



HAL
open science

Extreme expectile estimation for heavy-tailed time series

Simone A. Padoan, Gilles Stupfler

► **To cite this version:**

Simone A. Padoan, Gilles Stupfler. Extreme expectile estimation for heavy-tailed time series. 2020. hal-02541663v1

HAL Id: hal-02541663

<https://hal.science/hal-02541663v1>

Preprint submitted on 14 Apr 2020 (v1), last revised 6 Apr 2023 (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.

Extreme expectile estimation for heavy-tailed time series

Simone A. Padoan* and Gilles Stupfler†

April 14, 2020

Abstract

Expectiles define a least squares analogue of quantiles. They have lately received substantial attention in actuarial and financial risk management contexts. Unlike quantiles, expectiles define coherent risk measures and are determined by tail expectations rather than tail probabilities; unlike the popular Expected Shortfall, they define elicitable risk measures. This has motivated the study of the behaviour and estimation of extreme expectiles in some of the recent statistical literature. The case of stationary but weakly dependent observations has, however, been left largely untouched, even though correctly accounting for the uncertainty present in typical financial applications requires the consideration of dependent data. We investigate here the theoretical and practical behaviour of two classes of extreme expectile estimators in a strictly stationary β -mixing context, containing the classes of ARMA, ARCH and GARCH models with heavy-tailed innovations that are of interest in financial applications. We put a particular emphasis on the construction of asymptotic confidence intervals adapted to the dependence framework, whose performance we contrast with that of the naive intervals obtained from the theory of independent and identically distributed data. The methods are showcased in a numerical simulation study and on real financial data.

Keywords: Asymmetric least squares, Expectiles, Extrapolation, Extremal dependence, Heavy-tailed distribution, Mixing, Tail copula, Weak dependence.

1 Introduction

A major problem in econometrics and statistical finance is to quantify the risk associated to a real-valued profit-loss variable X . The class of quantiles, which plays an essential role in extreme value theory, is one of the basic tools in risk management. Arguably the best-known quantile-based risk measure in banking, and the financial sector generally, is Value-at-Risk

*Simone A. Padoan works at the Department of Decision Sciences, Bocconi University of Milan, via Roentgen 1, 20136 Milano, Italy, E-mail: simone.padoan@unibocconi.it.

†Gilles Stupfler works at Univ Rennes, Ensai, CNRS, CREST - UMR 9194, F-35000 Rennes, France, E-mail: gilles.stupfler@ensai.fr.

(VaR) at a confidence level $\tau \in (0, 1)$, defined as the negative τ th quantile $-q(\tau)$ of the real-valued profit-loss distribution, with τ being close to zero representing the situations carrying the greatest risk. The quantile can be obtained by minimising asymmetrically weighted mean absolute deviations (Koenker and Bassett, 1978):

$$q_\tau \in \arg \min_{q \in \mathbb{R}} \mathbb{E}(\rho_\tau(X - q) - \rho_\tau(X)), \quad (1)$$

where $\rho_\tau(x) = |\tau - \mathbb{1}\{x \leq 0\}||x|$ is the so-called quantile check function and $\mathbb{1}\{\cdot\}$ the indicator function. This property is important in the context of evaluation and validation of VaR forecasts, since it corresponds to the existence of a natural backtesting methodology. The latter is formalised in Gneiting (2011), who introduced the general notion of elicibility for a functional defined through the minimisation of a suitable loss function. The relevance of this elicibility property for the backtesting problem has been extensively discussed, including recently in Bellini and Di Bernardino (2017). The VaR, however, suffers from certain serious drawbacks. It does not, in general, define a coherent risk measure in the sense of the influential paper by Artzner et al. (1999), because it is not necessarily subadditive. This means that it does not abide by the intuitive diversification principle. From the point of view of extreme value theory, quantiles are often criticised for missing out on important information about the tail of the underlying distribution, since they only depend on the frequency of tail losses and not on their actual values. Unlike the VaR, the most popular quantile-based risk measure, called Expected Shortfall (ES, or Conditional VaR), is coherent, takes into account the actual values of the risk variable on the tail event, but is not elicitable.

An alternative risk measure which addresses these issues is given by the concept of expectiles, introduced by Newey and Powell (1987). Expectiles are found by substituting the absolute deviations in (1) with squared deviations:

$$\xi_\tau = \arg \min_{\theta \in \mathbb{R}} \mathbb{E}(\eta_\tau(X - \theta) - \eta_\tau(X)), \quad (2)$$

where $\eta_\tau(x) = |\tau - \mathbb{1}\{x \leq 0\}|x^2$. The special case $\tau = 0.5$ leads to the expectation of X as its best L^2 -predictor. For each $\tau \in (0, 1)$, the τ th expectile exists and is uniquely defined by its convex problem, and satisfies $\tau = \mathbb{E}[\mathbb{1}\{X \leq \xi_\tau\}]/\mathbb{E}|X - \xi_\tau|$. This interpretation of expectiles is in fact intimately connected to the notion of gain-loss ratio, which is a popular performance measure in portfolio management and is well-known in the literature on no good deal valuation in incomplete markets (see Bellini and Di Bernardino, 2017, and references therein). More generally, it implies that the τ th expectile is determined by tail expectations rather than tail probabilities and as such, unlike the VaR, it depends on tail realisations of the loss variable and their probability. This motivated Kuan et al. (2009) to introduce a notion of expectile-based Value-at-Risk as $-\xi(\tau)$ for real-valued profit-loss distributions. From an axiomatic viewpoint, the advantages of the expectile include that it induces a law-invariant, coherent and elicitable risk measure, see Bellini et al. (2014) and Ziegel (2016). It is actually the only risk measure, apart from the simple expectation, satisfying these three properties. Further results, both theoretical and numerical, obtained by Ehm et al. (2016)

and Bellini and Di Bernardino (2017) among others, indicate that expectiles define sensible alternatives to the standard VaR and ES.

On the statistical side, expectile estimation has recently regained growing interest in the context of nonparametric, semiparametric and more complex models, see for example Sobotka and Kneib (2012) and the references therein, as well as the two recent contributions by Holzmann and Klar (2016) and Krättschmer and Zähle (2017) for advanced theoretical results. However, a theory for extreme expectiles is still in full development. The probabilistic aspects of extreme expectiles, with $\tau \downarrow 0$ or $\tau \uparrow 1$, have been examined by Bellini et al. (2014) and Bellini and Di Bernardino (2017). Inference on extreme expectiles has been considered even more recently in Daouia et al. (2018, 2019, 2020), where the case $\tau = \tau_n \rightarrow 1$, as the available sample size n goes to infinity, is considered. These estimation methods are studied in the context of heavy-tailed distributions, which are ubiquitous in the modelling of extreme actuarial and financial losses, as argued on p.9 of Embrechts et al. (1997), as well as in Chavez-Demoulin et al. (2014) and the references therein.

This literature on extreme expectile estimation has so far mostly been restricted to independent and identically distributed (i.i.d.) data. The exception is Daouia et al. (2019), which considers data coming from a strictly stationary ϕ -mixing time series. This extension is, in fact, only of minor interest in practice because standard processes in financial and econometric modelling, such as ARCH and GARCH processes, are not in general ϕ -mixing. This is a serious gap in the current theory, to be addressed if expectiles are to be used widely in financial risk management. The derivation of the asymptotic distribution of extreme expectile estimators, in a much more general stationary but weakly dependent context, is therefore of great interest, as it would allow the construction of asymptotic confidence intervals that take the weak dependence between observations into account. Note that this problem is different from the prediction of extreme expectiles in dynamic time series models, where the interest is in estimating conditional extreme expectile levels for tomorrow with our knowledge of today, with a particular view on accommodating specific features of time series data such as heteroscedasticity and volatility clustering. From a purely quantile-based perspective, this other problem is tackled in, for instance, McNeil and Frey (2000) using a combination of filtering techniques and standard extreme value theory.

This paper focuses on the rigorous study of estimation techniques for extreme expectiles in the stationary and weakly dependent context, by working on two classes of extreme expectile estimators whose convergence is mathematically shown within a framework of β -mixing and heavy-tailed observations. Given a strictly stationary time series (X_t) having a one-dimensional marginal heavy-tailed distribution, whose serial dependence satisfies suitable general conditions, we begin by estimating an intermediate tail expectile of order $\tau_n \rightarrow 1$ such that $n(1 - \tau_n) \rightarrow \infty$ as $n \rightarrow \infty$. The resulting estimates are then extrapolated to proper extreme levels τ'_n converging to 1 at an arbitrarily fast rate in the sense that $n(1 - \tau'_n) \rightarrow c$ as $n \rightarrow \infty$, for some nonnegative constant c . Two estimation methods are considered: the Least Asymmetrically Weighted Squares (LAWS) estimator, defined as the direct empirical counterpart of the expectile through a minimisation formulation, and the indirect Quantile-Based (QB) estimator obtained using an asymptotic proportionality relationship that links

high expectiles to their quantile counterparts. We successfully develop their asymptotic theory to a β -mixing framework considered in the context of extreme quantile estimation by Drees (2003) (see also Drees, 2000, 2002), which provides probabilistic tools making it possible to examine the asymptotic properties of a wide class of statistical indicators of extremes of, among others, ARMA, ARCH and GARCH processes under reasonably general and mild conditions. We then discuss the construction of asymptotic confidence intervals for extreme expectiles that take into account the dependence between observations, and compare these intervals to the naive intervals obtained via the asymptotic theory of i.i.d. observations, as a way to illustrate the importance of accounting for dependence when extreme expectiles of a financial time series are estimated.

