper extreme level α_n , which satisfies $n(1 - \alpha_n) \rightarrow c \in [0, \infty)$, using the heavy-tailed assumption. We will, as a theoretical byproduct, demonstrate how our results make it possible to recover or improve upon known results in the literature on Conditional Tail Expectation. Our final step will be to assess the finite-sample behaviour of the suggested estimators. Our focus will not be to consider heavy-tailed models with a small value of γ : for these models and more generally in models with finite fourth moment, it is unlikely that any improvement will be brought on the CTE, whether in terms of quality of estimation or interpretability. Our view is rather that the use of tail L^p -medians with $p \in (1, 2)$ will be beneficial for very heavy-tailed models, in which γ is higher than the finite fourth moment threshold $\gamma = 1/4$, and possibly higher than the finite variance threshold $\gamma = 1/2$. For such values of γ and with an extreme level set to be $1 - 1/n$ (a typical choice in applications, see *e.g.* recently [3, 11, 12, 29]), we shall then evaluate the finite-sample performance of our estimators on simulated data sets, as well as on a real set of fire insurance data from the R package “CASdatasets”, featuring an estimated value of γ larger than $1/2$.

The paper is organised as follows. We first give a rigorous definition of our concept of tail L^p -median and state some of its elementary properties in Section 2. Section 3 then focuses on the analysis of asymptotic properties of the population tail L^p -median, as $\alpha \rightarrow 1$, in heavy-tailed models. Estimators of an extreme tail L^p -median, obtained by first constructing two distinct estimators at intermediate levels which are then extrapolated to extreme levels, will be discussed in Section 4. A simulation study of the finite-sample performance of our estimators is presented in Section 5, and an application to real fire insurance data is discussed in Section 6. Proofs and auxiliary results are deferred to an online Supplementary Material document.

2 Definition and first properties

Let X be a real-valued random variable. Its distribution function, which is assumed to be continuous throughout, is denoted by $F(x) = \mathbb{P}(X \leq x)$ and the associated α -quantile is given for $\alpha \in (0, 1)$ by $q(\alpha) := \inf\{x \in \mathbb{R} \mid F(x) \geq \alpha\}$. Our construction is motivated by the following three properties. Firstly, the median of X can be obtained by minimising the expected absolute deviation (see [39, 40]):

$$q(1/2) = \arg \min_{m \in \mathbb{R}} \mathbb{E} (|X - m| - |X|).$$

Note that the minimiser on the right-hand side may not be unique (but it is if F is strictly increasing); our convention throughout this paper will be that in such a situation, the minimiser is taken as the smallest possible minimiser. It is readily checked that equality then holds in this case as well. Secondly, when X has a finite first moment, it follows from a straightforward calculation that the expectation of X can be obtained by minimising the expected squared deviation:

$$\mathbb{E}(X) = \arg \min_{m \in \mathbb{R}} \mathbb{E} (|X - m|^2 - |X|^2).$$

Thirdly and finally, since F is continuous, we have $F(q(\alpha)) = \alpha$, and therefore the Median Shortfall $\text{MS}(\alpha) = q([1 + \alpha]/2)$ is nothing but the median of X given $X > q(\alpha)$ (see Example 3 in [42]). As such

$$\text{MS}(\alpha) = \arg \min_{m \in \mathbb{R}} \mathbb{E} (|X - m| - |X| \mid X > q(\alpha))$$

and, because $\text{CTE}(\alpha)$ is the expectation of X given $X > q(\alpha)$, we have, if $\mathbb{E}(|X| \mid X > q(\alpha)) < \infty$,

$$\text{CTE}(\alpha) = \arg \min_{m \in \mathbb{R}} \mathbb{E} (|X - m|^2 - |X|^2 \mid X > q(\alpha)).$$

Our construction now encompasses these two quantities by replacing the absolute or squared deviations by power deviations; more precisely, for $p \geq 1$, we define a *tail L^p -median* of order $\alpha \in (0, 1)$ (when it exists) as

$$m_p(\alpha) = \arg \min_{m \in \mathbb{R}} \mathbb{E} (|X - m|^p - |X|^p \mid X > q(\alpha)). \tag{1}$$

Let us highlight the following important connection between $m_p(\alpha)$ and the notion of L^p -quantiles: recall, from [13, 14], that an L^p -quantile of order $\tau \in (0, 1)$ of a univariate random variable Y , with $\mathbb{E}|Y|^{p-1} < \infty$, is defined as

$$q_\tau(p) = \arg \min_{q \in \mathbb{R}} \mathbb{E} (|\tau - \mathbb{1}_{\{Y \leq q\}}| \cdot |Y - q|^p - |\tau - \mathbb{1}_{\{Y \leq 0\}}| \cdot |Y|^p).$$

Consequently, the L^p -median of Y , obtained for $\tau = 1/2$, is

$$q_{1/2}(p) = \arg \min_{q \in \mathbb{R}} \mathbb{E} (|Y - q|^p - |Y|^p).$$

For an arbitrary $p \geq 1$, the tail L^p -median $m_p(\alpha)$ of X is then exactly an L^p -median of X given that $X > q(\alpha)$. This construction of $m_p(\alpha)$ as an L^p -median given the tail event $\{X > q(\alpha)\}$ is what motivated the name ‘‘tail L^p -median’’ for $m_p(\alpha)$.

Note that subtracting $|X|^p$ inside the expectation in (1) above makes the cost function well-defined whenever $\mathbb{E}(|X|^{p-1} | X > q(\alpha))$ is finite. This is a straightforward consequence of the triangle inequality when $p = 1$; when $p > 1$, this is a consequence of the fact that the function $x \mapsto |x|^p$ is continuously differentiable with derivative $x \mapsto p|x|^{p-1} \text{sign}(x)$, together with the mean value theorem. If X is moreover assumed to satisfy $\mathbb{E}(|X|^p | X > q(\alpha)) < \infty$, then the definition of $m_p(\alpha)$ is simply

$$m_p(\alpha) = \arg \min_{m \in \mathbb{R}} \mathbb{E} (|X - m|^p | X > q(\alpha)).$$

The following result shows in particular that if $\mathbb{E}(|X|^{p-1} | X > q(\alpha)) < \infty$ for some $p \geq 1$, then the tail L^p -median always exists and is characterised by a simple equation. Especially, for $p \in (1, 2)$, the tail L^p -median $m_p(\alpha)$ exists and is unique under a weaker integrability condition than the assumption of a finite first tail moment which is necessary for the existence of $\text{CTE}(\alpha)$.

Proposition 1. *Let $p \geq 1$. Pick $\alpha \in (0, 1)$ and assume that $\mathbb{E}(|X|^{p-1} \mathbb{1}_{\{X > q(\alpha)\}}) < \infty$. Then:*

- (i) *The tail L^p -median $m_p(\alpha)$ exists and is such that $m_p(\alpha) > q(\alpha)$.*
- (ii) *The tail L^p -median is equivariant with respect to increasing affine transformations: if $m_p^X(\alpha)$ is the tail L^p -median of X and $Y = aX + b$ with $a > 0$ and $b \in \mathbb{R}$, then the tail L^p -median $m_p^Y(\alpha)$ of Y is $m_p^Y(\alpha) = a m_p^X(\alpha) + b$.*
- (iii) *When $p > 1$, the tail L^p -median $m_p(\alpha)$ is the unique $m \in \mathbb{R}$ solution of the equation*

$$\mathbb{E}[(m - X)^{p-1} \mathbb{1}_{\{q(\alpha) < X < m\}}] = \mathbb{E}[(X - m)^{p-1} \mathbb{1}_{\{X > m\}}].$$

- (iv) *The tail L^p -median defines a monotonic functional with respect to first-order stochastic dominance: if Y is another random variable with $\mathbb{E}(|Y|^{p-1} \mathbb{1}_{\{Y > q(\alpha)\}}) < \infty$ then*

$$(\forall t > 0, \mathbb{P}(X > t) \leq \mathbb{P}(Y > t)) \Rightarrow m_p^X(\alpha) \leq m_p^Y(\alpha).$$

(v) When $\mathbb{E}|X|^{p-1} < \infty$, the function $\alpha \mapsto m_p(\alpha)$ is nondecreasing on $(0, 1)$.

Let us highlight that, in addition to these properties, we have most importantly

$$\text{MS}(\alpha) = m_1(\alpha) \quad \text{and} \quad \text{CTE}(\alpha) = m_2(\alpha). \quad (2)$$

In other words, the Median Shortfall is the tail L^1 -median and the Conditional Tail Expectation is the tail L^2 -median: the class of tail L^p -medians encompasses the notions of MS and CTE.

3 Asymptotic properties of an extreme tail L^p -median

Our construction of a tail L^p -median $m_p(\alpha)$ is primarily intended to provide a middle way between the risk measures $\text{MS}(\alpha)$ and $\text{CTE}(\alpha)$, for a level α close to 1. This is of value in the cases where the number of finite moments of X is low (typically, less than 4) because the estimation of $\text{CTE}(\alpha)$, for high α , is then a difficult task in practice. It has been found in the literature that heavy-tailed distributions constitute appropriate models for the kind of extreme value behaviour observed in actuarial and financial data (see, *e.g.*, [19, p.9] and [47, p.1]). Our first focus will therefore be to analyse whether $m_p(\alpha)$, for $p \in (1, 2)$, does indeed provide such a middle way as $\alpha \rightarrow 1$, in the class of heavy-tailed distributions. This is done by introducing the following regular variation assumption on the survival function $\bar{F} := 1 - F$ of X .

$\mathcal{C}_1(\gamma)$ There exists $\gamma > 0$ such that for all $x > 0$,

$$\lim_{t \rightarrow \infty} \frac{\bar{F}(tx)}{\bar{F}(t)} = x^{-1/\gamma}.$$

In other words, condition $\mathcal{C}_1(\gamma)$ means that the function \bar{F} is regularly varying with index $-1/\gamma$ in a neighbourhood of $+\infty$, see [8]. Note that assuming this condition is equivalent to supposing that the tail quantile function $t \mapsto U(t) := q(1 - t^{-1})$ of X is regularly varying with the so-called tail index γ . We also remark that when $\mathcal{C}_1(\gamma)$ is satisfied then X does not have any finite conditional tail moments of order larger than $1/\gamma$; similarly, all conditional tail moments of X of order smaller than $1/\gamma$ are finite (a rigorous statement is Exercise 1.16 in [30]). Since the existence of a tail L^p -median requires finite conditional tail moments of order $p - 1$, this means that our minimal working condition on the pair (p, γ) should be $\gamma < 1/(p - 1)$, a condition that will indeed appear in many of our asymptotic results.

We now provide some insight into what asymptotic result on $m_p(\alpha)$ we should aim for under condition $\mathcal{C}_1(\gamma)$. A consequence of this condition is that above a high threshold u , the variable X/u is approximately Pareto distributed with tail index γ , or, in other words:

$$\forall x > 1, \quad \mathbb{P} \left(\frac{X}{q(\alpha)} > x \mid \frac{X}{q(\alpha)} > 1 \right) \rightarrow x^{-1/\gamma} \quad \text{as } \alpha \rightarrow 1.$$

This is exactly the first statement of Theorem 1.2.1 in [30]. The above conditional Pareto approximation then suggests that when α is close enough to 1, the optimisation criterion (1) can be approximately written as follows:

$$\frac{m_p(\alpha)}{q(\alpha)} \approx \arg \min_{M \in \mathbb{R}} \mathbb{E}[|Z_\gamma - M|^p - |Z_\gamma|^p]$$

where Z_γ has a Pareto distribution with tail index γ . For the variable Z_γ , we can differentiate the cost function and use change-of-variables formulae to get that the minimiser M of the right-hand side should satisfy the following equation:

$$\int_{1/M}^1 (1-u)^{p-1} u^{-1/\gamma-1} du = \int_0^1 v^{p-1} (1-v)^{1/\gamma-p} dv,$$

or equivalently

$$g_{p,\gamma}(1/M) = B(p, \gamma^{-1} - p + 1)$$

where $g_{p,\gamma}(t) := \int_t^1 (1-u)^{p-1} u^{-1/\gamma-1} du$ on $(0, 1)$ and, for $x, y > 0$, $B(x, y) = \int_0^1 v^{x-1} (1-v)^{y-1} dv$ denotes the Beta function. In other words, and for a Pareto random variable, we have

$$\frac{m_p(\alpha)}{q(\alpha)} = \frac{1}{\kappa(p, \gamma)} \quad \text{where } \kappa(p, \gamma) := g_{p,\gamma}^{-1}(B(p, \gamma^{-1} - p + 1))$$

and $g_{p,\gamma}^{-1}$ denotes the inverse of the decreasing function $g_{p,\gamma}$ on $(0, 1)$. Our first asymptotic result states that this proportionality relationship is still valid asymptotically under condition $\mathcal{C}_1(\gamma)$.

Proposition 2. *Suppose that $p \geq 1$ and $\mathcal{C}_1(\gamma)$ holds with $\gamma < 1/(p-1)$. Then:*

$$\lim_{\alpha \rightarrow 1} \frac{m_p(\alpha)}{q(\alpha)} = \frac{1}{\kappa(p, \gamma)}.$$

It follows from Proposition 2 that a tail L^p -median above a high exceedance level is approximately a multiple of this exceedance level. This first-order result is similar in spirit to other asymptotic proportionality relationships linking extreme risk measures to extreme quantiles: we refer to [15] for a result on extreme expectiles and [14] on the general class of L^p -quantiles, as well as to [22, 57, 58] for a similar analysis of extreme Wang distortion risk measures.

The asymptotic proportionality constant $\kappa(p, \gamma) \in (0, 1)$ does not have a simple closed form in general, due to the complicated expression of the function $g_{p,\gamma}$. It does however have a nice explicit expression in the two particular cases $p = 1$ and $p = 2$. For $p = 1$, we have $\kappa(1, \gamma) = 2^{-\gamma}$, see Lemma 10(ii) in the Supplementary Material document. This clearly yields the same equivalent as the one obtained using the regular variation of the tail quantile function U with index γ , *i.e.* in virtue of (2):

$$\frac{m_1(\alpha)}{q(\alpha)} = \frac{q((1+\alpha)/2)}{q(\alpha)} = \frac{U(2(1-\alpha)^{-1})}{U((1-\alpha)^{-1})} \rightarrow 2^\gamma \text{ as } \alpha \rightarrow 1.$$

When $p = 2$ and $\gamma \in (0, 1)$, $\kappa(2, \gamma) = 1 - \gamma$, see Lemma 11(ii) in the Supplementary Material document. Since in this case, $m_2(\alpha)$ is nothing but $\text{CTE}(\alpha)$ by (2), Proposition 2 agrees here with the asymptotic equivalent of $\text{CTE}(\alpha)$ in terms of the exceedance level $q(\alpha)$, see *e.g.* [35].