The outline of the paper is the following. Section 2 explains in detail our statistical context. Section 3 contains the main results of the paper, first on intermediate expectile estimation and then on expectile estimation at extreme levels. Section 4 explores the implications of our results on asymptotic confidence interval construction. Section 5 discusses the important question of the selection of expectile level in practice. The finite-sample performance of the methods is examined on simulated data sets in Section 6 and on real financial data in Section 7. Section 8 concludes with a short discussion.

The methods and data considered in this article have been incorporated into the **R** package **ExtremeRisks**, freely available on CRAN. The **R** code for the simulation study and real data analysis is available in the “Software” Section of S. Padoan’s personal webpage at the URL <http://mypage.unibocconi.it/simonepadoan/>. A Supplementary Material document discusses in more depth our technical conditions, gives all necessary mathematical proofs, and contains further finite-sample results.

2 Statistical model and time series framework

Let $(X_t, t \geq 1)$ be a strictly stationary time series having a **continuous** one-dimensional marginal heavy-tailed distribution F : in other words, F is the distribution function of $X := X_1$ as well as of each X_t . Let $\bar{F} := 1 - F$ be the related survival function, and $U : s \mapsto \inf\{x \in \mathbb{R} \mid 1/\bar{F}(x) \geq s\}$ be its tail quantile function. Throughout, X should be seen as $-Y$, where Y denotes a generic financial position, so that large positive values of X represent extreme losses associated to Y .

Our target in the present paper is the estimation of high and extreme expectiles of X having order tending to 1, motivated by applications to financial risk management. To this end, we focus on heavy-tailed distributions with tail index $0 < \gamma < 1$, which are found to model the tail structure of many financial data examples quite well. The survival function of a heavy-tailed distribution can be expressed as $\bar{F}(x) = x^{-1/\gamma}L(x)$, for $x > 0$ large enough, where L is a slowly varying function at infinity. In other words,

$$\forall x > 0, \lim_{z \rightarrow \infty} \frac{\bar{F}(zx)}{\bar{F}(z)} = x^{-1/\gamma} \quad \text{or equivalently} \quad \lim_{z \rightarrow \infty} \frac{U(zx)}{U(z)} = x^\gamma. \quad (3)$$

The tail index γ specifies the tail heaviness of \bar{F} : the tail of \bar{F} gets heavier as γ increases.

Together with condition $\mathbb{E}|X_-| < \infty$, where $X_- := \min(X, 0)$, the assumption $\gamma < 1$ then ensures that the first moment of X exists, which entails that expectiles of X of any order are well-defined. These conditions shall therefore be part of our minimal assumptions.

An extension of the results of Daouia et al. (2018), on extreme expectile estimation for heavy tails in the i.i.d. setup, is given in Daouia et al. (2019) in a ϕ -mixing dependence framework. To the best of our knowledge, the latter is the only work in the literature considering the estimation of extreme expectiles in heavy-tailed models for weakly dependent data. The ϕ -mixing framework is the following. For any $m \geq 1$, let $\mathcal{F}_{1,m} = \sigma(X_1, \dots, X_m)$ and $\mathcal{F}_{m,\infty} = \sigma(X_m, X_{m+1}, \dots)$ denote the past and future σ -fields generated by the sequence (X_t) . The ϕ -mixing coefficients of this sequence are then defined by:

$$\forall l \geq 1, \phi(l) = \sup_{m \geq 1} \sup_{A \in \mathcal{F}_{1,m}} \sup_{B \in \mathcal{F}_{m+l,\infty}} |\mathbb{P}(B|A) - \mathbb{P}(B)|.$$

The time series (X_t) is said to be ϕ -mixing (or uniformly strongly mixing) if $\phi(l) \rightarrow 0$ as $l \rightarrow \infty$. This is in fact a very stringent assumption. For instance, even the simple AR(1) process with heavy-tailed innovations is never ϕ -mixing (see the Introduction of Rio, 2017). We work here in the more general context of β -mixing, defined through the coefficients

$$\forall l \geq 1, \beta(l) = \sup_{m \geq 1} \mathbb{E} \left(\sup_{B \in \mathcal{F}_{m+l,\infty}} |\mathbb{P}(B|\mathcal{F}_{1,m}) - \mathbb{P}(B)| \right).$$

The time series (X_t) is then said to be β -mixing (or absolutely regular) if $\beta(l) \rightarrow 0$ as $l \rightarrow \infty$. Roughly speaking, the β -mixing property brings a form of memorylessness much weaker than its ϕ -mixing version: the β -mixing property is written in an L^1 sense, while the ϕ -mixing property is written in the much stronger L^∞ sense. That β -mixing is weaker than ϕ -mixing can be seen by noting that $\beta(l) \leq \phi(l)$ for any l , see Doukhan (1994, Section 1.1).

Our motivation for making the β -mixing assumption is twofold. On the one hand, β -mixing is satisfied in a much wider class of models than ϕ -mixing: for instance, Doukhan (1994, Section 2.4) shows that a large class of Markov processes, among which ARMA processes, nonlinear autoregressive processes, ARCH and GARCH models are in fact geometrically β -mixing (*i.e.* there is $a < 1$ such that $\beta(l) \leq a^l$ for l large enough) under reasonably general conditions. On the other hand, there is a general theory of extremes for strictly stationary and β -mixing processes, developed in a series of papers by Drees (2000, 2002, 2003). This body of work provides probabilistic tools for the statistical analysis of extremes of strictly stationary and β -mixing observations through a general approximation result for the tail quantile process by a Gaussian process. Mixing conditions have more generally played a substantial role in later research on the extremes of a time series: see among others Robert (2008, 2009) who designed inference procedures for the clustering phenomenon in the extremes of time series, de Haan et al. (2016) for the development of bias-reduced estimators in the dependent setting, and Rootzén (2009) and Drees and Rootzén (2010) for complements to the asymptotic theory of extremes in mixing time series. The β -mixing assumption thus strikes a good balance between theoretical applicability and modelling strength, and as such constitutes a reasonable framework for our objective of estimating extreme expectiles in heavy-tailed time series. This motivates our basic modelling assumption below.

Condition A. The time series (X_t) is strictly stationary, β -mixing and its one-dimensional marginal distribution function F is continuous and heavy-tailed.

Condition A will be key to our development of an asymptotic theory for extreme expectile estimation, which we provide in the next section.

3 Extreme expectile estimation in time series

Suppose that we observe a random sample (X_1, \dots, X_n) extracted from a time series (X_t) satisfying Condition A. Denote by $X_{1,n} \leq \dots \leq X_{n,n}$ the ascending order statistics of (X_1, \dots, X_n) . The objective in this section is to estimate a marginal, unconditional extreme expectile ξ_{τ_n} of the random variable X , where $\tau_n \rightarrow 1$ as $n \rightarrow \infty$. We shall start by the case of an intermediate level τ_n , meaning that $\tau_n \rightarrow 1$ and $n(1 - \tau_n) \rightarrow \infty$ as $n \rightarrow \infty$. Intermediate expectile estimates will then be extrapolated to estimate expectiles at properly extreme levels τ'_n , satisfying $n(1 - \tau'_n) \rightarrow c > 0$ as $n \rightarrow \infty$, using a semiparametric approach warranted by the heavy-tailed assumption.

3.1 At intermediate levels

Direct asymmetric least squares estimator Let τ_n be an intermediate level. We first consider estimating the expectile ξ_{τ_n} of the marginal distribution F by its direct empirical estimator

$$\tilde{\xi}_{\tau_n} = \arg \min_{\theta \in \mathbb{R}} \sum_{t=1}^n \eta_{\tau_n}(X_t - \theta). \quad (4)$$

This LAWS estimator can easily be computed, for example using an iteratively reweighted least squares minimisation procedure. To find the asymptotic distribution of $\tilde{\xi}_{\tau_n}$, we make the following assumption on the dependence within the time series (X_t) .

Condition B. For the time series (X_t) , assume that

- (i) There are sequences of integers (l_n) and (r_n) such that

$$l_n \rightarrow \infty, r_n \rightarrow \infty, \frac{l_n}{r_n} \rightarrow 0, \frac{r_n}{n} \rightarrow 0 \text{ and } \frac{n \beta(l_n)}{r_n} \rightarrow 0 \text{ as } n \rightarrow \infty.$$

- (ii) For any $t \geq 1$, there is a function R_t on $\mathbb{D} := [0, \infty]^2 \setminus \{(\infty, \infty)\}$ such that

$$\forall (x, y) \in \mathbb{D}, \lim_{s \rightarrow \infty} s \mathbb{P} \left(\overline{F}(X_1) \leq \frac{x}{s}, \overline{F}(X_{t+1}) \leq \frac{y}{s} \right) = R_t(x, y).$$

- (iii) There exist $D \geq 0$ and a nonnegative sequence $\rho(t)$ satisfying $\sum_{t \geq 1} \rho(t) < \infty$ and such that, for s large enough, we have

$$s \mathbb{P} \left(\frac{u'}{s} < \overline{F}(X_1) \leq \frac{u}{s}, \frac{v'}{s} < \overline{F}(X_{t+1}) \leq \frac{v}{s} \right) \leq \rho(t) \sqrt{(u - u')(v - v')} + \frac{D}{s} (u - u')(v - v'),$$

for any $t \geq 1$ and all $u, u', v, v' \in [0, 1]$ with $u' < u$ and $v' < v$.

Condition B(i) and B(ii) are standard in the emerging literature on extreme value analysis with mixing conditions, see *e.g.* Drees (2002, 2003), Rootzén (2009) and Drees and Rootzén (2010) (we thank Holger Drees for pointing out in private communication that there is a typo in the first term of condition (C1) in Drees, 2003). In Assumption B(i), the sequences (l_n) and (r_n) are small-block and big-block sequences used to develop the kind of “big blocks separated by small blocks” arguments that are successfully employed in the literature on mixing time series. Condition B(iii) is slightly more precise than condition (C3) in Drees (2003).