For other values of p , an accurate numerical computation of the constant $\kappa(p, \gamma)$ can be carried out instead. Results of such numerical computations on the domain $(p, \gamma) \in [1, 2] \times (0, 1)$ are included in Figure 1. One can observe from this Figure that the functions $p \mapsto \kappa(p, \gamma)$ and $\gamma \mapsto \kappa(p, \gamma)$ appear to be both decreasing. This and (2) entail in particular that, for all $p_1, p_2 \in (1, 2)$ such that $p_1 < p_2$, we have, for α close enough to 1:

$$\text{MS}(\alpha) = m_1(\alpha) < m_{p_1}(\alpha) < m_{p_2}(\alpha) < m_2(\alpha) = \text{CTE}(\alpha).$$

The tail L^p -median $m_p(\alpha)$ can therefore be seen as a risk measure interpolating monotonically between $\text{MS}(\alpha)$ and $\text{CTE}(\alpha)$, at a high enough level α . Actually, Proposition 2 yields, for all $p \in [1, 2]$ and $\gamma < 1/(p - 1)$ that, as $\alpha \rightarrow 1$,

$$m_p(\alpha) \approx \lambda(p, \gamma)\text{MS}(\alpha) + [1 - \lambda(p, \gamma)]\text{CTE}(\alpha), \quad (3)$$

where the weighting constant $\lambda(p, \gamma) \in [0, 1]$ is defined by

$$\lambda(p, \gamma) := \lim_{\alpha \rightarrow 1} \frac{m_p(\alpha) - \text{CTE}(\alpha)}{\text{MS}(\alpha) - \text{CTE}(\alpha)} = \frac{1 - (1 - \gamma)/\kappa(p, \gamma)}{1 - 2^\gamma(1 - \gamma)}. \quad (4)$$

Extreme tail L^p -medians of heavy-tailed models can then be interpreted, for $p \in (1, 2)$, as weighted averages of extreme Median Shortfall and extreme Conditional Expectation at the same level. Because the latter two risk measures are Wang distortion risk measures (see [52]), this (asymptotic) convex combination property ensures that a tail L^p -median $m_p(\alpha)$ can also be interpreted, as $\alpha \rightarrow 1$, as a Wang distortion risk measure, *viz.*

$$m_p(\alpha) \approx \int_0^1 q(1 - \beta) dg_{\lambda, \alpha}(\beta),$$

where $\lambda = \lambda(p, \gamma)$ and the distortion function $g_{\lambda, \alpha}$ is defined by

$$g_{\lambda, \alpha}(x) = \lambda \mathbb{1}_{\{x \geq (1 - \alpha)/2\}} + (1 - \lambda) \min\left(\frac{x}{1 - \alpha}, 1\right).$$

It should be noted that, by contrast, the monotonic interpolation property (and hence the weighted average interpretation) is demonstrably false in general for L^p -quantiles, as is most easily seen from Figure 1 in [14]: this Figure suggests that high L^p -quantiles define, for γ close to $1/2$, a decreasing function of p when it is close to 1 and an increasing function of p when it is close to 2.

The fact that $\gamma \mapsto \kappa(p, \gamma)$ is decreasing, meanwhile, can be proven rigorously by noting that (see the expression of the logarithmic partial derivative of $\kappa(p, \gamma)$ in Theorem 2 below):

$$\frac{\partial \kappa}{\partial \gamma}(p, \gamma) = \frac{\int_{\kappa(p, \gamma)}^1 (1 - u)^{p-1} u^{-1/\gamma-1} \log(u) du - B(p, \gamma^{-1} - p + 1)[\Psi(\gamma^{-1} + 1) - \Psi(\gamma^{-1} - p + 1)]}{\gamma^2 [1 - \kappa(p, \gamma)]^{p-1} [\kappa(p, \gamma)]^{-1/\gamma-1}},$$

where $\Psi(x) = \Gamma'(x)/\Gamma(x)$ is Euler's digamma function (*i.e.* the logarithmic derivative of Euler's Gamma function $\Gamma(x) = \int_0^\infty e^{-t} t^{x-1} dt$), which is an increasing function by Formula 6.3.21 in [1]. More intuitively, the monotonicity of $\gamma \mapsto \kappa(p, \gamma)$ can be seen as a consequence of the heavy-tailedness of the distribution function F . Indeed, we saw that heuristically, as $\alpha \rightarrow 1$,

$$\frac{m_p(\alpha)}{q(\alpha)} \approx \arg \min_{M \in \mathbb{R}} \mathbb{E} [|Z_\gamma - M|^p - |Z_\gamma|^p]$$

where Z_γ is a Pareto random variable with tail index γ . When γ increases, the random variable Z_γ tends to return higher values because its survival function $\mathbb{P}(Z_\gamma > z) = z^{-1/\gamma}$ (for $z > 1$) is an increasing function of γ . We can therefore expect that, as γ increases, a higher value of $M = m_p(\alpha)/q(\alpha)$ will be needed in order to minimise the above cost function.

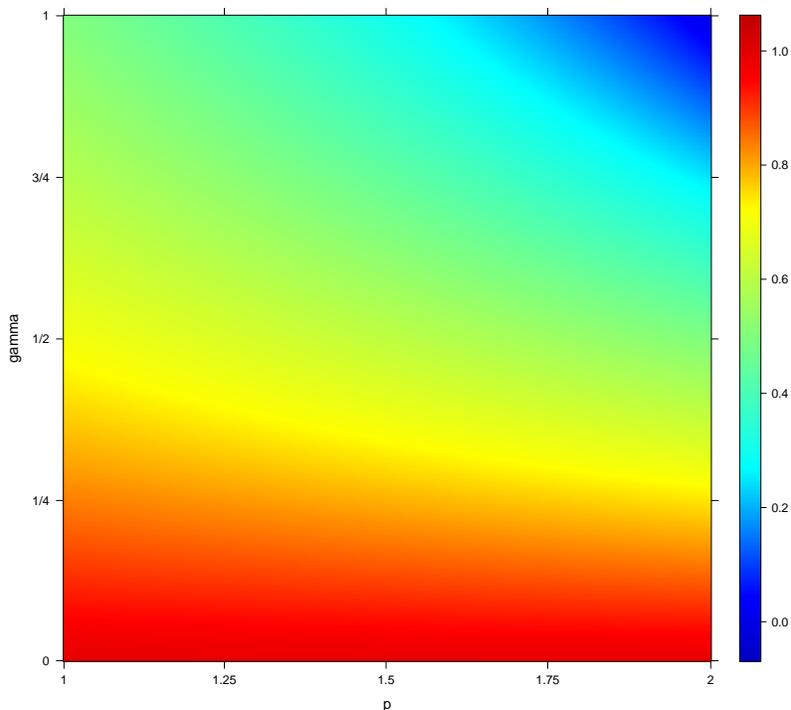


Figure 1: Behaviour of $(p, \gamma) \mapsto \kappa(p, \gamma)$ for $p \in [1, 2]$ and $\gamma \in (0, 1)$.

Our next goal is to derive an asymptotic expansion of the tail L^p -median $m_p(\alpha)$, relatively to the high exceedance level $q(\alpha)$. This will be the key theoretical tool making it possible to analyse the asymptotic properties of estimators of an extreme tail L^p -median. For this, we need to quantify precisely the error term in the convergence given by Proposition 2, and this prompts us to introduce the following second-order regular variation condition:

$\mathcal{C}_2(\gamma, \rho, \mathbf{A})$ The function \bar{F} is second-order regularly varying in a neighbourhood of $+\infty$ with index $-1/\gamma < 0$, second-order parameter $\rho \leq 0$ and an auxiliary function A having constant sign

and converging to 0 at infinity, that is,

$$\forall x > 0, \lim_{t \rightarrow \infty} \frac{1}{A(1/\bar{F}(t))} \left[\frac{\bar{F}(tx)}{\bar{F}(t)} - x^{-1/\gamma} \right] = x^{-1/\gamma} \frac{x^{\rho/\gamma} - 1}{\gamma\rho},$$

where the right-hand side should be read as $\frac{x^{-1/\gamma} \log x}{\gamma^2}$ when $\rho = 0$.

This standard condition on \bar{F} controls the rate of convergence in $\mathbf{C}_1(\gamma)$: the larger $|\rho|$ is, the faster the function $|A|$ converges to 0 (since $|A|$ is regularly varying with index ρ , in virtue of Theorems 2.3.3 and 2.3.9 in [30]) and the smaller the error in the approximation of the right tail of X by a purely Pareto tail will be. Further interpretation of this assumption can be found in [6] and [30] along with various examples of commonly used continuous distributions satisfying it: we mention, among others, that the (Generalised) Pareto, Generalised Extreme Value, Burr, Fréchet, Student, Fisher and Inverse-Gamma distributions all satisfy this condition. More generally, so does any distribution whose survival function \bar{F} satisfies the asymptotic power-type expansion:

$$\bar{F}(x) = x^{-1/\gamma} (a + bx^{-c} + o(x^{-c})) \quad \text{as } x \rightarrow +\infty,$$

where $a > 0$, $b \in \mathbb{R} \setminus \{0\}$ and $c > 0$ are constants. This contains in particular the Hall-Weiss class of models [35]. Let us finally mention that it is a consequence of Theorem 2.3.9 in [30] that $\mathbf{C}_2(\gamma, \rho, \mathbf{A})$ is actually equivalent to the following, perhaps more usual extremal assumption on the tail quantile function U :

$$\forall x > 0, \lim_{t \rightarrow \infty} \frac{1}{A(t)} \left[\frac{U(tx)}{U(t)} - x^\gamma \right] = x^\gamma \frac{x^\rho - 1}{\rho}. \quad (5)$$

The next result is the desired refinement of Proposition 2 giving the error term in the asymptotic proportionality relationship linking $m_p(\alpha)$ to $q(\alpha)$ when α is close to 1.

Proposition 3. *Suppose that $p \geq 1$ and $\mathbf{C}_2(\gamma, \rho, \mathbf{A})$ holds with $\gamma < 1/(p-1)$. Then, as $\alpha \rightarrow 1$,*

$$\frac{m_p(\alpha)}{q(\alpha)} = \frac{1}{\kappa(p, \gamma)} (1 + [R(p, \gamma, \rho) + o(1)]A((1-\alpha)^{-1}))$$

with

$$R(p, \gamma, \rho) = \frac{[\kappa(p, \gamma)]^{-(\rho-1)/\gamma}}{(1 - \kappa(p, \gamma))^{p-1}} \times \frac{1-\rho}{\gamma\rho} [B(p, (1-\rho)\gamma^{-1} - p + 1) - g_{p, \gamma/(1-\rho)}(\kappa(p, \gamma))].$$

This result is similar in spirit to second-order results that have been shown for other extreme risk measures: we refer again to [57, 58], as well as to [14, 15, 22] for analogue results used as a basis to carry out extreme-value based inference on other types of indicators. It should, however, be underlined that the asymptotic expansion of an extreme tail L^p -median depends solely on the extreme parameters γ , ρ and A , along with the power p . By contrast, the asymptotic expansion

of an extreme L^p -quantile depends on the expectation and left-tail behaviour of X , which are typically considered to be irrelevant to the understanding of the right tail of X . From an extreme value point of view, the asymptotic expansion of extreme tail L^p -medians is therefore easier to understand than that of L^p -quantiles. Statistically speaking, it also implies that there are less sources of potential bias in extreme tail L^p -median estimation than in extreme L^p -quantile estimation. Both of these statements can be explained by the fact that the tail L^p -median is constructed exclusively on the event $\{X > q_p(\alpha)\}$, while the equation defining an L^p -quantile $q_p(\alpha)$ is

$$(1 - \alpha)\mathbb{E}((q_p(\alpha) - X)^{p-1}\mathbb{1}_{\{X < q_p(\alpha)\}}) = \alpha\mathbb{E}((X - q_p(\alpha))^{p-1}\mathbb{1}_{\{X > q_p(\alpha)\}})$$

(see [13, Section 2]). The left-hand side term ensures that the central and left-tail behaviour of X will necessarily have an influence on the value of any L^p -quantile, even at an extreme level. The construction of a tail L^p -median as an L^p -median in the right tail of X removes this issue.

We close this section by noting that, like the asymptotic proportionality constant $\kappa(p, \gamma)$ on which it depends, the remainder term $R(p, \gamma, \rho)$ does not have an explicit form in general. That being said, $R(1, \gamma, \rho)$ and $R(2, \gamma, \rho)$ have simple explicit values: for $p = 1$, $R(1, \gamma, \rho) = (2^\rho - 1)/\rho$ (with the convention $(x^\rho - 1)/\rho = \log x$ for $\rho = 0$), see Lemma 10(iii) in the Supplementary Material document. We then find back the result that comes as an immediate consequence of (2) and our second-order condition $\mathcal{C}_2(\gamma, \rho, \mathbf{A})$:

$$\frac{m_1(\alpha)}{q(\alpha)} = \frac{U(2(1 - \alpha)^{-1})}{U((1 - \alpha)^{-1})} = 2^\gamma \left(1 + A((1 - \alpha)^{-1}) \frac{2^\rho - 1}{\rho} (1 + o(1)) \right).$$

For $p = 2$ and $\gamma \in (0, 1)$, condition $\mathcal{C}_2(\gamma, \rho, \mathbf{A})$ and (2) suggest that:

$$\frac{m_2(\alpha)}{q(\alpha)} = \int_1^\infty \frac{U((1 - \alpha)^{-1}x)}{U((1 - \alpha)^{-1})} \frac{dx}{x^2} = \frac{1}{1 - \gamma} \left(1 + \frac{1}{1 - \gamma - \rho} A((1 - \alpha)^{-1}) (1 + o(1)) \right)$$

which coincides with Proposition 3, given that $R(2, \gamma, \rho) = 1/(1 - \gamma - \rho)$ from Lemma 11(iii) in the Supplementary Material document.