Now, we state our first main result, on the asymptotic normality of the estimator $\tilde{\xi}_{\tau_n}$. Recall the notation $X_- := \min(X, 0)$ for the negative part of X .

Theorem 3.1. *Assume that conditions A and B are satisfied. Assume further that there is $\delta > 0$ such that $\mathbb{E}|X_-|^{2+\delta} < \infty$, $0 < \gamma < 1/(2 + \delta)$ and $\sum_{l \geq 1} [\beta(l)]^{\delta/(2+\delta)} < \infty$. Let $\tau_n \uparrow 1$ be such that $n(1 - \tau_n) \rightarrow \infty$, $r_n(1 - \tau_n) \rightarrow 0$ and $r_n(r_n/\sqrt{n(1 - \tau_n)})^\delta \rightarrow 0$ as $n \rightarrow \infty$. Then*

$$\sqrt{n(1 - \tau_n)} \left(\frac{\tilde{\xi}_{\tau_n}}{\xi_{\tau_n}} - 1 \right) \xrightarrow{d} \mathcal{N} \left(0, \frac{2\gamma^3}{1 - 2\gamma} (1 + \sigma^2(\gamma, \mathbf{R})) \right),$$

with $\sigma^2(\gamma, \mathbf{R}) := (1 - \gamma)(1 - 2\gamma) \iint_{(0,1)^2} \sum_{t=1}^{\infty} R_t(u, v) \frac{du}{u^{\gamma+1}} \frac{dv}{v^{\gamma+1}}$.

The family of functions R_t specifies the extremal dependence within the time series between different time points; when $R_t \equiv 0$ for any $t \geq 1$, which is for instance the case when (X_t) is an i.i.d. sequence, the asymptotic variance is $2\gamma^3/(1 - 2\gamma)$. The quantity $\sigma^2(\gamma, \mathbf{R})$ represents the proportion of increase of this asymptotic variance due to the mixing setting. The conditions $\mathbb{E}|X_-|^{2+\delta} < \infty$ and $0 < \gamma < 1/(2 + \delta)$ already appear in Theorem 2 of Daouia et al. (2018) for the i.i.d. case, of which the present result can be considered a generalisation. Our Theorem 3.1 represents a substantial theoretical step compared to Theorem 1 of Daouia et al. (2019) in the ϕ -mixing case. While the latter is essentially shown by simply updating a couple of correlation calculations in the proof of Theorem 2 of Daouia et al. (2018), the proof of Theorem 3.1 uses some rather delicate arguments involving a tailored central limit theory for tail array sums in the time-dependent setting, developed by Rootzén et al. (1998).

Although an exhaustive discussion of our hypotheses is somewhat involved, our assumptions are in fact very mild when $\beta(l)$ converges to 0 geometrically fast as $l \rightarrow \infty$. In that case, one may choose, for instance, $l_n = \lfloor C \log n \rfloor$, $r_n = \lfloor \log^2(n) \rfloor$ and $\tau_n = 1 - n^{-\tau}$, for any $\tau \in (0, 1)$ and sufficiently large C . The case of geometrically strong β -mixing covers many cases widely used in the modelling of financial time series, such as ARMA processes, ARCH/GARCH processes and solutions of stochastic difference equations (we refer to Doukhan, 1994; Drees, 2000, 2003; Francq et al., 2006; Boussama et al., 2011). Besides, in this geometrically mixing case, condition $\sum_{l \geq 1} [\beta(l)]^{\delta/(2+\delta)} < \infty$ is satisfied for any $\delta > 0$. The following corollary of Theorem 3.1 can be viewed as a synthesis of this theoretical discussion of our assumptions in this important geometrically mixing case.

Corollary 3.2. *Assume that conditions A and B(ii)-(iii) are satisfied, and that $\beta(l) = O(a^l)$ for some $a \in (0, 1)$. Assume further that there is $\delta > 0$ such that $\mathbb{E}|X|^{2+\delta} < \infty$. Let $\tau_n = 1 - n^{-\tau}$, for some $\tau \in (0, 1)$. Then*

$$\sqrt{n(1 - \tau_n)} \begin{pmatrix} \tilde{\xi}_{\tau_n} \\ \xi_{\tau_n} \end{pmatrix} - 1 \xrightarrow{d} \mathcal{N} \left(0, \frac{2\gamma^3}{1 - 2\gamma} (1 + \sigma^2(\gamma, \mathbf{R})) \right),$$

with the notation of Theorem 3.1.

We turn to the consideration of a different estimator, built on the asymptotic proportionality between high expectiles and their quantile counterparts.

Indirect quantile-based estimator A competitor is obtained by exploiting an asymptotic proportionality relationship between high expectiles and quantiles: within our heavy-tailed model,

$$\frac{\xi_\tau}{q_\tau} \rightarrow (\gamma^{-1} - 1)^{-\gamma} \text{ as } \tau \uparrow 1. \quad (5)$$

This was first noted by Bellini et al. (2014). An indirect QB estimator of ξ_{τ_n} can then be obtained through the asymptotic proportionality relationship (5):

$$\widehat{\xi}_{\tau_n} = (\widehat{\gamma}_n^{-1} - 1)^{-\widehat{\gamma}_n} \widehat{q}_{\tau_n},$$

where $\widehat{q}_{\tau_n} = X_{n - \lfloor n(1 - \tau_n) \rfloor, n}$ is the empirical counterpart of q_{τ_n} and $\widehat{\gamma}_n$ is a consistent estimator of γ . We provide first, under the appropriate second-order condition below, a high-level result making it possible to quantify the bias incurred by using convergence (5).

Condition C. The function \overline{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 measurable function A having constant sign and converging to 0 at infinity, that is,

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

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

This second-order condition on \overline{F} controls the rate of convergence in the heavy-tailed assumption (3). More precisely, the larger $|\rho|$ is, the faster the function $|A|$ converges to 0 (since $|A|$ is regularly varying with index ρ , see Theorems 2.3.3 and 2.3.9 in de Haan and Ferreira, 2006) and the smaller the error in the approximation of the tail of \overline{F} by a purely Pareto tail will be. Further interpretation of this assumption can be found in Beirlant et al. (2004) and de Haan and Ferreira (2006) along with numerous examples of commonly used continuous distributions satisfying this condition.

Proposition 3.3. *Assume that $\mathbb{E}|X_-| < \infty$ and condition C holds with $0 < \gamma < 1$. Assume further that $\tau_n \uparrow 1$ and $n(1 - \tau_n) \rightarrow \infty$ as $n \rightarrow \infty$. Suppose also that there is a nondegenerate limiting random pair (Γ, Θ) such that*

$$\sqrt{n(1 - \tau_n)} \left(\widehat{\gamma}_n - \gamma, \frac{\widehat{q}_{\tau_n}}{q_{\tau_n}} - 1 \right) \xrightarrow{d} (\Gamma, \Theta). \quad (6)$$

If $\sqrt{n(1 - \tau_n)}A((1 - \tau_n)^{-1}) \rightarrow \lambda_1 \in \mathbb{R}$ and $\sqrt{n(1 - \tau_n)}q_{\tau_n}^{-1} \rightarrow \lambda_2 \in \mathbb{R}$ as $n \rightarrow \infty$, then

$$\sqrt{n(1 - \tau_n)} \begin{pmatrix} \widehat{\xi}_{\tau_n} \\ \xi_{\tau_n} \end{pmatrix} - 1 \xrightarrow{d} m(\gamma)\Gamma + \Theta - \lambda$$

with $m(\gamma) := (1 - \gamma)^{-1} - \log(\gamma^{-1} - 1)$ and

$$\lambda := \left(\frac{(\gamma^{-1} - 1)^{-\rho}}{1 - \gamma - \rho} + \frac{(\gamma^{-1} - 1)^{-\rho} - 1}{\rho} \right) \lambda_1 + \gamma(\gamma^{-1} - 1)^\gamma \mathbb{E}(X) \lambda_2.$$

The above result does not make any assumption about the dependence within the time series (X_t) and, contrary to the related Theorem 1 in Daouia et al. (2018), does not feature the unnecessary assumption of an increasing distribution function F . In practice, conditions on the type of dependence featured in the series (X_t) will of course be required to check this high-level condition for given estimators.

Our next main contribution is to give, in our mixing time series framework, two examples of estimators $\widehat{\gamma}_n$ for which Proposition 3.3 can be applied. We start by the Hill estimator (Hill, 1975), which is also the maximum likelihood estimator in a purely Pareto model and is arguably the most popular semiparametric estimator in the analysis of heavy tails:

$$\widehat{\gamma}_n^H = \frac{1}{[n(1 - \tau_n)]} \sum_{i=1}^{\lfloor n(1 - \tau_n) \rfloor} \log \left(\frac{X_{n-i+1,n}}{X_{n-\lfloor n(1 - \tau_n) \rfloor, n}} \right).$$

The key result in this case, which is of interest in its own right, consists in a joint Gaussian approximation of the processes $s \mapsto \widehat{q}_{1-(1-\tau_n)s} = X_{n-\lfloor n(1-\tau_n)s \rfloor, n}$ and $s \mapsto \log \widehat{q}_{1-(1-\tau_n)s}$ (that is, the tail empirical quantile process and its logarithm) in our mixing framework. See Theorem 2.4.8 in de Haan and Ferreira (2006) for an analogue result in the independent case, as well as Proposition A.1 in de Haan et al. (2016) and Proposition 1 in Chavez-Demoulin and Guillou (2018) for related statements on the tail empirical quantile process with dependent observations.

Theorem 3.4. *Assume that conditions A, B and C are satisfied. Assume that $\tau_n \uparrow 1$, $n(1 - \tau_n) \rightarrow \infty$, $r_n(1 - \tau_n) \rightarrow 0$, $r_n \log^2(n(1 - \tau_n))/\sqrt{n(1 - \tau_n)} \rightarrow 0$ and $\sqrt{n(1 - \tau_n)}A((1 - \tau_n)^{-1}) = O(1)$ as $n \rightarrow \infty$. Suppose finally that $n(1 - \tau_n)$ is a sequence of integers and pick $s_0 > 0$. Then there exist appropriate versions of the process $s \mapsto \widehat{q}_{1-(1-\tau_n)s}$ and a continuous, centred Gaussian process W having covariance function*

$$r(x, y) := \min(x, y) + \sum_{t=1}^{\infty} R_t(x, y) + R_t(y, x)$$

such that, for any $\varepsilon > 0$ sufficiently small, we have, uniformly in $s \in (0, s_0]$,

$$\frac{\widehat{q}_{1-(1-\tau_n)s}}{q_{\tau_n}} = s^{-\gamma} \left(1 + \frac{1}{\sqrt{n(1-\tau_n)}} \gamma s^{-1} W(s) + \frac{s^{-\rho} - 1}{\rho} A((1-\tau_n)^{-1}) + o_{\mathbb{P}} \left(\frac{s^{-1/2-\varepsilon}}{\sqrt{n(1-\tau_n)}} \right) \right)$$

and

$$\log \frac{\widehat{q}_{1-(1-\tau_n)s}}{q_{\tau_n}} = -\gamma \log s + \frac{1}{\sqrt{n(1-\tau_n)}} \gamma s^{-1} W(s) + \frac{s^{-\rho} - 1}{\rho} A((1-\tau_n)^{-1}) + o_{\mathbb{P}} \left(\frac{s^{-1/2-\varepsilon}}{\sqrt{n(1-\tau_n)}} \right).$$

As the statement of Theorem 3.4 makes clear, this approximation result only applies to appropriate versions of the tail empirical process, equal in distribution to the original process, on a rich enough probability space. Our weak convergence results following from Theorem 3.4 are of course unaffected by this choice. Our next such result gives the asymptotic behaviour of the pair $(\widehat{\gamma}_n^H, \widehat{q}_{\tau_n})$, which we then use to find the limiting distribution of the estimator $\widehat{\xi}_{\tau_n}$ constructed using $\widehat{\gamma}_n^H$ as the tail index estimator.

Corollary 3.5. *Assume that conditions A, B and C are satisfied. Let $\tau_n \uparrow 1$ be such that $n(1-\tau_n) \rightarrow \infty$, $r_n(1-\tau_n) \rightarrow 0$, $r_n \log^2(n(1-\tau_n))/\sqrt{n(1-\tau_n)} \rightarrow 0$ and $\sqrt{n(1-\tau_n)}A((1-\tau_n)^{-1}) \rightarrow \lambda_1 \in \mathbb{R}$ as $n \rightarrow \infty$. Then*

$$\sqrt{n(1-\tau_n)} \left(\widehat{\gamma}_n^H - \gamma, \frac{\widehat{q}_{\tau_n}}{q_{\tau_n}} - 1 \right) \xrightarrow{d} (\Gamma, \Theta)$$

where (Γ, Θ) is a Gaussian random pair having expectation $(\lambda_1/(1-\rho), 0)^\top$ and covariance matrix defined by

$$\begin{aligned} \text{Var}(\Gamma) = \text{Var}(\Theta) &= \gamma^2 \left(1 + 2 \sum_{t=1}^{\infty} R_t(1, 1) \right) \\ \text{and } \text{Cov}(\Gamma, \Theta) &= \gamma^2 \int_0^1 \sum_{t=1}^{\infty} \left[\frac{R_t(s, 1) + R_t(1, s)}{s} - 2R_t(1, 1) \right] ds. \end{aligned}$$

If moreover $\mathbb{E}|X_-| < \infty$ and $0 < \gamma < 1$ as well as $\sqrt{n(1-\tau_n)}q_{\tau_n}^{-1} \rightarrow \lambda_2 \in \mathbb{R}$ as $n \rightarrow \infty$, then for $\widehat{\gamma}_n = \widehat{\gamma}_n^H$ in the estimator $\widehat{\xi}_{\tau_n}$,

$$\sqrt{n(1-\tau_n)} \left(\frac{\widehat{\xi}_{\tau_n}}{\xi_{\tau_n}} - 1 \right) \xrightarrow{d} \mathcal{N} \left(\frac{m(\gamma)}{1-\rho} \lambda_1 - \lambda, \gamma^2 \times v^H(\gamma, \mathbf{R}) \right)$$

with λ as in Proposition 3.3 and

$$\begin{aligned} v^H(\gamma, \mathbf{R}) &:= (1 + [m(\gamma)]^2) \left(1 + 2 \sum_{t=1}^{\infty} R_t(1, 1) \right) \\ &+ 2m(\gamma) \int_0^1 \sum_{t=1}^{\infty} \left[\frac{R_t(s, 1) + R_t(1, s)}{s} - 2R_t(1, 1) \right] ds. \end{aligned}$$

Corollary 3.5 is, to the best of our knowledge, the first result on the QB estimator at intermediate levels under weak dependence assumptions. This result contains Corollary 2 in Daouia et al. (2018), restricted to the i.i.d. setup, in which case the asymptotic variance of $\widehat{\xi}_{\tau_n}$ is $\gamma^2(1 + [m(\gamma)]^2)$, as our result indeed shows by taking $R_t \equiv 0$ for any $t \geq 1$.

There are of course many other ways to estimate the tail index γ . We briefly present here an alternative, novel estimator based on the use of intermediate expectiles. The asymptotic proportionality relationship (5) can be equivalently rephrased as $\overline{F}(\xi_\tau)/(1 - \tau) \rightarrow \gamma^{-1} - 1$ as $\tau \uparrow 1$, which implies

$$\gamma = \lim_{\tau \uparrow 1} \left(1 + \frac{\overline{F}(\xi_\tau)}{1 - \tau} \right)^{-1}.$$

Taking $\tau = \tau_n \rightarrow 1$, and estimating $\overline{F}(\xi_\tau)$ by $\widehat{F}_n(\tilde{\xi}_{\tau_n})$, where $\widehat{F}_n(u) = n^{-1} \sum_{t=1}^n \mathbb{1}\{X_t > u\}$ is the empirical survival function, suggests the Expectile-Based (EB) estimator

$$\widehat{\gamma}_n^E = \left(1 + \frac{\widehat{F}_n(\tilde{\xi}_{\tau_n})}{1 - \tau_n} \right)^{-1}.$$

It can be seen there that, due to asymptotic variance considerations, this estimator will tend to be less variable than the Hill estimator (in the i.i.d. case, when $\gamma < 0.38$). This may make it a valuable device in the construction of confidence intervals requiring an estimate of γ .

3.2 At extreme levels

We now consider the important problem of the estimation of extreme expectiles $\xi_{\tau'_n}$, whose level $\tau'_n \rightarrow 1$ satisfies $n(1 - \tau'_n) \rightarrow c \in [0, \infty)$ as $n \rightarrow \infty$. A typical choice in applications is $\tau'_n = 1 - p_n$ for an exceedance probability p_n not greater than $1/n$, see *e.g.* Cai et al. (2015). The idea of the semiparametric approach we present here is to define an estimator of an extreme expectile through a Weissman-type construction (Weissman, 1978). This is motivated by a combination of the heavy-tailed assumption with Equation (5), resulting in

$$\frac{\xi_{\tau'_n}}{\xi_{\tau_n}} \approx \frac{q_{\tau'_n}}{q_{\tau_n}} = \frac{U((1 - \tau'_n)^{-1})}{U((1 - \tau_n)^{-1})} \approx \left(\frac{1 - \tau'_n}{1 - \tau_n} \right)^{-\gamma} \quad \text{as } n \rightarrow \infty.$$

This suggests to consider the following class of plug-in estimators of $\xi_{\tau'_n}$:

$$\overline{\xi}_{\tau'_n}^* \equiv \overline{\xi}_{\tau'_n}^W(\tau_n) := \left(\frac{1 - \tau'_n}{1 - \tau_n} \right)^{-\widehat{\gamma}_n} \overline{\xi}_{\tau_n}$$

where $\widehat{\gamma}_n$ and $\overline{\xi}_{\tau_n}$ are consistent estimators of γ and of the intermediate expectile ξ_{τ_n} , respectively. We say that $\overline{\xi}_{\tau'_n}^*$ is the extrapolating LAWS estimator when $\overline{\xi}_{\tau_n} = \tilde{\xi}_{\tau_n}$, and we denote it by $\widetilde{\xi}_{\tau'_n}^*$. We call it the extrapolating QB estimator when $\overline{\xi}_{\tau_n} = \widehat{\xi}_{\tau_n}$, and we denote it by $\widehat{\xi}_{\tau'_n}^*$. Our next main result gives high-level conditions for the convergence of $\overline{\xi}_{\tau'_n}^*$.