4 Estimation of an extreme tail L^p -median

Suppose that we observe a random sample (X_1, \dots, X_n) of independent copies of X , and denote by $X_{1,n} \leq \dots \leq X_{n,n}$ the corresponding set of order statistics arranged in increasing order. Our goal in this section is to estimate an extreme tail L^p -median $m_p(\alpha_n)$, where $\alpha_n \rightarrow 1$ as $n \rightarrow \infty$. The final aim is to allow α_n to approach 1 at any rate, covering the cases of an intermediate tail L^p -median with $n(1 - \alpha_n) \rightarrow \infty$ and of a proper extreme tail L^p -median with $n(1 - \alpha_n) \rightarrow c$, where c is some finite positive constant.

4.1 Intermediate case: direct estimation by empirical L^p -optimisation

Recall that the tail L^p -median $m_p(\alpha_n)$ is, by (1),

$$m_p(\alpha_n) = \arg \min_{m \in \mathbb{R}} \mathbb{E} (|X - m|^p - |X|^p \mid X > q(\alpha_n)).$$

Assume here that $n(1 - \alpha_n) \rightarrow \infty$, so that $\lfloor n(1 - \alpha_n) \rfloor > 0$ eventually. The direct empirical analogue of the cost function on the right-hand side above is:

$$\frac{1}{\lfloor n(1 - \alpha_n) \rfloor} \sum_{i=1}^{\lfloor n(1 - \alpha_n) \rfloor} (|X_{n-i+1,n} - m|^p - |X_{n-i+1,n}|^p).$$

We can therefore define a direct empirical tail L^p -median estimator of $m_p(\alpha_n)$ by minimising the above empirical cost function:

$$\hat{m}_p(\alpha_n) = \arg \min_{m \in \mathbb{R}} \frac{1}{\lfloor n(1 - \alpha_n) \rfloor} \sum_{i=1}^{\lfloor n(1 - \alpha_n) \rfloor} (|X_{n-i+1,n} - m|^p - |X_{n-i+1,n}|^p). \quad (6)$$

We now pave the way for a theoretical study of this estimator. The key point is that since normalising constants and shifts are irrelevant in the definition of the empirical criterion, we clearly have the equivalent definition

$$\hat{m}_p(\alpha_n) = \arg \min_{m \in \mathbb{R}} \frac{1}{p[m_p(\alpha_n)]^p} \sum_{i=1}^{\lfloor n(1 - \alpha_n) \rfloor} (|X_{n-i+1,n} - m|^p - |X_{n-i+1,n} - m_p(\alpha_n)|^p).$$

Consequently

$$\sqrt{n(1 - \alpha_n)} \left(\frac{\hat{m}_p(\alpha_n)}{m_p(\alpha_n)} - 1 \right) = \arg \min_{u \in \mathbb{R}} \psi_n(u; p)$$

with

$$\psi_n(u; p) = \frac{1}{p[m_p(\alpha_n)]^p} \sum_{i=1}^{\lfloor n(1 - \alpha_n) \rfloor} \left(\left| X_{n-i+1,n} - m_p(\alpha_n) - \frac{u m_p(\alpha_n)}{\sqrt{n(1 - \alpha_n)}} \right|^p - |X_{n-i+1,n} - m_p(\alpha_n)|^p \right). \quad (7)$$

Note that the empirical criterion $\psi_n(u; p)$ is a continuous and convex function of u , so that the asymptotic properties of the minimiser follow directly from those of the criterion itself by the convexity lemma of [27] (see also [38]). The empirical criterion, then, is analysed by using its continuous differentiability (for $p > 1$) in order to formulate an L^p -analogue of Knight's identity [37] and divide the work between, on the one hand, the study of a $\sqrt{n(1 - \alpha_n)}$ -consistent and asymptotically Gaussian term which is an affine function of u and, on the other hand, a bias term which converges to a nonrandom multiple of u^2 . Further technical details are provided in the Supplementary Material document, and especially Lemmas 3, 5, 7 and 9.

This programme of work is broadly similar to that of [14] for the convergence of the direct intermediate L^p -quantile estimator. The difficulty in this particular case, however, is twofold: first, the

affine function of u is a generalised L -statistic (in the sense of *e.g.* [9]) whose analysis requires delicate arguments relying on the asymptotic behaviour of the tail quantile process *via* Theorem 5.1.4 p.161 in [30]. For L^p -quantiles, this is not necessary because the affine term is actually a sum of independent, identically distributed and centred variables. Second, the bias term is essentially a doubly integrated oscillation of a power function with generally noninteger exponent. The examination of its convergence requires certain precise real analysis arguments which do not follow from those developed in [14] for the asymptotic analysis of intermediate L^p -quantiles.

With this in mind, the asymptotic normality result for the direct intermediate tail L^p -median estimator $\widehat{m}_p(\alpha_n)$ is the following.

Theorem 1. *Suppose that $p \geq 1$ and $\mathcal{C}_2(\gamma, \rho, \mathbf{A})$ holds with $\gamma < 1/[2(p-1)]$. Assume further that $\alpha_n \rightarrow 1$ is such that $n(1-\alpha_n) \rightarrow \infty$ and $\sqrt{n(1-\alpha_n)}A((1-\alpha_n)^{-1}) = O(1)$. Then we have, as $n \rightarrow \infty$:*

$$\sqrt{n(1-\alpha_n)} \left(\frac{\widehat{m}_p(\alpha_n)}{m_p(\alpha_n)} - 1 \right) \xrightarrow{d} \mathcal{N}(0, V(p, \gamma)).$$

Here $V(p, \gamma) = V_1(p, \gamma)/V_2(p, \gamma)$ with

$$V_1(p, \gamma) = \frac{[\kappa(p, \gamma)]^{1/\gamma}}{\gamma} (B(2p-1, \gamma^{-1} - 2(p-1)) + g_{2p-1, \gamma}(\kappa(p, \gamma))) + [1 - \kappa(p, \gamma)]^{2(p-1)}$$

while $V_2(p, \gamma)$ is defined by: $V_2(1, \gamma) = 1/\gamma^2$ and, for $p > 1$,

$$V_2(p, \gamma) = \left(\frac{p-1}{\gamma} [\kappa(p, \gamma)]^{1/\gamma} [B(p-1, \gamma^{-1} - p + 2) + g_{p-1, \gamma}(\kappa(p, \gamma))] \right)^2.$$

Moreover, the functions $V_2(\cdot, \gamma)$ and $V(\cdot, \gamma)$ defined this way are right-continuous at 1.

The variance $V(p, \gamma)$ has a rather involved expression. We include in Figures 2 and 3 some plots to understand its behaviour. The top picture in Figure 2 provides a graphical representation of the function $\gamma \mapsto V(p, \gamma)$, on the interval $(0, 1/2)$ and on evenly spaced values of $p \in [1, 2]$. It can be seen here that the function $\gamma \mapsto V(p, \gamma)$ appears to be increasing. This reflects the increasing tendency of the underlying distribution to generate extremely high observations when the tail index increases (see, among others, Section 3.1 in [23]), thus increasing the variability of the empirical criterion $\psi_n(\cdot; p)$ and consequently that of its minimiser. The bottom panel of Figure 2 provides further evidence of that with a plot of the function $(p, \gamma) \mapsto V(p, \gamma)$ on the domain $[1, 1.8] \times [0.25, 0.5]$. It is not, however, clear from any of these two graphs that the function $p \mapsto V(p, \gamma)$ is monotonic, like the proportionality constant κ was. It turns out that, somewhat surprisingly, the function $p \mapsto V(p, \gamma)$ is not in general a monotonic function of p , and an illustration is provided for $\gamma = 0.3$ in Figure 3. A numerical study, which is not reported here, actually shows that we have $V(p, \gamma) < V(1, \gamma)$ for any $(p, \gamma) \in (1, 1.2] \times [0.25, 0.5]$. This suggests that for all heavy-tailed distributions having only a second moment (an already difficult case as

far as estimation in heavy-tailed models is concerned), a direct L^p -tail median estimator with $p \in (1, 1.2]$ will have a smaller asymptotic variance than the empirical L^1 -tail median estimator, or, in other words, the empirical Median Shortfall.

We conclude this section by noting that, like the constants appearing in our previous asymptotic results, the variance term $V(p, \gamma)$ has a simple expression when $p = 1$ or $p = 2$. In the case $p = 1$, we have $V(1, \gamma) = 2\gamma^2$ by Lemma 10(iv) in the Supplementary Material document. Statement (2) suggests that this should be identical to the asymptotic variance of the high quantile estimator $\widehat{\text{MS}}(\alpha_n) = X_{n-\lfloor n(1-\alpha_n)/2 \rfloor, n}$, and indeed

$$\sqrt{n(1-\alpha_n)} \left(\frac{\widehat{\text{MS}}(\alpha_n)}{\text{MS}(\alpha_n)} - 1 \right) = \sqrt{2} \sqrt{\frac{n(1-\alpha_n)}{2}} \left(\frac{X_{n-\lfloor n(1-\alpha_n)/2 \rfloor, n}}{q(1-(1-\alpha_n)/2)} - 1 \right) \xrightarrow{d} \mathcal{N}(0, 2\gamma^2)$$

by Theorem 2.4.8 p.52 in [30]. For $p = 2$, we have $V(2, \gamma) = 2\gamma^2(1-\gamma)/(1-2\gamma)$ by Lemma 10(iv) in the Supplementary Material document. By (2), we should expect this constant to coincide with the asymptotic variance of the empirical counterpart of the Conditional Tail Expectation, namely

$$\widehat{\text{CTE}}(\alpha_n) = \frac{1}{\lfloor n(1-\alpha_n) \rfloor} \sum_{i=1}^{\lfloor n(1-\alpha_n) \rfloor} X_{n-i+1, n}.$$

This is indeed the case as Corollary 1 in [20] shows; see also Theorem 2 in [21]. In particular, the function $\gamma \mapsto V(2, \gamma)$ tends to infinity as $\gamma \uparrow 1/2$, reflecting the increasing difficulty of estimating a high Conditional Tail Expectation by its direct empirical counterpart as the right tail of X gets heavier; see also the top panel of Figure 2 for a plot of $\gamma \mapsto V(2, \gamma)$ on the interval $(0, 1/2)$.

4.2 Intermediate case: indirect quantile-based estimation

We can also design an estimator of $m_p(\alpha_n)$ based on the asymptotic equivalence between $m_p(\alpha_n)$ and $q(\alpha_n)$ that is provided by Proposition 2. Indeed, since this result suggests that $m_p(\alpha)/q(\alpha) \approx 1/\kappa(p, \gamma)$ when α is close to 1, it makes sense to build a plug-in estimator of $m_p(\alpha_n)$ in the following way:

$$\tilde{m}_p(\alpha_n) = \frac{\hat{q}(\alpha_n)}{\kappa(p, \hat{\gamma}_n)},$$

where $\hat{q}(\alpha_n)$ and $\hat{\gamma}_n$ are respectively two consistent estimators of the high quantile $q(\alpha_n)$ and of the tail index γ . Since we work here in the intermediate case $n(1-\alpha_n) \rightarrow \infty$, we know that the sample counterpart $X_{\lfloor n\alpha_n \rfloor, n}$ of $q(\alpha_n)$ is a relatively consistent estimator of $q(\alpha_n)$, see Theorem 2.4.1 in [30]. This suggests to use the estimator

$$\tilde{m}_p(\alpha_n) := \frac{X_{\lfloor n\alpha_n \rfloor, n}}{\kappa(p, \hat{\gamma}_n)}. \quad (8)$$

Our next result analyses the asymptotic distribution of this estimator, assuming that the pair $(\hat{\gamma}_n, X_{\lfloor n\alpha_n \rfloor, n})$ is jointly $\sqrt{n(1-\alpha_n)}$ -consistent.

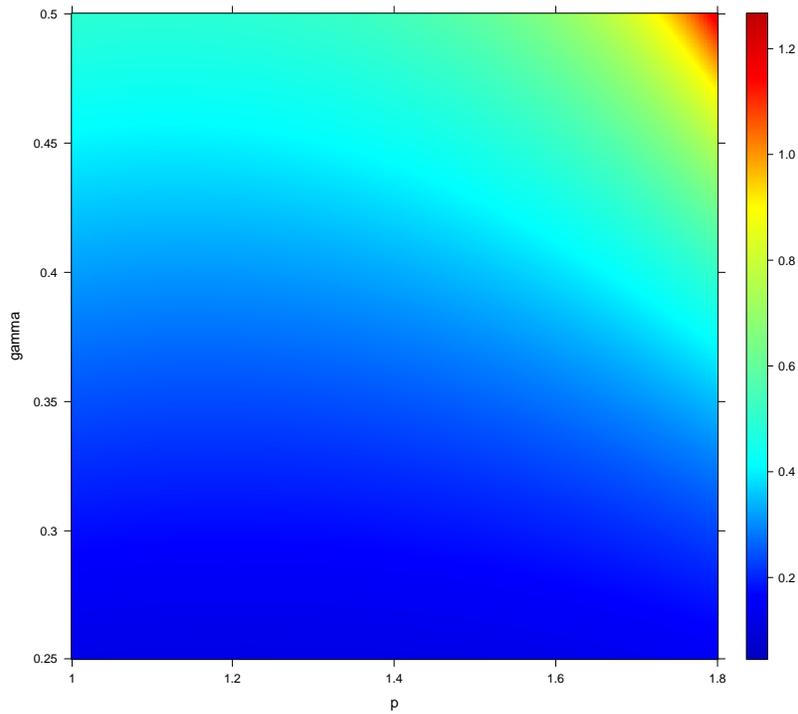
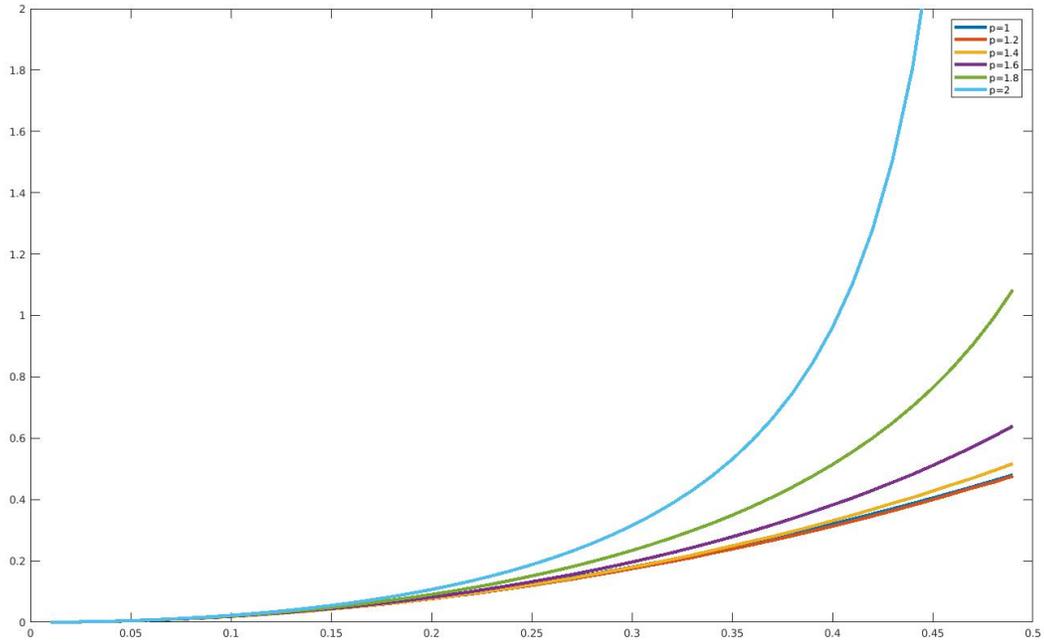


Figure 2: Behaviour of the asymptotic variance V of the direct empirical tail L^p -median estimator. Top panel: a plot of the function $\gamma \mapsto V(p, \gamma)$, on the interval $(0, 1/2)$, for $p \in \{1, 1.2, 1.4, 1.6, 1.8, 2\}$. Bottom panel: a plot of the function $(p, \gamma) \mapsto V(p, \gamma)$ on the domain $[1, 1.8] \times [0.25, 0.5]$.