Proposition 3.6. *Assume that $\mathbb{E}|X_-| < \infty$ and condition C holds with $0 < \gamma < 1$ and $\rho < 0$. Assume further that $\tau_n, \tau'_n \uparrow 1$ with $n(1 - \tau_n) \rightarrow \infty$, $n(1 - \tau'_n) \rightarrow c \in [0, \infty)$ and $\sqrt{n(1 - \tau_n)}/\log[(1 - \tau_n)/(1 - \tau'_n)] \rightarrow \infty$ as $n \rightarrow \infty$. Suppose also that there are nondegenerate limiting random variables Γ, Δ such that*

$$\sqrt{n(1 - \tau_n)}(\hat{\gamma}_n - \gamma) \xrightarrow{d} \Gamma \quad \text{and} \quad \sqrt{n(1 - \tau_n)} \begin{pmatrix} \bar{\xi}_{\tau_n} \\ \xi_{\tau_n} \end{pmatrix} - 1 \xrightarrow{d} \Delta.$$

If moreover $\sqrt{n(1 - \tau_n)}A((1 - \tau_n)^{-1}) \rightarrow \lambda_1 \in \mathbb{R}$ and $\sqrt{n(1 - \tau_n)}q_{\tau_n}^{-1} \rightarrow \lambda_2 \in \mathbb{R}$ as $n \rightarrow \infty$, then

$$\frac{\sqrt{n(1 - \tau_n)}}{\log[(1 - \tau_n)/(1 - \tau'_n)]} \begin{pmatrix} \bar{\xi}_{\tau'_n}^* \\ \xi_{\tau'_n} \end{pmatrix} - 1 \xrightarrow{d} \Gamma.$$

Proposition 3.6 extends Theorem 3 in Daouia et al. (2018) by dropping the unnecessary assumption of an increasing function F . When $\hat{\gamma}_n$ is chosen to be the Hill estimator $\hat{\gamma}_n^H$, we have the following corollary of Proposition 3.6.

Corollary 3.7. *Assume that $\mathbb{E}|X_-| < \infty$, and that conditions A, B and C are satisfied with $0 < \gamma < 1$ and $\rho < 0$. Let $\tau_n, \tau'_n \uparrow 1$ with $n(1 - \tau_n) \rightarrow \infty$, $n(1 - \tau'_n) \rightarrow c \in [0, \infty)$ and $\sqrt{n(1 - \tau_n)}/\log[(1 - \tau_n)/(1 - \tau'_n)] \rightarrow \infty$ as $n \rightarrow \infty$. Assume also that $r_n(1 - \tau_n) \rightarrow 0$, $r_n \log^2(n(1 - \tau_n))/\sqrt{n(1 - \tau_n)} \rightarrow 0$, $\sqrt{n(1 - \tau_n)}A((1 - \tau_n)^{-1}) \rightarrow \lambda_1 \in \mathbb{R}$ and $\sqrt{n(1 - \tau_n)}q_{\tau_n}^{-1} \rightarrow \lambda_2 \in \mathbb{R}$ as $n \rightarrow \infty$. Suppose finally that*

$$\sqrt{n(1 - \tau_n)} \begin{pmatrix} \bar{\xi}_{\tau_n} \\ \xi_{\tau_n} \end{pmatrix} - 1 \xrightarrow{d} \Delta.$$

Then, if $\hat{\gamma}_n = \hat{\gamma}_n^H$ in $\bar{\xi}_{\tau'_n}^$, one has*

$$\frac{\sqrt{n(1 - \tau_n)}}{\log[(1 - \tau_n)/(1 - \tau'_n)]} \begin{pmatrix} \bar{\xi}_{\tau'_n}^* \\ \xi_{\tau'_n} \end{pmatrix} - 1 \xrightarrow{d} \mathcal{N} \left(\frac{\lambda_1}{1 - \rho}, \gamma^2 \left[1 + 2 \sum_{t=1}^{\infty} R_t(1, 1) \right] \right).$$

Corollary 3.7 makes it possible to construct confidence intervals for our extreme expectile estimators, and applies indifferently to the estimators extrapolated from the direct and indirect intermediate expectile estimators. When $R_t \equiv 0$ for any $t \geq 1$, the asymptotic variance is γ^2 . The quantity $2 \sum_{t=1}^{\infty} R_t(1, 1)$ represents the proportion of increase of this asymptotic variance, compared to the i.i.d. case, due to the temporal dependence. An analogue result is of course possible for the Expectile-Based (EB) estimator $\hat{\gamma}_n^E$, although we shall not pursue this for the sake of brevity.

4 Asymptotic confidence interval construction

On the basis of the theory developed in Section 3, we propose here asymptotic confidence interval estimators for inferring extreme expectiles. They can be derived from the asymptotic

behaviour of either the extrapolating LAWS or QB estimators (see Corollary 3.7), and we examine both options.

First, we recall that an estimator of the expectile at the extreme level is given by

$$\bar{\xi}_{\tau'_n}^* = \left(\frac{1 - \tau'_n}{1 - \tau_n} \right)^{-\hat{\gamma}_n^H} \bar{\xi}_{\tau_n}, \quad (7)$$

where $\hat{\gamma}_n^H$ is the Hill estimator and $\bar{\xi}_{\tau_n}$ is an estimator of the expectile at the intermediate level. Second, for the asymptotic variance of the estimator in (7), deduced in Corollary 3.7 and that we denote hereafter by $w(\gamma, \mathbf{R})$, we propose the following estimator. By Proposition 2.1 in Drees (2003) we have that, when $n(1 - \tau_n) \rightarrow \infty$, $r_n \rightarrow \infty$ and $r_n(1 - \tau_n) \rightarrow 0$,

$$\frac{1}{r_n(1 - \tau_n)} \text{Var} \left(\sum_{i=1}^{r_n} \mathbb{1}\{F(X_i) > \tau_n\} \right) \rightarrow 1 + 2 \sum_{t=1}^{\infty} R_t(1, 1) \text{ as } n \rightarrow \infty.$$

We then adopt a ‘‘big-block/small-block’’ technique where we split the data into big blocks of size r_n separated by small blocks of size l_n , respectively, and we define the sequence of random variables

$$Z_j = \sum_{t=1+j\ell_n}^{r_n+j\ell_n} \mathbb{1}\{\hat{F}_n(X_t) > \tau_n\}$$

for $j = 0, 1, \dots, m_n - 1$, where $m_n = \lfloor n/\ell_n \rfloor$ and $\ell_n = r_n + l_n$, and \hat{F}_n is the empirical distribution function of all the observations. Therefore, we compute the sample variance of the sequence (Z_0, \dots, Z_{m_n-1}) , say Σ_n . As a result, an empirical estimator of the asymptotic variance $w(\gamma, \mathbf{R}) := \gamma^2 [1 + 2 \sum_{t=1}^{\infty} R_t(1, 1)]$ is

$$\hat{w}_n(\gamma, \mathbf{R}) = \frac{(\hat{\gamma}_n^H)^2}{r_n(1 - \tau_n)} \Sigma_n. \quad (8)$$

Next, note that Corollary 3.7 is equivalent to its (in practice more accurate) log-scale version

$$\frac{\sqrt{n(1 - \tau_n)}}{\log[(1 - \tau_n)/(1 - \tau'_n)]} \log \frac{\bar{\xi}_{\tau'_n}^*}{\bar{\xi}_{\tau_n}} \xrightarrow{d} \mathcal{N} \left(\frac{\lambda_1}{1 - \rho}, w(\lambda, \mathbf{R}) \right) \text{ as } n \rightarrow \infty.$$

Hence, we propose the following interval estimator

$$\left[\bar{\xi}_{\tau'_n}^* \left(\frac{1 - \tau_n}{1 - \tau'_n} \right)^{z_{\alpha/2} \sqrt{\hat{w}_n(\gamma, \mathbf{R})/[n(1 - \tau_n)]}}, \bar{\xi}_{\tau'_n}^* \left(\frac{1 - \tau_n}{1 - \tau'_n} \right)^{z_{1-\alpha/2} \sqrt{\hat{w}_n(\gamma, \mathbf{R})/[n(1 - \tau_n)]}} \right], \quad (9)$$

where $\bar{\xi}_{\tau'_n}^*$ is the estimator in (7), $\hat{w}_n(\gamma, \mathbf{R})$ is the estimator in (8), $z_{\alpha/2}$ and $z_{1-\alpha/2}$ are the $(\alpha/2)$ th and $(1 - \alpha/2)$ th quantiles of the standard normal distribution, with $\alpha \in (0, 1)$. For simplicity we have ignored the bias term $\lambda_1/(1 - \rho)$. We call the estimator in (9) the LAWS-D-based estimator when $\bar{\xi}_{\tau_n} = \tilde{\xi}_{\tau_n}$ in (7), and QB-D-based estimator when $\bar{\xi}_{\tau_n} = \hat{\xi}_{\tau_n}$. Finally,

and for comparison purposes, the asymptotic variance in the i.i.d. case is $w(\gamma, \mathbf{R}) = \gamma^2$, estimated by $(\widehat{\gamma}_n^H)^2$. In this case, an interval estimator is simply

$$\left[\bar{\xi}_{\tau'_n}^* \left(\frac{1 - \tau_n}{1 - \tau'_n} \right)^{z_{\alpha/2} \widehat{\gamma}_n^H / \sqrt{n(1-\tau_n)}}, \bar{\xi}_{\tau'_n}^* \left(\frac{1 - \tau_n}{1 - \tau'_n} \right)^{z_{1-\alpha/2} \widehat{\gamma}_n^H / \sqrt{n(1-\tau_n)}} \right] \quad (10)$$

and we call it the LAWS-IID-based estimator or QB-IID-based estimator depending on whether $\bar{\xi}_{\tau_n} = \widetilde{\xi}_{\tau_n}$ or $\bar{\xi}_{\tau_n} = \widehat{\xi}_{\tau_n}$.

5 Extreme expectile level selection

A crucial practical question in actuarial and financial risk management is the choice of the level of prudence of the risk measure under consideration. When working with quantile-based risk measures, as we mentioned in Section 3.2, one usually chooses extreme tail probabilities $\alpha_n \uparrow 1$ with $n(1 - \alpha_n) \rightarrow c$, a finite constant, as $n \rightarrow \infty$, to allow for more prudent risk management. In the case when expectiles are of interest, a reasonable idea is to select τ'_n so that $\xi_{\tau'_n} \equiv q_{\alpha_n}$ for a given relative frequency α_n . This had already been suggested by Bellini and Di Bernardino (2017), albeit for a normally distributed Y .

For heavy-tailed and i.i.d. data, Daouia et al. (2018) instead suggest a nonparametric estimator of the level τ'_n that satisfies $\xi_{\tau'_n} \equiv q_{\alpha_n}$, without recourse to any parametric distributional specification. Here we briefly recall, and then extend, this theory to serially dependent data. Recall that the interpretation of the expectile $\xi_{\tau'_n}$ in terms of average distances is

$$\tau'_n = \frac{\mathbb{E}[|X - \xi_{\tau'_n}| \mathbb{1}\{X \leq \xi_{\tau'_n}\}]}{\mathbb{E}|X - \xi_{\tau'_n}|}.$$

Using standard theory of conditional tail expectations, Daouia et al. (2018) find (in their Proposition 3) that the extreme expectile level $\tau'_n(\alpha_n) = \tau'_n$ such that $\xi_{\tau'_n} \equiv q_{\alpha_n}$ then satisfies

$$\frac{1 - \tau'_n(\alpha_n)}{1 - \alpha_n} \rightarrow \frac{\gamma}{1 - \gamma} \quad \text{as } n \rightarrow \infty.$$

If $\widehat{\gamma}_n$ is a consistent estimator of γ , one can then define a natural estimator of $\tau'_n(\alpha_n)$ as

$$\widehat{\tau}'_n(\alpha_n) = 1 - (1 - \alpha_n) \frac{\widehat{\gamma}_n}{1 - \widehat{\gamma}_n}.$$

By substituting this estimated value in place of $\tau'_n \equiv \tau'_n(\alpha_n)$ in the extrapolating LAWS estimator $\widetilde{\xi}_{\tau'_n}^*$ and in the extrapolating QB estimator $\widehat{\xi}_{\tau'_n}^*$, we obtain composite estimators of q_{α_n} . It is interesting to note that if one uses the exact same estimator $\widehat{\gamma}_n$ in the extrapolation step and the calculation of $\widehat{\tau}'_n(\alpha_n)$, the composite extrapolating LAWS estimator has a natural interpretation. Indeed, in this case,

$$\widetilde{\xi}_{\widehat{\tau}'_n(\alpha_n)}^* = \left(\frac{1 - \widehat{\tau}'_n(\alpha_n)}{1 - \tau_n} \right)^{-\widehat{\gamma}_n} \widetilde{\xi}_{\tau_n} = (\widehat{\gamma}_n^{-1} - 1)^{\widehat{\gamma}_n} \widetilde{\xi}_{\alpha_n}^*.$$

In other words, rewriting Equation (5) as the approximation $q_{\alpha_n} \approx (\gamma^{-1} - 1)^\gamma \xi_{\alpha_n}$, the composite extrapolating LAWS estimator can be constructed by plugging in an estimator $\hat{\gamma}_n$ and the extrapolating LAWS estimator at level α_n in the right-hand side of this approximation. This way of constructing the composite extrapolating LAWS estimator had not, to the best of our knowledge, been appreciated in the literature. By contrast, Daouia et al. (2018) had already mentioned the interpretation of the extrapolating QB estimator as a Weissman estimator of the quantile.

The available theory of the composite LAWS and QB estimators in Daouia et al. (2018) is limited to i.i.d. data. We give below a result showing their asymptotic normality in our dependent setting, when $\hat{\gamma}_n$ is the Hill estimator.

Theorem 5.1. *Suppose the conditions of Corollary 3.7 hold with α_n in place of τ'_n . Then, if $\hat{\gamma}_n = \hat{\gamma}_n^H$ and $\bar{\xi}^*$ is either $\hat{\xi}^*$ or $\tilde{\xi}^*$, we have*

$$\frac{\sqrt{n(1-\tau_n)}}{\log[(1-\tau_n)/(1-\alpha_n)]} \left(\frac{\bar{\xi}_{\hat{\gamma}_n^H(\alpha_n)}^*}{q_{\alpha_n}} - 1 \right) \xrightarrow{d} \mathcal{N} \left(\frac{\lambda_1}{1-\rho}, \gamma^2 \left[1 + 2 \sum_{t=1}^{\infty} R_t(1,1) \right] \right).$$

6 Simulation experiments

Here we investigate the finite sample performance of the point and interval expectile estimators, at the extreme level, through a simulation study. We consider AR, ARMA, ARCH and GARCH models. Specifically, we first consider the AR(1) family $X_{t+1} = \phi X_t + \varepsilon_t$, where the innovations ε_t are i.i.d. and have a common distribution F_ε which is a Student- t with $\nu > 0$ degrees of freedom. The tail index of F_ε is $\gamma = 1/\nu$, and we choose

(i) $\phi = 0.8, \nu = 3;$

(ii) $\phi = 0.8, \nu = 4.$

These models exhibit fairly strong linear dependence. We then consider the ARMA(1,1) family $X_{t+1} = \phi X_t + \varepsilon_{t+1} + \theta \varepsilon_t$, where the innovations ε_t are i.i.d. and have a common distribution F_ε which is a symmetric Pareto with shape parameter $\zeta > 0$. The tail index of F_ε is $\gamma = 1/\zeta$, and we consider

(iii) $\phi = 0.95, \theta = 0.9, \zeta = 3;$

(iv) $\phi = 0.95, \theta = -0.6, \zeta = 3;$

(v) $\phi = 0.95, \theta = -0.9, \zeta = 3;$

(vi) $\phi = 0.3, \theta = 0.9, \zeta = 3.$

A very strong linear dependence is present in the first two models and a weak one in the second two models. We recall that for standard linear time series families with a heavy-tailed innovation satisfying the so-called tail balance condition, which is the case for models

(i)-(vi), we have that $\overline{F}_X(x) \approx C \overline{F}_\varepsilon(x)$ as $x \rightarrow \infty$, where C is a positive constant (see *e.g.* Drees, 2003, Section 3.2 and the references therein). In other words, the tail index of X in models (i)-(vi) is always $1/3$ except in model (ii), where it is $1/4$. Finally, we consider the nonlinear GARCH(1,1) family $X_t = \sigma_{t+1}\varepsilon_t$, where $\sigma_{t+1}^2 = \alpha_0 + \alpha_1 X_t^2 + \beta \sigma_t^2$, and (ε_t) is a sequence of i.i.d. Gaussian innovations. We work on two ARCH models and then two proper GARCH models:

$$(vii) \quad \alpha_0 = 0.0001, \alpha_1 = 0.9, \beta = 0;$$

$$(viii) \quad \alpha_0 = 0.4, \alpha_1 = 0.6, \beta = 0;$$

$$(ix) \quad \alpha_0 = 0.0001, \alpha_1 = 0.4, \beta = 0.5;$$

$$(x) \quad \alpha_0 = 0.1, \alpha_1 = 0.4, \beta = 0.4.$$

It is known that the marginal distribution F_X is heavy-tailed in such ARCH/GARCH models, under suitable conditions (see *e.g.* Embrechts et al., 1997, Chapter 8). The actual expression of the tail index, however, is quite involved. Numerical experiments suggest that the tail index in models (vii)-(x) is respectively approximately 0.4, 0.25, 0.3 and 0.25. Thus, the tails of the models (vii) and (ix) are quite heavy, while those of the models (viii) and (x) are slightly lighter. In addition, these four models feature quadratic serial dependence. Finally, let us point out that the time series models (iii)-(vii) and (ix) have also been investigated by Drees (2003).

For each of the models (i)-(x), we simulate 10^4 samples of size $n = 2500$. We consider the extreme levels $\tau'_n = 0.9995$ and $\tau'_n = 0.9999$; the former is slightly smaller than $1 - 1/n$ while the latter is larger. The true value $\xi_{\tau'_n}$ is estimated to a high degree of accuracy via Monte-Carlo simulation. We use the extreme expectile estimator as defined in (7). On each simulated dataset, and for both the extrapolating LAWS and QB estimators, we repeat the estimation exercise for $k \in \{6, 8, \dots, 700\}$ and with the intermediate level $\tau_n = 1 - k/n$. Then, we estimate an asymptotic confidence interval using the LAWS-D-based interval estimator in (9), with 95% nominal coverage probability, and we check whether it contains the true expectile value. The big- and small-block sequences are chosen as $r_n = \lfloor \log^2(n) \rfloor$ and $l_n = \lfloor C \log n \rfloor$, respectively, where C is selected such that l_n is greater than or equal to a lag after which the value of the sample autocorrelation is small, *e.g.* smaller than 0.1. This allows us to compute a Monte Carlo approximation of the coverage probability. We carry out the same exercise for the QB-D-based, QB-IID-based and LAWS-IID-based estimators. Results are reported in Figure 1 at level $\tau'_n = 0.9995$; results at level $\tau'_n = 0.9999$ are similar.

It is readily seen that, for time series data, our proposed confidence intervals, derived from the theory of serial dependent data, behave substantially better than the previously available i.i.d.-based intervals. In particular, the QB-IID-based and LAWS-IID-based intervals are overall far too permissive, and are acceptable only within models where serial dependence is weak, such as the ARMA(1,1) model (v). By contrast, in terms of coverage, our QB-D-based and LAWS-D-based intervals provide far better intervals, which seem to perform overall well in linear time series models (although they tend to be conservative in certain cases, see for

instance model (*iv*)). The ARCH and GARCH cases are more difficult, but once again our intervals represent significant improvement compared to the i.i.d. case, by bringing down the non-coverage probability by about a half to around 15% in the ARCH cases, and by getting very close to the nominal level when k is in a neighbourhood of approximately 60.