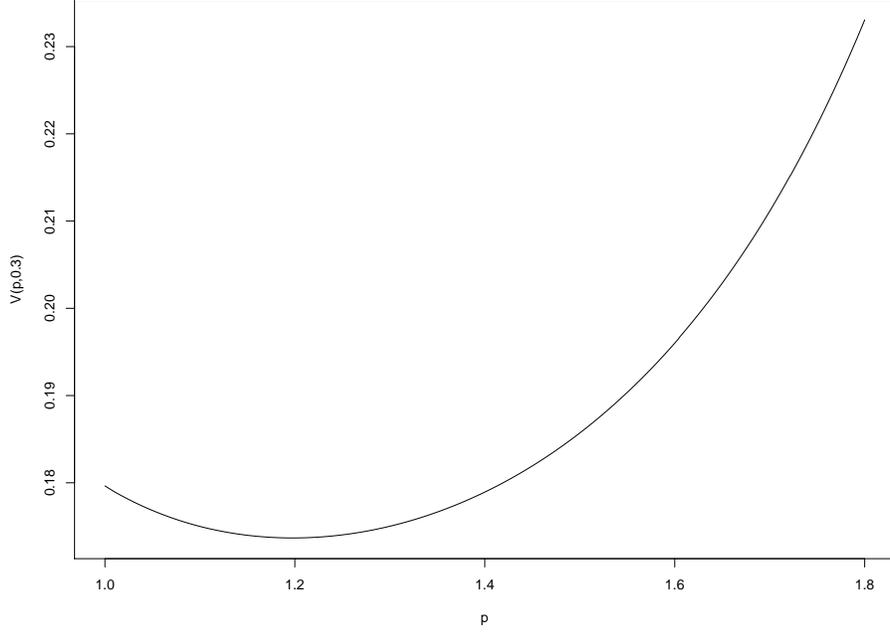


Figure 3: Behaviour of the function $p \mapsto V(p, 0.3)$ on the interval $[1, 1.8]$.

Theorem 2. *Suppose that $p \geq 1$ and $\mathbf{C}_2(\gamma, \rho, \mathbf{A})$ holds with $\gamma < 1/(p-1)$. Assume further that $\alpha_n \rightarrow 1$ is such that $n(1-\alpha_n) \rightarrow \infty$ and $\sqrt{n(1-\alpha_n)}A((1-\alpha_n)^{-1}) \rightarrow \lambda \in \mathbb{R}$, and that*

$$\sqrt{n(1-\alpha_n)} \left(\hat{\gamma}_n - \gamma, \frac{X_{[n\alpha_n],n}}{q(\alpha_n)} - 1 \right) \xrightarrow{d} (\xi_1, \xi_2)$$

where (ξ_1, ξ_2) is a pair of nondegenerate random variables. Then we have, as $n \rightarrow \infty$:

$$\sqrt{n(1-\alpha_n)} \left(\frac{\tilde{m}_p(\alpha_n)}{m_p(\alpha_n)} - 1 \right) \xrightarrow{d} \sigma(p, \gamma)\xi_1 + \xi_2 - \lambda R(p, \gamma, \rho),$$

with the positive constant $\sigma(p, \gamma)$ being $\sigma(p, \gamma) = -\frac{1}{\kappa(p, \gamma)} \frac{\partial \kappa}{\partial \gamma}(p, \gamma)$. In other words,

$$\sigma(p, \gamma) = \frac{B(p, \gamma^{-1} - p + 1)[\Psi(\gamma^{-1} + 1) - \Psi(\gamma^{-1} - p + 1)] - \int_{\kappa(p, \gamma)}^1 (1-u)^{p-1} u^{-1/\gamma-1} \log(u) du}{\gamma^2 [1 - \kappa(p, \gamma)]^{p-1} [\kappa(p, \gamma)]^{-1/\gamma}}$$

where $\Psi(x) = \Gamma'(x)/\Gamma(x)$ is Euler's digamma function.

Again, the constant $\sigma(p, \gamma)$ does not generally have a simple explicit form, but we can compute it when $p = 1$ or $p = 2$. Lemma 10(v) shows that $\sigma(1, \gamma) = \log 2$, while Lemma 11(v) entails $\sigma(2, \gamma) = 1/(1-\gamma)$ (see the Supplementary Material document). Contrary to our previous analyses, it is more difficult to relate these constants to pre-existing results in high quantile or high CTE estimation because Theorem 2 is a general result that applies to a wide range of estimators $\hat{\gamma}_n$. To

the best of our knowledge, there is no general analogue of this result in the literature for the case $p = 1$. In the case $p = 2$, we find back the asymptotic distribution result in Theorem 1 of [22]:

$$\sqrt{n(1-\alpha_n)} \left(\frac{\tilde{m}_2(\alpha_n)}{\text{CTE}(\alpha_n)} - 1 \right) \xrightarrow{d} \frac{1}{1-\gamma} \xi_1 + \xi_2 - \frac{\lambda}{1-\gamma-\rho}.$$

It should be highlighted that, for $p = 2$, Theorem 2 in the present paper is a stronger result than Theorem 1 of [22], since the condition on γ is less stringent than that of Theorem 1 therein. As the above convergence is valid for any $\gamma < 1$, one may therefore think that the estimator $\tilde{m}_2(\alpha_n)$ is a widely applicable estimator of the CTE at high levels. Since it is also more robust than the direct empirical CTE estimator due to its reliance on the sample quantile $X_{[n\alpha_n],n}$, this would defeat the point of looking for a middle ground solution between the sensitivity of CTE to high values and the robustness of VaR-type measures in very heavy-tailed models. The simulation study in [22] shows however that in general, the estimator \tilde{m}_2 fares worse than the direct CTE estimator \hat{m}_2 , and increasingly so as γ increases within the range $(0, 1/4]$. We will also confirm this in our simulation study by considering several cases with higher values of γ and showing that the estimator \tilde{m}_2 should in general not be preferred to \hat{m}_2 . The benefit of using the indirect estimator \tilde{m}_p will rather be found for values of p away from 2, when a genuine compromise between sensitivity and robustness is sought.

Theorem 2 applies whenever $\hat{\gamma}_n$ is a consistent estimator of γ that satisfies a joint convergence condition together with the intermediate order statistic $X_{[n\alpha_n],n}$. This is not a restrictive requirement. For instance, if $\hat{\gamma}_n = \hat{\gamma}_n^H$ is the widely used Hill estimator [31]:

$$\hat{\gamma}_n^H := \frac{1}{[n(1-\alpha_n)]} \sum_{i=1}^{[n(1-\alpha_n)]} \log(X_{n-i+1,n}) - \log(X_{n-[n(1-\alpha_n)],n}), \quad (9)$$

then, under the bias condition $\sqrt{n(1-\alpha_n)}A((1-\alpha_n)^{-1}) \rightarrow \lambda$, we have the joint convergence

$$\sqrt{n(1-\alpha_n)} \left(\hat{\gamma}_n^H - \gamma, \frac{X_{[n\alpha_n],n}}{q(\alpha_n)} - 1 \right) \xrightarrow{d} \left(\gamma N_1 + \frac{\lambda}{1-\rho}, \gamma N_2 \right)$$

where (N_1, N_2) is a pair of independent standard Gaussian random variables (for a proof, combine Theorem 2.4.1, Lemma 3.2.3 and Theorem 3.2.5 in [30]). As a corollary of Theorem 2, we then get the following asymptotic result on $\tilde{m}_p(\alpha_n)$ when the estimator $\hat{\gamma}_n$ is the Hill estimator $\hat{\gamma}_n^H$.

Corollary 1. *Suppose that $p \geq 1$ and $\mathcal{C}_2(\gamma, \rho, \mathbf{A})$ holds with $\gamma < 1/(p-1)$. Assume further that $\alpha_n \rightarrow 1$ is such that $n(1-\alpha_n) \rightarrow \infty$ and $\sqrt{n(1-\alpha_n)}A((1-\alpha_n)^{-1}) \rightarrow \lambda \in \mathbb{R}$. Then, as $n \rightarrow \infty$:*

$$\sqrt{n(1-\alpha_n)} \left(\frac{\tilde{m}_p(\alpha_n)}{m_p(\alpha_n)} - 1 \right) \xrightarrow{d} \mathcal{N} \left(\lambda \left(\frac{\sigma(p, \gamma)}{1-\rho} - R(p, \gamma, \rho) \right), \mathbf{v}(p, \gamma) \right),$$

where $\mathbf{v}(p, \gamma) = \gamma^2 ([\sigma(p, \gamma)]^2 + 1)$.

The asymptotic variance $\mathfrak{v}(p, \gamma)$ is plotted, for several values of p , on Figure 4, and a comparison of the asymptotic variance $V(p, \gamma)$ of the direct estimator with $\mathfrak{v}(p, \gamma)$ is depicted on Figure 5, for $\gamma \in [1/4, 1/2]$. It appears from Figure 5 that the indirect estimator has a lower variance than the direct one. The difference between the two variances becomes sizeable when the quantity $2\gamma(p-1)$ gets closer to 1, as should be expected since the asymptotic variance of the direct estimator asymptotically increases to infinity (see Theorem 1), while the asymptotic variance of the indirect estimator is kept under control (see Corollary 1). This seems to indicate that the indirect estimator should be preferred to the direct estimator in terms of variability. However, the indirect estimator is asymptotically biased (due to the use of the approximation $m_p(\alpha)/q(\alpha) \approx 1/\kappa(p, \gamma)$ in its construction), while the direct estimator is not, and we will see that this can make one prefer the direct estimator in terms of mean squared error on finite samples, even for large values of γ and p .

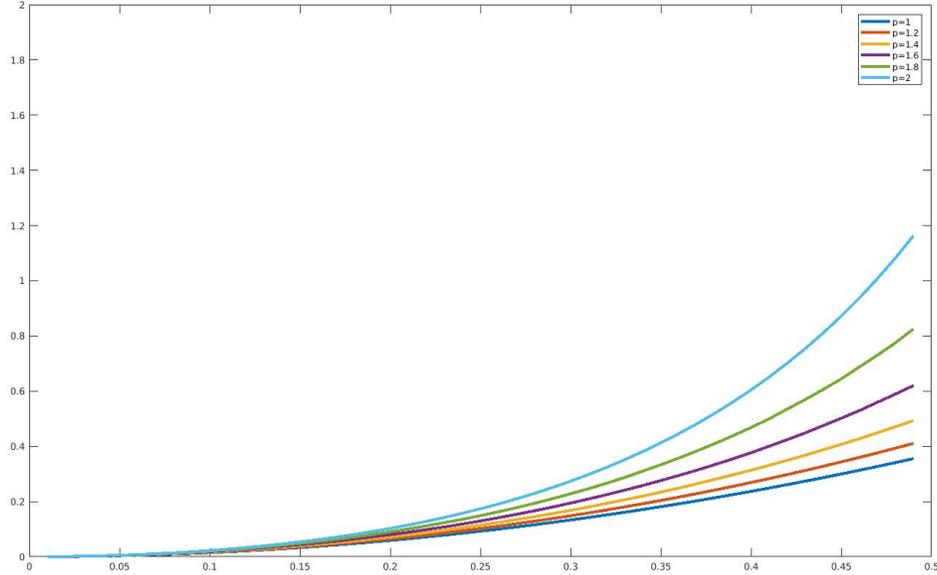


Figure 4: Behaviour of the asymptotic variance \mathfrak{v} of the indirect tail L^p -median estimator: plot of the function $\gamma \mapsto \mathfrak{v}(p, \gamma)$, on the interval $(0, 1/2)$, for $p \in \{1, 1.2, 1.4, 1.6, 1.8, 2\}$.

We conclude this section by mentioning that similar joint convergence results on $(\widehat{\gamma}_n, X_{[n\alpha n], n})$, and therefore analogues of Corollary 1, can be found for a wide range of other estimators of γ . We mention for instance the Pickands estimator [46], the maximum likelihood estimator in an approximate Generalised Pareto models for exceedances [18, 49] and probability-weighted moment estimators [17, 34]; we refer to *e.g.* Sections 3 and 4 of [30] for an asymptotic analysis of these alternatives. This can be used to construct several other indirect tail L^p -median estimators.

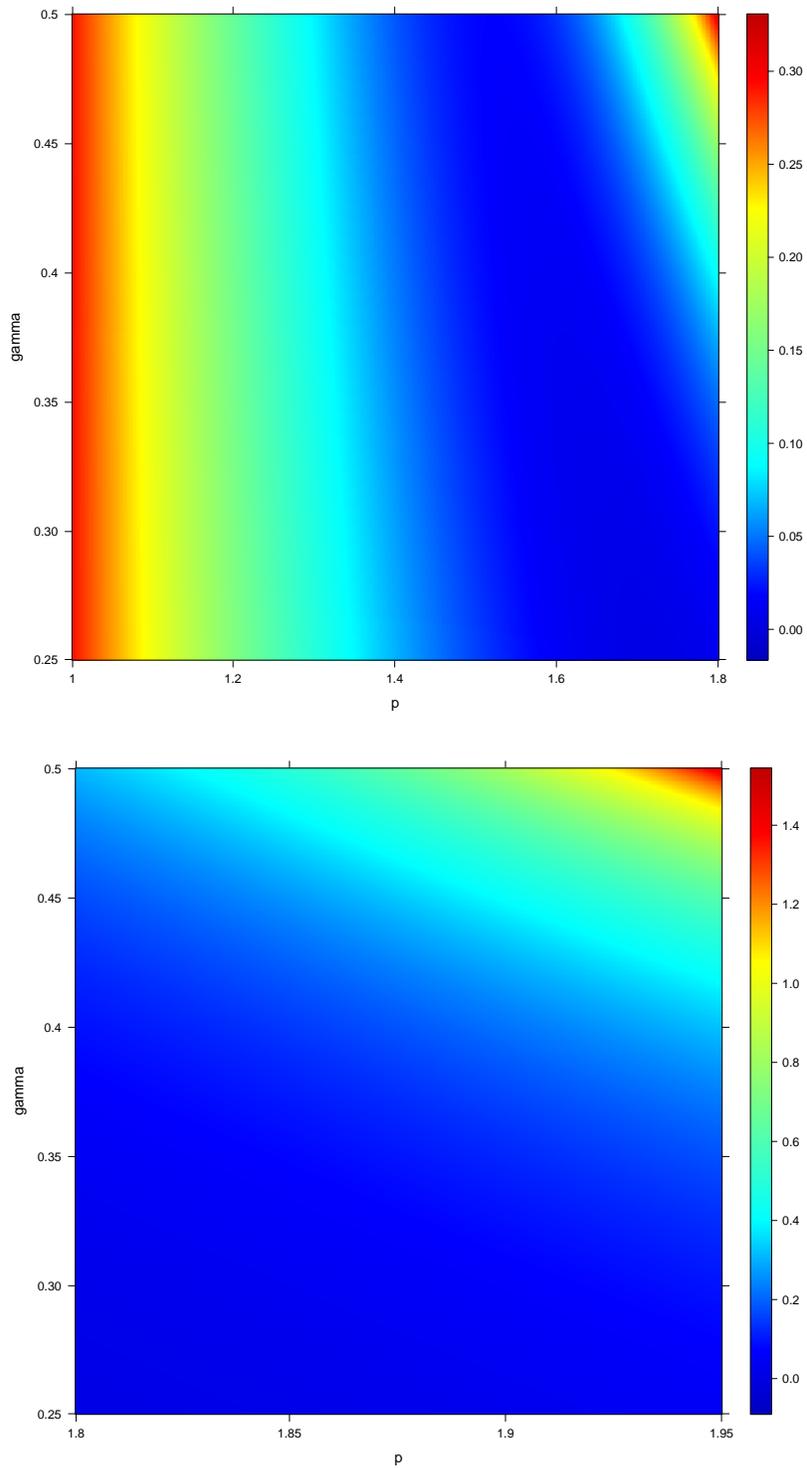


Figure 5: A comparison between the asymptotic variances V and \mathbf{v} with plots of the function $(p, \gamma) \mapsto \log(V(p, \gamma)/\mathbf{v}(p, \gamma))$. Top: on the domain $(p, \gamma) \in [1, 1.8] \times [0.25, 0.5]$, bottom: on the domain $(p, \gamma) \in [1.8, 1.95] \times [0.25, 0.5]$.

4.3 Extreme case: an extrapolation device

Both the direct and indirect estimators constructed so far are only consistent for intermediate sequences α_n such that $n(1 - \alpha_n) \rightarrow \infty$. Our purpose is now to extrapolate these intermediate tail L^p -median estimators to proper extreme levels $\beta_n \rightarrow 1$ with $n(1 - \beta_n) \rightarrow c < \infty$ as $n \rightarrow \infty$. The extrapolation argument is based on the fact that under the regular variation condition $\mathcal{C}_1(\gamma)$, the function $t \mapsto U(t) = q(1 - t^{-1})$ is regularly varying with index γ . In particular, we have:

$$\frac{q(\beta_n)}{q(\alpha_n)} = \frac{U((1 - \beta_n)^{-1})}{U((1 - \alpha_n)^{-1})} \approx \left(\frac{1 - \beta_n}{1 - \alpha_n} \right)^{-\gamma}$$

when $\alpha_n, \beta_n \rightarrow 1$. This approximation is at the heart of the construction of the classical Weissman extreme quantile estimator $\hat{q}^W(\beta_n)$, introduced in [55]:

$$\hat{q}^W(\beta_n) := \left(\frac{1 - \beta_n}{1 - \alpha_n} \right)^{-\hat{\gamma}_n} X_{\lceil n\alpha_n \rceil, n}.$$

The key point is then that, when $\gamma < 1/(p - 1)$, the quantity $m_p(\alpha)$ is asymptotically proportional to $q(\alpha)$, by Proposition 2. Combining this with the above approximation on ratios of high quantiles suggests the following extrapolation formula:

$$m_p(\beta_n) \approx \left(\frac{1 - \beta_n}{1 - \alpha_n} \right)^{-\gamma} m_p(\alpha_n).$$

An estimator of the extreme tail L^p -median $m_p(\beta_n)$ is obtained from this approximation by plugging in a consistent estimator $\hat{\gamma}_n$ of γ and a consistent estimator of $m_p(\alpha_n)$. In our context, the latter can be the direct, empirical L^p -estimator $\hat{m}_p(\alpha_n)$, yielding the extrapolated estimator

$$\hat{m}_p^W(\beta_n) := \left(\frac{1 - \beta_n}{1 - \alpha_n} \right)^{-\hat{\gamma}_n} \hat{m}_p(\alpha_n).$$

Alternatively, we can use the indirect, intermediate quantile-based estimator $\tilde{m}_p(\alpha_n)$, which gives the extrapolated estimator

$$\tilde{m}_p^W(\beta_n) := \left(\frac{1 - \beta_n}{1 - \alpha_n} \right)^{-\hat{\gamma}_n} \tilde{m}_p(\alpha_n).$$

We note, moreover, that the latter estimator is precisely the estimator deduced by plugging the Weissman extreme quantile estimator $\hat{q}^W(\beta_n)$ in the asymptotic proportionality relationship $m_p(\beta_n)/q(\beta_n) \approx 1/\kappa(p, \gamma)$, since

$$\tilde{m}_p^W(\beta_n) = \left(\frac{1 - \beta_n}{1 - \alpha_n} \right)^{-\hat{\gamma}_n} \times \left\{ \frac{X_{\lceil n\alpha_n \rceil, n}}{\kappa(p, \hat{\gamma}_n)} \right\} = \frac{\hat{q}^W(\beta_n)}{\kappa(p, \hat{\gamma}_n)}.$$

The asymptotic behaviour of the two extrapolated estimators $\hat{m}_p^W(\beta_n)$ or $\tilde{m}_p^W(\beta_n)$ is analysed in our next main result.

Theorem 3. Suppose that $p \geq 1$ and $\mathbf{C}_2(\gamma, \rho, \mathbf{A})$ holds with $\rho < 0$. Assume also that $\alpha_n, \beta_n \rightarrow 1$ are such that $n(1-\alpha_n) \rightarrow \infty$ and $n(1-\beta_n) \rightarrow c < \infty$, with $\sqrt{n(1-\alpha_n)}/\log([1-\alpha_n]/[1-\beta_n]) \rightarrow \infty$. Assume finally that

$$\sqrt{n(1-\alpha_n)}(\hat{\gamma}_n - \gamma) \xrightarrow{d} \xi$$

where ξ is a nondegenerate limiting random variable.

(i) If $\gamma < 1/[2(p-1)]$ and $\sqrt{n(1-\alpha_n)}A((1-\alpha_n)^{-1}) = O(1)$ then, as $n \rightarrow \infty$:

$$\frac{\sqrt{n(1-\alpha_n)}}{\log[(1-\alpha_n)/(1-\beta_n)]} \left(\frac{\hat{m}_p^W(\beta_n)}{m_p(\beta_n)} - 1 \right) \xrightarrow{d} \xi.$$

(ii) If $\gamma < 1/(p-1)$ and $\sqrt{n(1-\alpha_n)}A((1-\alpha_n)^{-1}) \rightarrow \lambda \in \mathbb{R}$ then, as $n \rightarrow \infty$:

$$\frac{\sqrt{n(1-\alpha_n)}}{\log[(1-\alpha_n)/(1-\beta_n)]} \left(\frac{\tilde{m}_p^W(\beta_n)}{m_p(\beta_n)} - 1 \right) \xrightarrow{d} \xi.$$

This result shows that both of the estimators $\hat{m}_p^W(\beta_n)$ and $\tilde{m}_p^W(\beta_n)$ have their asymptotic properties governed by those of the tail index estimator $\hat{\gamma}_n$. This is not an unusual phenomenon for extrapolated estimators: actually, the very fact that these two estimators are built on an intermediate tail L^p -median estimator and a tail index estimator $\hat{\gamma}_n$ sharing the same rate of convergence guarantees that the asymptotic behaviour of $\hat{\gamma}_n$ will dominate. A brief, theoretical justification for this is that while the intermediate tail L^p -median estimator is $\sqrt{n(1-\alpha_n)}$ -relatively consistent, the (estimated) extrapolation factor $([1-\beta_n]/[1-\alpha_n])^{-\hat{\gamma}_n}$, whose asymptotic behaviour only depends on that of $\hat{\gamma}_n$, converges relatively to $([1-\beta_n]/[1-\alpha_n])^{-\gamma}$ with the slower rate of convergence $\sqrt{n(1-\alpha_n)}/\log[(1-\alpha_n)/(1-\beta_n)]$. This is explained in detail in the proof of Theorem 3, and we also refer to Theorem 4.3.8 of [30] and its proof for a detailed exposition regarding the Weissman quantile estimator. In particular, if $\hat{\gamma}_n$ is the Hill estimator (9), then the common asymptotic distribution of our extrapolated estimators will be Gaussian with mean $\lambda/(1-\rho)$ and variance γ^2 , provided $\sqrt{n(1-\alpha_n)}A((1-\alpha_n)^{-1}) \rightarrow \lambda \in \mathbb{R}$, see Theorem 3.2.5 in [30].

Let us highlight though that while the asymptotic behaviour of $\hat{\gamma}_n$ is crucial, we should anticipate that in finite-sample situations, an accurate estimation of the intermediate tail L^p -median $m_p(\alpha_n)$ is also important. A mathematical reason for this is that in the typical situation when $1-\beta_n = 1/n$ (considered recently by *e.g.* [3, 11, 12, 29]), the logarithmic term $\log[(1-\alpha_n)/(1-\beta_n)]$ has order at most $\log(n)$, and thus for a moderately high sample size n , the quantity $\sqrt{n(1-\alpha_n)}/\log[(1-\alpha_n)/(1-\beta_n)]$ representing the rate of convergence of the extrapolation factor may only be slightly lower than the quantity $\sqrt{n(1-\alpha_n)}$ representing the rate of convergence of the estimator at the intermediate step. Hence the idea that, while for n very large the difference in finite-sample behaviour between any two estimators of the tail L^p -median at the basic intermediate level α_n will be eventually wiped out by the performance of the estimator $\hat{\gamma}_n$, there may still be a significant

impact of the quality of the intermediate tail L^p -median estimator used on the overall accuracy of the extrapolated estimator when n is moderately large. This is illustrated in the simulation study below.

5 Simulation study

Our goal in the present section is to assess the finite-sample performance of our direct and indirect estimators of an extreme tail L^p -median, for $p \in [1, 2]$. In addition, we shall do so in a way that provides guidance as to how an extreme tail L^p -median, and its estimates, can be used and interpreted in practical setups. Let us recall that our focus is not to consider cases with low γ , as in such cases the easily interpretable CTE risk measure can be used and estimated with good accuracy, including at extreme levels. We shall rather consider cases with $\gamma > 1/4$, where the fact that the tail extreme tail L^p -median realises a compromise between the robust MS and the sensitive CTE should be expected to result in estimators with an improved finite-sample performance compared to that of the classical empirical CTE estimator. It was actually highlighted in (3) and (4) that, for $p \in [1, 2]$ and $\gamma < 1/(p - 1)$, an extreme tail L^p -median $m_p(\alpha)$ can be understood asymptotically as a weighted average of $\text{MS}(\alpha)$ and $\text{CTE}(\alpha)$. In other words, defining the interpolating risk measure

$$R_\lambda(\alpha) := \lambda \text{MS}(\alpha) + (1 - \lambda) \text{CTE}(\alpha),$$

we have

$$m_p(\alpha) \approx R_\lambda(\alpha) \quad \text{as } \alpha \rightarrow 1$$

with $\lambda = \lambda(p, \gamma)$ as in (4). It then turns out that at the population level, we have two distinct (but asymptotically equivalent) possibilities to interpolate, and thus create a compromise, between extreme Median Shortfall and extreme Conditional Tail Expectation:

- Consider the family of measures $m_p(\alpha)$, $p \in [1, 2]$;
- Consider the family of measures $R_\lambda(\alpha)$, $\lambda \in [0, 1]$.

In the rest of this section, based on a sample of data (X_1, \dots, X_n) of size n , we consider the estimation of the tail L^p -median $m_p(\alpha_n)$ for both intermediate and extreme levels α_n , and how this estimation compares with direct estimation of the interpolating measure $R_\lambda(\alpha_n)$.