7 Real data analysis

Stock market index data. We consider here daily negative log-returns of the S&P 500 (GSPC) and Dow Jones Industrial Average (DJIA) index from January 29, 1985 to December 12, 2019 (freely available from Yahoo! Finance). These samples of size $n = 8785$ are plotted on the rightmost panels of Figure 2. They show evidence of the stylised facts such as heteroscedasticity, fat-tailedness, and asymmetric behaviour of positive and negative returns, which are typically found in financial time series (Embrechts et al., 1997). Throughout this section, we use the Hill estimator in our estimates and confidence intervals.

In the literature, the analysis of tail risk of loss returns is typically based on VaR at the 99.9% level (see *e.g.* Drees, 2003; de Haan et al., 2016) or on using a quantile at level $\alpha_n = 1 - p_n$ where p_n is not larger than $1/n$. Bellini and Di Bernardino (2017) showed that forecasts based on the expectile provide capital requirements that are similar to those obtained with VaR for some econometric models. This is as long as the extreme level τ'_n of the expectile is appropriately selected so that it is close to a pre-specified extreme VaR. For this selection, we use the method of Section 5 which, contrary to the method of Bellini and Di Bernardino (2017), does not depend on a specific parametric model. Precisely, here we fix $p_n = 1/n = 1/8784$ and $\alpha_n = 1 - p_n = 0.9998862$, and according to Section 5 we estimate first $\tau'_n(\alpha_n) = \tau'_n(0.9998862)$. Then we compute an estimate of the expectile at the extreme level $\hat{\tau}'_n(\alpha_n)$ using the composite extrapolating LAWS estimator $\tilde{\xi}_{\tau'_n}^*$ and the composite extrapolating QB estimator $\hat{\xi}_{\tau'_n}^*$ for $\tau'_n = \hat{\tau}'_n(\alpha_n)$. This accordingly produces estimators of $\xi_{\tau'_n(\alpha_n)}$, which is also the quantile-VaR $q_{\alpha_n} = q_{1-1/n}$. Let us reiterate here that our focus in this real data application is the estimation of an extreme marginal expectile; we are therefore not tackling the completely different problem of dynamically predicting extreme expectiles, where the interest is in estimating conditional extreme expectile levels for tomorrow with our knowledge of today.

The first panels of Figure 2 display $\hat{\tau}'_n(\alpha_n)$ against k for $k \leq 700$, where as before $\tau_n = 1 - k/n$. After large fluctuations, the estimates stabilise around a (close for both series) common value, and then drift away due to bias from the centre of the distribution. The choice $k = 200$ seems to be a reasonable compromise. We further confirm this choice by calculating the composite extrapolating LAWS and composite extrapolating QB estimators, and the corresponding LAWS-D and QB-D confidence intervals of Section 4 at level $\tau'_n = \hat{\tau}'_n(\alpha_n)$. Again, after the high-variance part of the sample path (see the green lines on the second and third panels of Figure 2), the estimates stabilise and then drift away due to bias. The choice $k = 200$ once again seems sensible and is therefore adopted. It can be seen that our confidence intervals taking the dependence into account are indeed wider than the i.i.d.-based estimators, as a way to reflect better the uncertainty about the estimation. With $k = 200$,

we find $\hat{\tau}'_n(\alpha_n) \approx 0.9999423$ for the S&P 500 data, and 0.9999402 for the Dow Jones data. These levels are indeed larger than the original $\alpha_n = 1 - p_n = 0.9998862$.

We compare our extrapolating LAWS and QB methods with the traditional Weissman extreme quantile estimator at level α_n . This is

$$\hat{q}_{\alpha_n}^* = \left(\frac{1 - \alpha_n}{1 - \tau_n} \right)^{-\hat{\gamma}_n^H} \hat{q}_{\tau_n} = \left(\frac{1 - \alpha_n}{1 - \tau_n} \right)^{-\hat{\gamma}_n^H} X_{n - \lfloor n(1 - \tau_n) \rfloor, n}.$$

Confidence intervals can also be constructed for the extreme quantile q_{α_n} on the basis of this estimator: here we use a method developed by Drees (2003) which, contrary to ours, does not rely on a big-block/small-block argument, see Formula (33) therein. The estimates are reported in Table 1 and can be visualised in the rightmost panel of Figure 2. Quite reassuringly the methods give point estimates that are similar: note that the composite extrapolating QB point estimator is indeed nothing but $\hat{q}_{\alpha_n}^*$. However, on the third and fourth panels of Figure 2, it can be seen that the confidence intervals constructed on the basis of $\hat{q}_{\alpha_n}^*$ and the method of Drees (2003) are in general much more volatile than the LAWS-D-based interval; moreover, in a neighbourhood of our selected value of k , they are very close to the intervals based on i.i.d. theory. For our selected $k = 200$, according to the fourth panels of Figure 2, they do not contain the maximum observation in the sample, even though one is estimating $q_{\alpha_n} = q_{1-1/n}$, whereas the LAWS-D-based interval does contain this maximum value.

Financial returns of individual banks. We carried out an analogue analysis of the financial returns of Goldman Sachs and Morgan Stanley. We consider the negative log-returns on their equity prices at a daily frequency from July 3, 2000, to June 30, 2010. These samples of data were already considered in Cai et al. (2015) and Daouia et al. (2018), albeit under the assumption of i.i.d. data and in the context of measuring systemic risk. Choosing $\alpha_n = 1 - 1/n = 0.9996021$, we display $\hat{\tau}'_n(\alpha_n)$ against k , as well as the composite extrapolating LAWS and composite extrapolating QB estimators, and the corresponding LAWS-D and QB-D confidence intervals at level $\tau'_n = \hat{\tau}'_n(\alpha_n)$, on the first three panels of Figure 3. The choice $k = 150$ seems reasonable here and we then find $\hat{\tau}'_n(\alpha_n) \approx 0.9997239$ for the Goldman Sachs data, and 0.9996626 for the Morgan Stanley data. Again, we compare our estimates to the Weissman extreme quantile estimator at level α_n . All estimates are reported in Table 2 and can be visualised in the rightmost panel of Figure 3. Like on our stock market index data, it can be seen that the confidence intervals constructed with the method of Drees (2003) are very volatile, and in this case, they appear to be far more conservative than the LAWS-D-based interval.

Table 1: Estimates for the negative daily log-returns of the S&P 500 and Dow Jones indices, obtained with $k = 200$. Here $\alpha_n = 1 - 1/n = 0.9998862$.

Estimator	S&P 500	Dow Jones
$\widehat{\gamma}_n = \widehat{\gamma}_n^H$	0.3364 [0.2198; 0.4530]	0.3442 [0.2219; 0.4665]
$\widetilde{\xi}_{\widehat{\gamma}_n(\alpha_n)}^*$	0.1358 [0.0676; 0.2727]	0.1360 [0.0657; 0.2813]
$\widehat{\xi}_{\widehat{\gamma}_n(\alpha_n)}^*$	0.1398 [0.0696; 0.2807]	0.1394 [0.0674; 0.2884]
$\widehat{q}_{\alpha_n}^*$	0.1398 [0.1124; 0.1739]	0.1394 [0.1025; 0.1896]

Table 2: Estimates for the loss returns of Goldman Sachs and Morgan Stanley, obtained with $k = 150$. Here $\alpha_n = 1 - 1/n = 0.9996021$.

Estimator	Goldman Sachs	Morgan Stanley
$\widehat{\gamma}_n = \widehat{\gamma}_n^H$	0.4096 [0.2815; 0.5377]	0.4589 [0.2966; 0.6212]
$\widetilde{\xi}_{\widehat{\gamma}_n(\alpha_n)}^*$	0.2523 [0.1267; 0.5024]	0.4036 [0.1742; 0.9350]
$\widehat{\xi}_{\widehat{\gamma}_n(\alpha_n)}^*$	0.2746 [0.1379; 0.5469]	0.4422 [0.1909; 1.0243]
$\widehat{q}_{\alpha_n}^*$	0.2746 [0.0890; 0.8474]	0.4422 [0.1565; 1.2498]

8 Discussion

We provide asymptotic theory and practical guidance for the estimation of extreme expectiles in a stationary but weakly dependent framework. Our model allows us to handle the ARMA and GARCH models that are of interest in financial practice. This represents a substantial advance compared to earlier procedures in which only i.i.d. or uniformly strongly mixing data could be considered. In particular, the performance of our proposed procedure for confidence interval construction is very encouraging, both on simulated and real data; this procedure has been implemented in an accessible way in our R package `ExtremeRisks`.

With this in mind, there remains a lot to be done. For the sake of simplicity, in the estimation of extreme expectiles and construction of our confidence intervals, we did not consider the question of estimating and correcting for the bias of the procedures. Bias can here be incurred at the extrapolation step either in the estimation of the tail index, or by the use of the extrapolation relationship itself. It is reasonable to think that one can use tools recently developed in the weakly dependent context, such as those of de Haan et al. (2016), to handle the bias of the Hill estimator of the tail index. In addition to obvious theoretical complications, this would also require modifying our method for the construction of confidence intervals: note indeed that according to Theorem 4.1 in de Haan et al. (2016), the asymptotic variance of their bias-reduced version is not the same as that of the Hill estimator. A different idea is to adapt the efficient bias-reduced estimator of Caeiro et al. (2005), whose asymptotic variance coincides with that of the Hill estimator, to our time-

dependent setting. This is, however, a highly non-trivial theoretical step.