5.1 Intermediate case

We first investigate the estimation of $m_p(\alpha_n)$, for $p \in [1, 2]$, and of $R_\lambda(\alpha_n)$, for $\lambda \in [0, 1]$, in the intermediate case when $\alpha_n \rightarrow 1$ and $n(1 - \alpha_n) \rightarrow \infty$. As far as the estimation of $m_p(\alpha_n)$

is concerned, we compare the direct estimator $\widehat{m}_p(\alpha_n)$ defined in (6) and the indirect estimator in (8). In the latter, $\widehat{\gamma}_n$ is taken as the Hill estimator defined in (9): in other words,

$$\begin{aligned}\widetilde{m}_p(\alpha_n) &= \frac{X_{n-\lfloor n(1-\alpha_n) \rfloor, n}}{\kappa(p, \widehat{\gamma}_n^H(\lfloor n(1-\alpha_n) \rfloor))} \\ \text{with } \widehat{\gamma}_n^H(k) &= \frac{1}{k} \sum_{i=1}^k \log(X_{n-i+1, n}) - \log(X_{n-k, n}).\end{aligned}$$

To further compare the performance of these two estimators, and therefore the practical applicability of the measure $m_p(\alpha_n)$ for interpolating between extreme MS and extreme CTE, the finite sample behaviour of these estimators are compared to that of the estimator of $R_\lambda(\alpha_n)$ given by

$$\begin{aligned}\widehat{R}_\lambda(\alpha_n) &:= \lambda X_{n-\lfloor n(1-\alpha_n)/2 \rfloor, n} + (1-\lambda) \widehat{\text{CTE}}(\alpha_n) \\ \text{with } \widehat{\text{CTE}}(\alpha_n) &= \widehat{m}_2(\alpha_n) = \frac{1}{\lfloor n(1-\alpha_n) \rfloor} \sum_{i=1}^{\lfloor n(1-\alpha_n) \rfloor} X_{n-i+1, n}.\end{aligned}$$

In order to be able to compare this estimator of $R_\lambda(\alpha_n)$ to our estimators of the tail L^p -median $m_p(\alpha_n)$, we take $\lambda = \lambda(p, \gamma)$ as in (4). This results in $m_p(\alpha_n) \approx R_{\lambda(p, \gamma)}(\alpha_n)$ as $n \rightarrow \infty$, and we can then compare the estimators $\widehat{m}_p(\alpha_n)$, $\widetilde{m}_p(\alpha_n)$ and $\widehat{R}_{\lambda(p, \gamma)}(\alpha_n)$ on the range $p \in [1, 2]$.

We do so on $N = 500$ simulated random samples of size $n = 500$, with $\alpha_n = 1 - 75/n = 0.85$. Two distributions satisfying condition $\mathcal{C}_2(\gamma, \rho, \mathbf{A})$ are considered:

- The Burr distribution having distribution function $F(x) = 1 - (1 + x^{3/(2\gamma)})^{-2/3}$ on $(0, \infty)$ (here $\gamma > 0$). This distribution has tail index γ and second-order parameter $\rho = -3/2$.
- The Student distribution with $1/\gamma$ degrees of freedom, where $\gamma > 0$. This distribution has tail index γ and second-order parameter $\rho = -2\gamma$.

For each of these two distributions, we examine the cases $\gamma \in \{1/4, 1/2, 3/4\}$, corresponding to, respectively, a borderline case for finite fourth moment, a borderline case for finite variance, and a case where there is no finite variance. The accuracy of each of the three estimators $\widehat{m}_p(\alpha_n)$, $\widetilde{m}_p(\alpha_n)$ and $\widehat{R}_{\lambda(p, \gamma)}(\alpha_n)$ is measured by their respective empirical Mean Log-Squared Error (MLSE) on the N simulated samples. This is defined by

$$\text{MLSE}(\widehat{m}_p(\alpha_n)) = \frac{1}{N} \sum_{j=1}^N \log^2 \left(\frac{\widehat{m}_p^{(j)}(\alpha_n)}{m_p(\alpha_n)} \right)$$

where $\widehat{m}_p^{(j)}(\alpha_n)$ is the estimate of $m_p(\alpha_n)$ calculated on the j th sample, and similarly for $\widetilde{m}_p(\alpha_n)$ and $\widehat{R}_{\lambda(p, \gamma)}(\alpha_n)$ (in the latter, we replace $m_p(\alpha_n)$ by $R_{\lambda(p, \gamma)}(\alpha_n)$). These quantities are represented, as functions of $p \in [1, 2]$, in the top panels of Figures 6–8.

It can be seen from these graphs that, on the Burr distribution, the indirect estimator $\tilde{m}_p(\alpha_n)$ has a lower MLSE than the direct estimator, except for γ close to 1 and p close to 2. This instability is very likely due to the fact that, for p close to 2, we have $1/\kappa(p, \gamma) \approx 1/(1 - \gamma)$, and thus if γ is also close to 1, the quantity $1/\kappa(p, \hat{\gamma}_n^H(\lfloor n(1 - \alpha_n) \rfloor))$ appearing in the estimator $\tilde{m}_p(\alpha_n)$ will be a highly unstable estimator of $1/\kappa(p, \gamma)$. By contrast, the direct estimator is generally more accurate than the indirect one when the underlying distribution is a Student distribution, especially for $p \geq 1.6$. Furthermore, it should be noted that the direct estimator $\hat{m}_p(\alpha_n)$ performs overall noticeably better than the estimator $\hat{R}_{\lambda(p, \gamma)}(\alpha_n)$. This confirms our theoretical expectation that estimation of $m_p(\alpha_n)$ should be easier than estimation of $R_{\lambda(p, \gamma)}(\alpha_n)$, since $\hat{m}_p(\alpha_n)$ is relatively $\sqrt{n(1 - \tau_n)}$ -consistent as soon as $\gamma < 1/[2(p-1)]$, which is a weaker condition than the assumption $\gamma < 1/2$ needed to ensure the relative $\sqrt{n(1 - \tau_n)}$ -consistency of $\hat{R}_{\lambda(p, \gamma)}(\alpha_n)$ (due to its reliance on the estimator $\widehat{\text{CTE}}(\alpha_n)$). In other words, on finite samples and at intermediate levels, it is preferable to interpolate between extreme MS and extreme CTE *via* tail L^p -medians rather than using direct linear interpolation, even though these two ideas are asymptotically equivalent at the population level.

5.2 Extreme case

We now focus on the estimation of $m_p(\beta_n)$ and $R_\lambda(\beta_n)$ for a proper extreme level $\beta_n \rightarrow 1$, such that $n(1 - \beta_n) \rightarrow c \in [0, \infty)$. The estimators we consider are, for the estimation of the extreme tail L^p -median $m_p(\beta_n)$, the two extrapolated estimators defined in Section 4.3: first, the extrapolated direct estimator, given by

$$\hat{m}_p^W(\beta_n, k) := \hat{m}_p^W(\beta_n) = \left(\frac{n(1 - \beta_n)}{k} \right)^{-\hat{\gamma}_n^H(k)} \hat{m}_p(1 - k/n),$$

where $k \in \{1, \dots, n-1\}$, and, second, the extrapolated indirect estimator

$$\tilde{m}_p^W(\beta_n, k) := \tilde{m}_p^W(\beta_n) = \left(\frac{n(1 - \beta_n)}{k} \right)^{-\hat{\gamma}_n^H(k)} \tilde{m}_p(1 - k/n).$$

Recall that here, $\hat{\gamma}_n^H(k)$ denotes the Hill estimator introduced in (9). These two estimators are compared to the following extrapolated version of the estimator \hat{R}_λ :

$$\hat{R}_\lambda^W(\beta_n, k) := \left(\frac{n(1 - \beta_n)}{k} \right)^{-\hat{\gamma}_n^H(k)} \left[\lambda X_{n - \lfloor k/2 \rfloor, n} + (1 - \lambda) \widehat{\text{CTE}}(1 - k/n) \right].$$

We take again $\lambda = \lambda(p, \gamma)$ so that the estimators $\hat{m}_p^W(\beta_n, k)$, $\tilde{m}_p^W(\beta_n, k)$ and $\hat{R}_\lambda^W(\beta_n, k)$ can be compared on the range $p \in [1, 2]$. In what follows, we also take $\beta_n = 1 - 1/n$ in all three estimators. These estimators depend on a tuning parameter k , which we take as

$$\hat{k}_{\text{opt}} := \left[\frac{1}{J} \sum_{j=1}^J \arg \min_{k \in \{4, \dots, \lfloor n/4 \rfloor\}} \int_{1-k/n}^{1-k/(4n)} \log^2 \left(\frac{\hat{m}_{p_j}^W(\alpha, k)}{\hat{m}_{p_j}(\alpha)} \right) d\alpha \right],$$

where $J \in \mathbb{N} \setminus \{0\}$ and $1 = p_1 < \dots < p_J = 2$. The idea behind this criterion is that for an intermediate order α , the empirical and extrapolated estimators should both be able to estimate accurately $m_p(\alpha)$, across the range $p \in [1, 2]$, and therefore the distance between these two estimators should be small at intermediate levels provided the parameter k is chosen properly. A related selection rule is used by [26] for extreme quantile estimation. We take here $J = 4$ and $p_j = (j + 2)/3$. Simulation settings (sample size, distributions, ...) are the same as those of the intermediate case; empirical MLSEs of the estimators $\widehat{m}_p^W(\beta_n, \widehat{k}_{\text{opt}})$, $\widetilde{m}_p^W(\beta_n, \widehat{k}_{\text{opt}})$ and $\widehat{R}_{\lambda(p,\gamma)}^W(\beta_n, \widehat{k}_{\text{opt}})$ are represented, as functions of $p \in [1, 2]$, in the bottom panels of Figures 6–8.

These graphs show that the conclusions reached in the intermediate case remain true in the extreme case $\beta_n = 1 - 1/n$: on the Burr distribution, the extrapolated indirect estimator of $m_p(\beta_n)$ performs better than the extrapolated direct estimator, except for γ close to 1 and p close to 2. The reverse conclusion holds true on the Student distribution. The extrapolated direct estimator $\widehat{m}_p^W(\beta_n, \widehat{k}_{\text{opt}})$ also performs generally better than the extrapolated linear interpolation estimator $\widehat{R}_{\lambda(p,\gamma)}^W(\beta_n, \widehat{k}_{\text{opt}})$. There is a noticeable improvement for $p \in [1.25, 1.75]$ in the case $\gamma = 1/2$ and even more so for $\gamma = 3/4$, which are the most relevant cases for our purpose. The accuracy of $\widehat{m}_p^W(\beta_n, \widehat{k}_{\text{opt}})$ is also comparable overall to that of $\widehat{R}_{\lambda(p,\gamma)}^W(\beta_n, \widehat{k}_{\text{opt}})$ for $\gamma = 1/4$, although this is not the case we originally constructed the tail L^p –median for.

As a conclusion, this simulation study indicates that, whether at intermediate or proper extreme levels, our tail L^p –median estimators provide a way of interpolating between extreme MS and extreme CTE that is more accurate in practice than a simple linear interpolation. This appears to be true on a wide range of values of γ , allowing for a flexible use of the class of tail L^p –medians in practice, although the improvement is clearer for $\gamma \geq 1/2$, where it is known that CTE estimation is a difficult problem.

6 Real data analysis

The data set we consider is made of $n = 1098$ commercial fire losses recorded between 1st January 1995 and 31st December 1996 by the FFSA (an acronym for the *Fédération Française des Sociétés d'Assurance*). This data set is available from the R package `CASdatasets` as `data(frecomfire)`. The data is converted into euros from French francs, and denoted by (X_1, \dots, X_n) . Insurance and financial companies have a strong interest in the analysis of this type of data set. For example, extremely high losses have to be taken into account in order to estimate, at company level, the capital requirements that have to be put in place so as to survive the upcoming calendar year with a probability not less than 0.995, as part of compliance with the Solvency II directive. At the same time, it is in the interest of companies to carry out a balanced assessment of risk: an underestimation of risk could threaten the company's survival, but an overestimation may lead

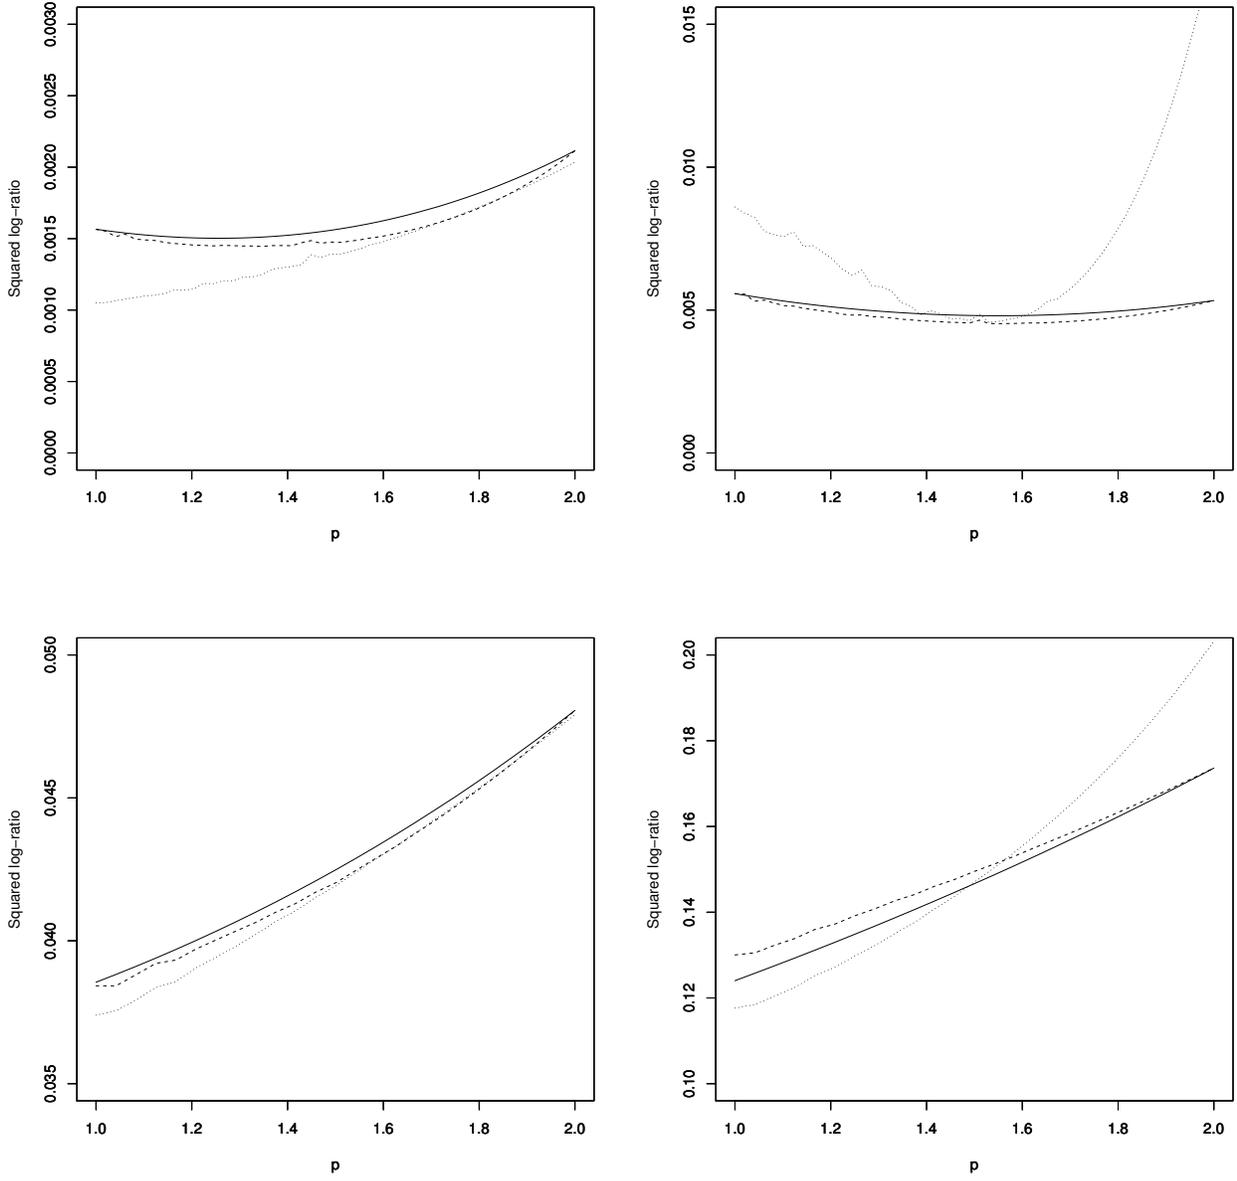


Figure 6: Simulation study, case $\gamma = 1/4$. **Top panels:** empirical MLSE of the estimator $\hat{R}_{\lambda(p,\gamma)}(\alpha_n)$ estimator (full line) and of the two tail L^p -median estimators $\hat{m}_p(\alpha_n)$ (dashed line) and $\tilde{m}_p(\alpha_n)$ (dotted line). **Bottom panels:** empirical MLSE of the extrapolated estimator $\hat{R}_{\lambda(p,\gamma)}^W(\beta_n, \hat{k}_{\text{opt}})$ (full line) and of the two extrapolated tail L^p -median estimators $\hat{m}_p^W(\beta_n, \hat{k}_{\text{opt}})$ (dashed line) and $\tilde{m}_p^W(\beta_n, \hat{k}_{\text{opt}})$ (dotted line). Left: Burr distribution, right: Student distribution.

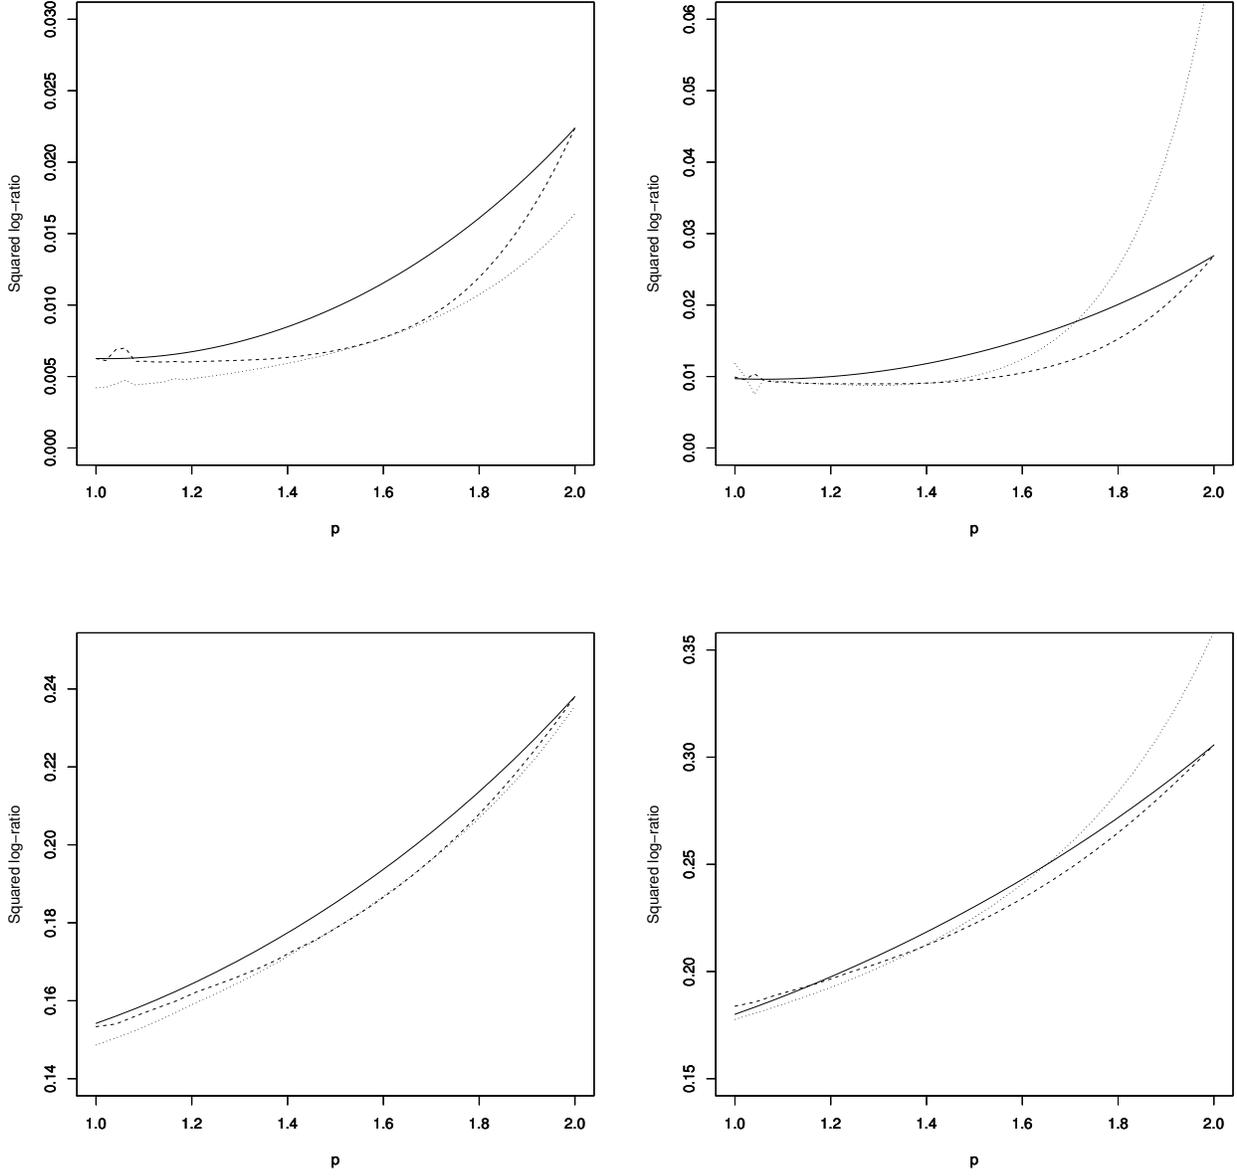


Figure 7: Simulation study, case $\gamma = 1/2$. **Top panels:** empirical MLSE of the estimator $\hat{R}_{\lambda(p,\gamma)}(\alpha_n)$ estimator (full line) and of the two tail L^p -median estimators $\hat{m}_p(\alpha_n)$ (dashed line) and $\tilde{m}_p(\alpha_n)$ (dotted line). **Bottom panels:** empirical MLSE of the extrapolated estimator $\hat{R}_{\lambda(p,\gamma)}^W(\beta_n, \hat{k}_{\text{opt}})$ (full line) and of the two extrapolated tail L^p -median estimators $\hat{m}_p^W(\beta_n, \hat{k}_{\text{opt}})$ (dashed line) and $\tilde{m}_p^W(\beta_n, \hat{k}_{\text{opt}})$ (dotted line). Left: Burr distribution, right: Student distribution.

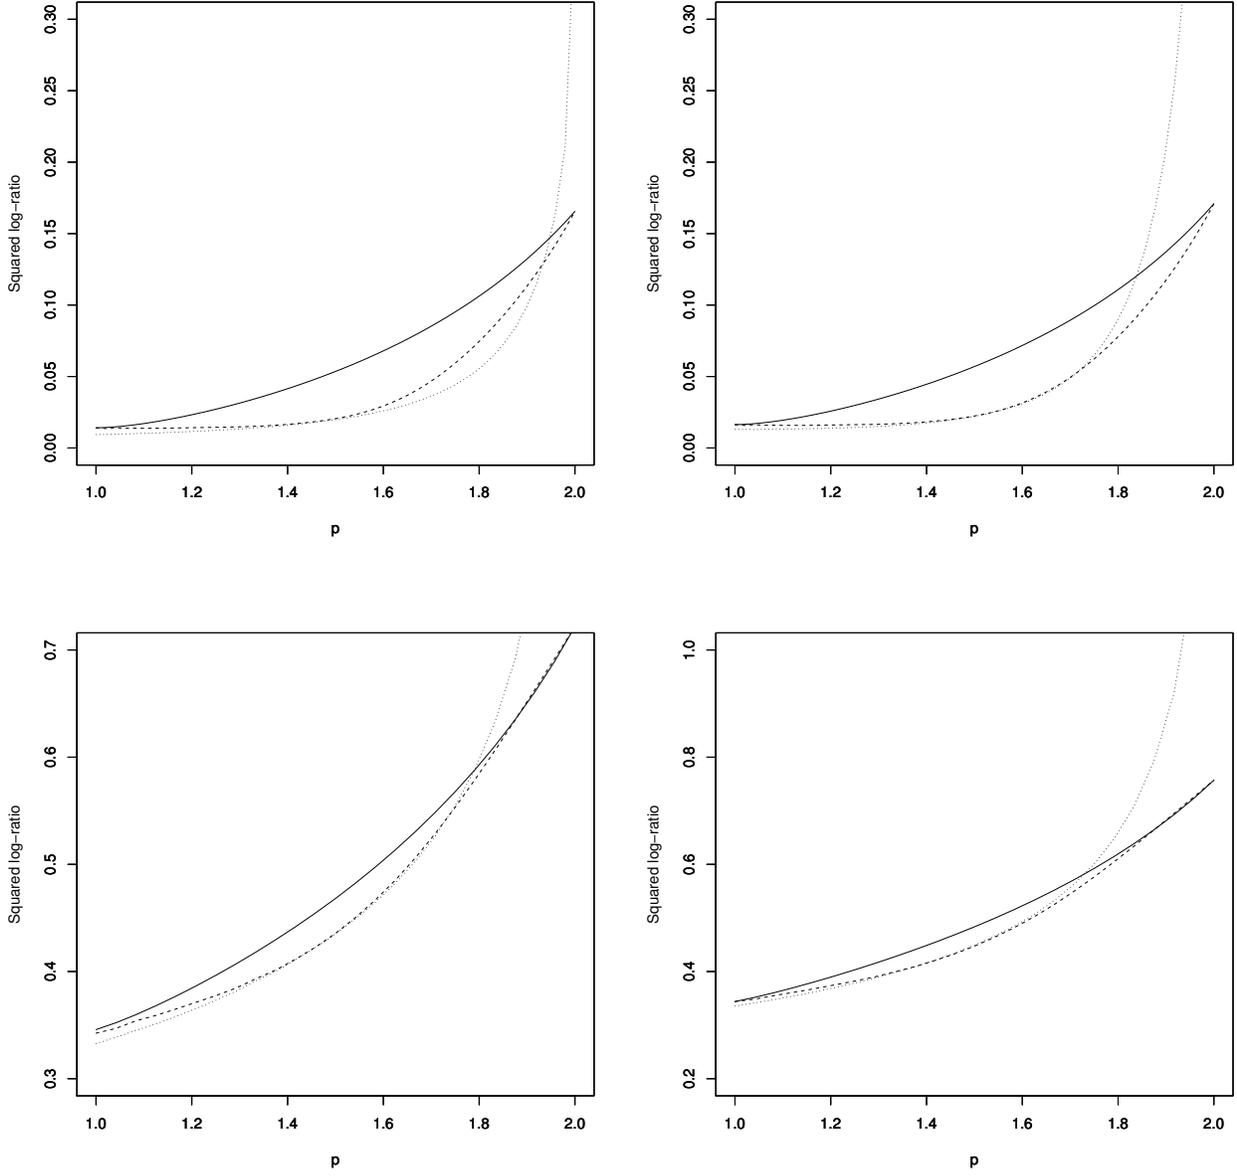


Figure 8: Simulation study, case $\gamma = 3/4$. **Top panels:** empirical MLSE of the estimator $\hat{R}_{\lambda(p,\gamma)}(\alpha_n)$ (full line) and of the two tail L^p -median estimators $\hat{m}_p(\alpha_n)$ (dashed line) and $\tilde{m}_p(\alpha_n)$ (dotted line). **Bottom panels:** empirical MLSE of the extrapolated estimator $\hat{R}_{\lambda(p,\gamma)}^W(\beta_n, \hat{k}_{\text{opt}})$ (full line) and of the two extrapolated tail L^p -median estimators $\hat{m}_p^W(\beta_n, \hat{k}_{\text{opt}})$ (dashed line) and $\tilde{m}_p^W(\beta_n, \hat{k}_{\text{opt}})$ (dotted line). Left: Burr distribution, right: Student distribution.

the insurer to ask for higher premiums and deductibles, reducing the attractiveness of its products to the consumer, thus negatively affecting the company's competitiveness on the market. A single quantile, as in the above Solvency II compliance example, cannot account for a detailed picture of risk, which is why we investigate here the use of the alternative class of tail L^p -medians.

Our first step in the analysis of the extreme losses in this data set is to estimate the tail index γ . To this end, the procedure of choice of \hat{k}_{opt} outlined in Section 5.2 is used: the value of k returned by this procedure is $\hat{k}_{\text{opt}} = 64$, and the corresponding estimate of γ provided by the Hill estimator is $\hat{\gamma}^H = 0.67$. This estimated value of γ is in line with the findings of [23], and it suggests that there is evidence for an infinite second moment of the underlying distribution. We know, in this context, that the estimation of the extreme CTE is going to be a difficult problem, although it would give a better understanding of risk in this data set than a single quantile such as the VaR or MS would do. It therefore makes sense, on this data set, to use the class of tail L^p -medians to find a middle ground between MS and CTE estimation by interpolation.

Estimates of the tail L^p -median $m_p(1 - 1/n)$, obtained through our extrapolated direct and extrapolated indirect estimators, are depicted on Figure 9. These estimates are fairly close overall on the range $p \in [1, 2]$. There is a difference for p close to 2, where the increased theoretical sensitivity of the direct estimator to the highest values in the sample makes it exceed the more robust indirect estimator; note that here γ is estimated to be 0.67, which is sufficiently far from 1 to ensure that the instability of the indirect estimator for very high γ is not an issue, and the estimate returned by the indirect extrapolated method can be considered to be a reasonable one.

With these estimates at hand comes the question of how to choose p . This is of course a difficult question that depends on the objective of the analysis from the perspective of risk assessment: should the analysis be conservative (*i.e.* return a higher, more prudent estimation) or not? We do not wish to enter into such a debate here, which would be better held within the financial and actuarial communities. Let us, however, illustrate a simple way to choose p based on our asymptotic interpretation of a tail L^p -median as a weighted average of MS and CTE. Recall from (3) and (4) that

$$m_p(1 - 1/n) \approx \lambda(p, \gamma)\text{MS}(1 - 1/n) + (1 - \lambda(p, \gamma))\text{CTE}(1 - 1/n)$$

$$\text{with } \lambda(p, \gamma) = \frac{1 - (1 - \gamma)/\kappa(p, \gamma)}{1 - 2^\gamma(1 - \gamma)}.$$

We also know from our simulation study that it is generally more accurate to estimate $m_p(1 - 1/n)$ rather than the corresponding linear combination of $\text{MS}(1 - 1/n)$ and $\text{CTE}(1 - 1/n)$. This suggests the following procedure: if a weighting λ_0 , specifying a compromise between robustness of MS and sensitivity of CTE, is fixed in advance, and γ is estimated by $\hat{\gamma}$, one can choose $p = \hat{p}$ as the unique

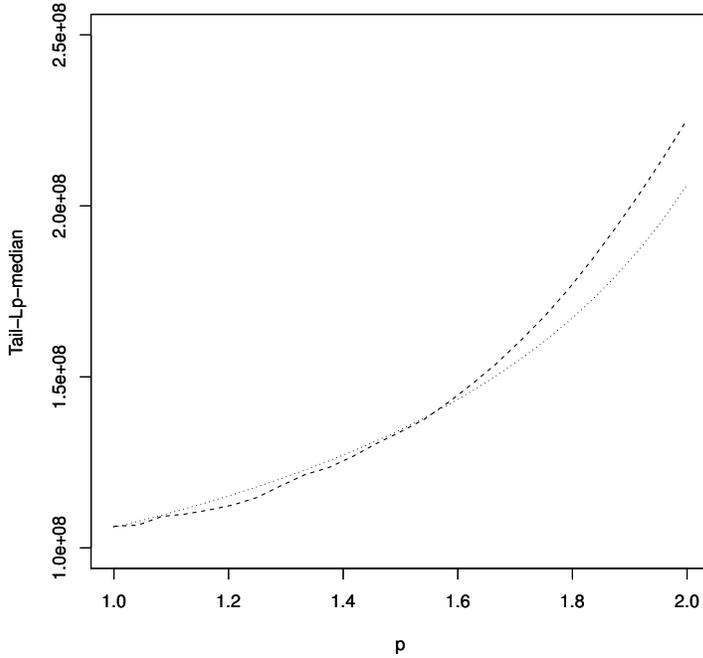


Figure 9: Real fire insurance data set. Extrapolated tail L^p -median estimators $\widehat{m}_p^W(1 - 1/n, \hat{k}_{\text{opt}})$ (dashed line) and $\widetilde{m}_p^W(1 - 1/n, \hat{k}_{\text{opt}})$ (dotted line) as functions of $p \in [1, 2]$.

root of the equation $\lambda(\widehat{p}, \widehat{\gamma}) = \lambda_0$. This results in an estimator of the linear combination

$$\lambda_0 \text{MS}(1 - 1/n) + (1 - \lambda_0) \text{CTE}(1 - 1/n).$$

With $\lambda_0 = 1/2$, representing the simple average between $\text{MS}(1 - 1/n)$ and $\text{CTE}(1 - 1/n)$, this method yields $\widehat{p} = 1.711$. The corresponding estimates, in million euros, are $\widehat{m}_p^W(1 - 1/n) = 160.8$ and $\widetilde{m}_p^W(1 - 1/n) = 155.4$. It is interesting to note that, although these quantities estimate the average of $\text{MS}(1 - 1/n) = m_1(1 - 1/n)$ and $\text{CTE}(1 - 1/n) = m_2(1 - 1/n)$, we also have

$$\widehat{m}_1^W(1 - 1/n) = 106.3 \quad \text{and} \quad \widehat{m}_2^W(1 - 1/n) = 225.2$$

so that

$$\widetilde{m}_p^W(1 - 1/n) = 155.4 < \widehat{m}_p^W(1 - 1/n) = 160.8 < \frac{1}{2} [\widehat{m}_1^W(1 - 1/n) + \widehat{m}_2^W(1 - 1/n)] = 165.8.$$

The estimate $\widehat{m}_p^W(1 - 1/n)$ of the simple average between $\text{MS}(1 - 1/n)$ and $\text{CTE}(1 - 1/n)$ obtained through extrapolating the direct tail L^p -median estimator is therefore itself a midway between the indirect estimator $\widetilde{m}_p^W(1 - 1/n)$, which relies on the VaR estimator $\widehat{q}(1 - 1/n)$, and the direct estimator of this average, which depends on the highly variable estimator $\widehat{\text{CTE}}(1 - 1/n)$.

Supplementary Material document

The online Supplementary Material document contains the proofs of our main results, as well as all auxiliary results and their proofs.

References

- [1] Abramovitz, M. and Stegun, I.A. (1972). *Handbook of Mathematical Functions*, National Bureau of Standards Applied Mathematics Series.
- [2] Acerbi, C. and Tasche, D. (2002). On the coherence of expected shortfall. *J. Banking & Finance* **26**, 1487–1503.
- [3] Acharya, V.V., Pedersen, L.H., Philippon, T. and Richardson, M. (2012). Measuring systemic risk. *Discussion Paper DP8824*, Centre for Economic Policy Research, London.
- [4] Artzner, P., Delbaen, F., Eber, J.-M. and Heath, D. (1999). Coherent measures of risk. *Math. Finance* **9**, 203–228.
- [5] Basel Committee on Banking Supervision (2013). *Fundamental review of the trading book: A revised market risk framework*.
- [6] Beirlant, J., Goegebeur, Y., Segers, J. and Teugels, J. (2004). *Statistics of Extremes: Theory and Applications*, Wiley.
- [7] Belles-Sampera, J., Guillén, M. and Santolino, M. (2014). GlueVaR risk measures in capital allocation applications. *Insurance Math. Econom.* **58**, 132–137.
- [8] Bingham, N.H., Goldie, C.M. and Teugels, J.L. (1987). *Regular Variation*, Cambridge University Press.
- [9] Borisov, I.S. and Baklanov, E.A. (2001). Probability inequalities for generalized L -statistics. *Sib. Math. J.* **42**, 217–231.
- [10] Brazauskas, V., Jones, B.L., Puri, M.L. and Zitikis, R. (2008). Estimating conditional tail expectation with actuarial applications in view. *J. Statist. Plann. Inference* **138**, 3590–3604.
- [11] Cai, J., Einmahl, J., de Haan, L. and Zhou, C. (2015). Estimation of the marginal expected shortfall: the mean when a related variable is extreme. *J. Roy. Statist. Soc. Ser. B* **77**, 417–442.
- [12] Chavez-Demoulin, V., Embrechts, P. and Sardy, S. (2014). Extreme-quantile tracking for financial time series. *J. Econometrics* **181**, 44–52.

- [13] Chen, Z. (1996). Conditional L_p -quantiles and their application to testing of symmetry in non-parametric regression. *Statist. Probab. Lett.* **29**, 107–115.
- [14] Daouia, A., Girard, S. and Stupfler, G. (2017). Extreme M-quantiles as risk measures: From L^1 to L^p optimization. *Bernoulli*, to appear.
- [15] Daouia, A., Girard, S. and Stupfler, G. (2018). Estimation of tail risk based on extreme expectiles. *J. Roy. Statist. Soc. Ser. B* **80**, 263–292.
- [16] Daniélsson, J., Jorgensen, B.N., Samorodnitsky, G., Sarma, M. and de Vries, C.G. (2013). Fat tails, VaR and subadditivity. *J. Econometrics* **172**, 283–291.
- [17] Diebolt, J., Guillou, A. and Rached, I. (2007). Approximation of the distribution of excesses through a generalized probability-weighted moments method. *J. Statist. Plann. Inference* **137**, 841–857.
- [18] Drees, H., Ferreira, A. and de Haan, L. (2004). On maximum likelihood estimation of the extreme value index. *Ann. Appl. Probab.* **14**, 1179–1201.
- [19] Embrechts, P., Klüppelberg, C. and Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance*, Springer.
- [20] El Methni, J., Gardes, L. and Girard, S. (2014). Nonparametric estimation of extreme risks from conditional heavy-tailed distributions. *Scand. J. Stat.* **41**, 988–1012.
- [21] El Methni, J., Gardes, L. and Girard, S. (2018). Kernel estimation of extreme regression risk measures. *Electron. J. Stat.* **12**, 359–398.
- [22] El Methni, J. and Stupfler, G. (2017a). Extreme versions of Wang risk measures and their estimation for heavy-tailed distributions. *Stat. Sinica* **27**, 907–930.
- [23] El Methni, J. and Stupfler, G. (2017b). Improved estimators of extreme Wang distortion risk measures for very heavy-tailed distributions, *Econometrics and Statistics*, to appear.
- [24] Embrechts, P., Lambrigger, D.P. and Wüthrich, M.V. (2009). Multivariate extremes and the aggregation of dependent risks: examples and counter-examples. *Extremes* **12**, 107–127.
- [25] Emmer, S., Kratz, M. and Tasche, D. (2015). What is the best risk measure in practice? A comparison of standard measures. *J. Risk* **18**, 31–60.
- [26] Gardes, L. and Stupfler, G. (2015). Estimating extreme quantiles under random truncation. *TEST* **24**, 207–227.

- [27] Geyer, C.J. (1996). On the asymptotics of convex stochastic optimization, unpublished manuscript.
- [28] Gneiting, T. (2011). Making and evaluating point forecasts. *J. Am. Statist. Assoc.* **106**, 746–762.
- [29] Gong, J., Li, Y., Peng, L. and Yao, Q. (2015). Estimation of extreme quantiles for functions of dependent random variables. *J. Roy. Statist. Soc. Ser. B* **77**, 1001–1024.
- [30] de Haan, L. and Ferreira, A. (2006). *Extreme Value Theory: An Introduction*, Springer.
- [31] Hill, B.M. (1975). A simple general approach to inference about the tail of a distribution. *Ann. Statist.* **3**, 1163–1174.
- [32] Hill, J.B. (2013). Least tail-trimmed squares for infinite variance autoregressions. *J. Time Series Anal.* **34**, 168–186.
- [33] Hill, J.B. (2015). Expected shortfall estimation and Gaussian inference for infinite variance time series. *J. Financ. Econom.* **13**, 1–44.
- [34] Hosking, J.R.M., Wallis, J.R. and Wood, E.F. (1987). Estimation of the generalized extreme value distribution by the method of probability-weighted moments. *Technometrics* **27**, 251–261.
- [35] Hua, L. and Joe, H. (2011). Second order regular variation and conditional tail expectation of multiple risks. *Insurance Math. Econom.* **49**, 537–546.
- [36] International Monetary Fund (2014). *Canada : Financial Sector Assessment Program—Insurance Core Principles—Detailed Assessment of Observance*, Washington, D.C.
- [37] Knight, K. (1998). Limiting distributions for L_1 regression estimators under general conditions. *Ann. Statist.* **26**, 755–770.
- [38] Knight, K. (1999). Epi-convergence in distribution and stochastic equi-semicontinuity, technical report, University of Toronto. Available at <http://www.utstat.toronto.edu/keith/EPFL/equi.ps>.
- [39] Koenker, R. (2005). *Quantile Regression*, Cambridge University Press.
- [40] Koenker, R. and Bassett, G.S. (1978). Regression quantiles. *Econometrica* **46**, 33–50.
- [41] Kou, S., Peng, X. and Heyde, C.C. (2013). External risk measures and Basel accords. *Math. Oper. Res.* **38**, 393–417.

- [42] Kou, S. and Peng, X. (2016). On the measurement of economic tail risk. *Oper. Res.* **64**, 1056–1072.
- [43] Ling, S. (2005). Self-weighted least absolute deviation estimation for infinite variance autoregressive models. *J. Roy. Statist. Soc. Ser. B* **67**, 381–393.
- [44] Newey, W.K. and Powell, J.L. (1987). Asymmetric least squares estimation and testing. *Econometrica* **55**, 819–847.
- [45] Organisation for Economic Co-operation and Development (2016). *Life Annuity Products and their Guarantees*, OECD Publishing, Paris.
- [46] Pickands, J. (1975). Statistical inference using extreme order statistics. *Ann. Statist.* **3**, 119–131.
- [47] Resnick, S. (2007). *Heavy-Tail Phenomena: Probabilistic and Statistical Modeling*, Springer.
- [48] Rockafellar, R.T. and Uryasev, S. (2002). Conditional Value-at-Risk for general loss distributions. *J. Banking & Finance* **26**, 1443–1471.
- [49] Smith, R.L. (1987). Estimating tails of probability distributions. *Ann. Statist.* **15**, 1174–1207.
- [50] Tasche, D. (2002). Expected shortfall and beyond. *J. Banking & Finance* **26**, 1519–1533.
- [51] Tasche, D. (2008). Capital allocation to business units and sub-portfolios: the Euler principle. In: Resti, A. (ed.), *Pillar II in the New Basel Accord: The Challenge of Economic Capital*, Risk Books, 423–453.
- [52] Wang, S. (1996). Premium calculation by transforming the layer premium density. *ASTIN Bull.* **1**, 71–92.
- [53] Wang, S. (2004). CAT bond pricing using probability transforms. In *Geneva Papers, special issue on insurance and the state of the art in CAT bond pricing* **278**, 19–29.
- [54] Wang, S., Young, V. and Panjer, H. (1997). Axiomatic characterization of insurance prices. *Insurance Math. Econom.* **2**, 173–183.
- [55] Weissman, I. (1978). Estimation of parameters and large quantiles based on the k largest observations. *J. Am. Statist. Assoc.* **73**, 812–815.
- [56] Wüthrich, M. and Merz, M. (2013). *Financial Modeling, Actuarial Valuation and Solvency in Insurance*, Springer.
- [57] Yang, F. (2015). First- and second-order asymptotics for the tail distortion risk measure of extreme risks. *Commun. Stat. Theory Methods* **44**, 520–532.

- [58] Zhu, L. and Li, H. (2012). Tail distortion risk and its asymptotic analysis. *Insurance Math. Econom.* **51**, 115–121.