Besides, in the kind of stock market data application we consider, a well-known problem is the estimation and use of the dependence between indices or prices of individual stocks. As an illustration, we have represented a scatterplot of the negative daily log-returns of the Dow Jones versus those of the S&P 500 in Figure 4, and similarly we have plotted a scatterplot of the loss returns of Morgan Stanley versus those of Goldman Sachs. It is clear that there is strong positive pairwise dependence in both cases. Our statistical theory, being developed for univariate time series, cannot yet harness that dependence structure to make more precise inferences (*i.e.* with shorter confidence intervals). This will require adapting our procedures to multivariate time series, and as such is of course far beyond the scope of this paper.

Acknowledgements

This research was in large part carried out when the authors were visiting each other at Bocconi University and the University of Nottingham where G. Stupfler was previously based. Support from the Nottingham PEF Fund is gratefully acknowledged. We would also like to thank Holger Drees for a detailed explanation of some of the technical details in Drees (2003) and Drees and Rootzén (2010). S. A. Padoan is supported by the Bocconi Institute for Data Science and Analytics (BIDSA), Italy. G. Stupfler is supported by the French National Research Agency under the grant ANR-19-CE40-0013-01/ExtremReg project.

References

- ARTZNER, P., DELBAEN, F., EBER, J.-M. AND HEATH, D. (1999). Coherent measures of risk. *Mathematical Finance* **9**, 203–228.
- BEIRLANT, J., GOEGEBEUR, Y., SEGERS, J. AND TEUGELS, J. (2004). *Statistics of Extremes: Theory and Applications*. John Wiley & Sons, Chichester.
- BELLINI, F. AND DI BERNARDINO, E. (2017). Risk management with expectiles. *The European Journal of Finance* **23**, 487–506.
- BELLINI, F., KLAR, B., MÜLLER, A. AND GIANIN, E.R. (2014). Generalized quantiles as risk measures. *Insurance: Mathematics and Economics* **54**, 41–48.
- BOUSSAMA, F., FUCHS, F. AND STELZER, R. (2011). Stationarity and geometric ergodicity of BEKK multivariate GARCH models. *Stochastic Processes and their Applications* **121**, 2331–2360.
- CAEIRO, F., GOMES, M.I. AND PESTANA, D.D. (2005). Direct reduction of bias of the classical Hill estimator. *Revstat* **3**, 111–136.

- CAI, J.-J., EINMAHL, J.H.J., DE HAAN, L. AND ZHOU, C. (2015). Estimation of the marginal expected shortfall: the mean when a related variable is extreme. *Journal of the Royal Statistical Society: Series B* **77**, 417–442.
- CHAVEZ-DEMOULIN, V., EMBRECHTS, P. AND SARDY, S. (2014). Extreme-quantile tracking for financial time series. *Journal of Econometrics* **181**, 44–52.
- CHAVEZ-DEMOULIN, V. AND GUILLOU, A. (2018). Extreme quantile estimation for β -mixing time series and applications. *Insurance: Mathematics and Economics* **83**, 59–74.
- DAOUIA, A., GIRARD, S. AND STUPFLER, G. (2018). Estimation of tail risk based on extreme expectiles. *Journal of the Royal Statistical Society: Series B* **80**, 263–292.
- DAOUIA, A., GIRARD, S. AND STUPFLER, G. (2019). Extreme M-quantiles as risk measures: From L^1 to L^p optimization. *Bernoulli* **25**, 264–309.
- DAOUIA, A., GIRARD, S. AND STUPFLER, G. (2020). Tail expectile process and risk assessment. *Bernoulli* **26**, 531–556.
- DOUKHAN, P. (1994). *Mixing: Properties and Examples*. Springer, New York.
- DREES, H. (2000). Weighted approximations of tail processes for β -mixing random variables. *Annals of Applied Probability* **10**, 1274–1301.
- DREES, H. (2002). Tail empirical processes under mixing conditions. In *Dehling, H.G., Mikosch, T. and Sørensen, M. (eds.), Empirical Process Techniques for Dependent Data*. Birkhäuser, Boston.
- DREES, H. (2003). Extreme quantile estimation for dependent data, with applications to finance. *Bernoulli* **9**, 617–657.
- DREES, H. AND ROOTZÉN, H. (2010). Limit theorems for empirical processes of cluster functionals. *Annals of Statistics* **38**, 2145–2186.
- EHM, W., GNEITING, T., JORDAN, A. AND KRÜGER, F. (2016). Of quantiles and expectiles: consistent scoring functions, Choquet representations and forecast rankings (with discussion). *Journal of the Royal Statistical Society: Series B* **78**, 505–562.
- EMBRECHTS, P., KLÜPPELBERG, C. AND MIKOSCH, T. (1997). *Modelling Extremal Events for Insurance and Finance*. Springer-Verlag, Berlin.
- FRANCQ, C. AND ZAKOÏAN, J.-M. (2006). Mixing properties of a general class of GARCH(1,1) models without moment assumptions on the observed process. *Econometric Theory* **22**, 815–834.
- GNEITING, T. (2011). Making and evaluating point forecasts. *Journal of the American Statistical Association* **106**, 746–762.

- DE HAAN, L. AND FERREIRA, A. (2006). *Extreme Value Theory: An Introduction*. Springer-Verlag, New York.
- DE HAAN, L., MERCADIER, C. AND ZHOU, C. (2016). Adapting extreme value statistics to financial time series: dealing with bias and serial dependence. *Finance and Stochastics* **20**, 321–354.
- HILL, B.M. (1975). A simple general approach to inference about the tail of a distribution. *Annals of Statistics* **3**, 1163–1174.
- HOLZMANN, H. AND KLAR, B. (2016). Expectile asymptotics. *Electronic Journal of Statistics* **10**, 2355–2371.
- KOENKER, R. AND BASSETT, G. (1978). Regression quantiles *Econometrica* **46**, 33–50.
- KRÄTSCHEMER, V. AND ZÄHLE, H. (2017). Statistical inference for expectile-based risk measures. *Scandinavian Journal of Statistics* **44**, 425–454.
- KUAN, C.-M., YEH, J.-H. AND HSU, Y.-C. (2009). Assessing value at risk with CARE, the Conditional Autoregressive Expectile models. *Journal of Econometrics* **150**, 261–270.
- MCNEIL, A.J. AND FREY, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance* **7**, 271–300.
- NEWKEY, W.K. AND POWELL, J.L. (1987). Asymmetric least squares estimation and testing. *Econometrica* **55**, 819–847.
- RIO, E. (2017). Asymptotic Theory of Weakly Dependent Random Processes. *In Probability Theory and Stochastic Modelling*, volume 80. Springer-Verlag, Berlin-Heidelberg.
- ROBERT, C.Y. (2008). Estimating the multivariate extremal index function. *Bernoulli* **14**, 1027–1064.
- ROBERT, C.Y. (2009). Inference for the limiting cluster size distribution of extreme values. *Annals of Statistics* **37**, 271–310.
- ROOTZÉN, H. (2009). Weak convergence of the tail empirical process for dependent sequences. *Stochastic Processes and their Applications* **119**, 468–490.
- ROOTZÉN, H., LEADBETTER, M.R. AND DE HAAN, L. (1998). On the distribution of tail array sums for strongly mixing stationary sequences. *The Annals of Applied Probability* **8**, 868–885.
- SOBOTKA, F. AND KNEIB, T. (2012). Geoadditive expectile regression. *Computational Statistics & Data Analysis* **56**, 755–767.

- WEISSMAN, I. (1978). Estimation of parameters and large quantiles based on the k largest observations. *Journal of the American Statistical Association* **73**, 812–815.
- ZIEGEL, J.F. (2016). Coherence and elicibility. *Mathematical Finance* **26**, 901–918.




Figure 1: Actual non-coverage probabilities (in %) for the confidence intervals of the expectile $\xi_{\tau'_n}$, with 95% nominal level and expectile level $\tau'_n = 0.9995$. From left to right, top row: models (i)-(iv), middle row: models (v) and (vi), bottom row: models (vii)-(x). In each panel, the horizontal dotted red line represents the 5% nominal non-coverage probability, and the solid black, blue dashed, dotted grey and dashed-dotted violet lines represent the actual non-coverage probabilities of the QB-D, LAWS-D, QB-IID and LAWS-IID confidence intervals.




Figure 2: GSPC and DJIA negative daily log-returns data. From left to right, first panel: estimate $\widehat{\tau}'_n(0.9998862)$ (green line). Second panel: composite extrapolated LAWS estimate (green line) at level $\tau'_n = \widehat{\tau}'_n(0.9998862)$, with 95% confidence intervals based on i.i.d. theory (grey dotted line) and our LAWS-D procedure (blue dashed line). Third panel: Weissman quantile estimate (green line) at level $\alpha_n = 0.9998862$ (equivalently, composite extrapolated QB estimate at level $\tau'_n = \widehat{\tau}'_n(0.9998862)$), with 95% confidence intervals based on i.i.d. theory (grey dotted line), our LAWS-D procedure (blue dashed line) and the method of Drees (2003) (purple dashed-dotted line). Fourth panel: Negative daily log-returns data, with the composite extrapolated LAWS estimate (green dashed line) at level $\tau'_n = \widehat{\tau}'_n(0.9998862)$, compared to the standard Weissman quantile estimate $\widehat{q}_{0.9998862}$ (purple dashed line), both with 95% confidence intervals (blue dotted line and purple dotted line respectively). Top row: GSPC data, bottom row: DJIA data.

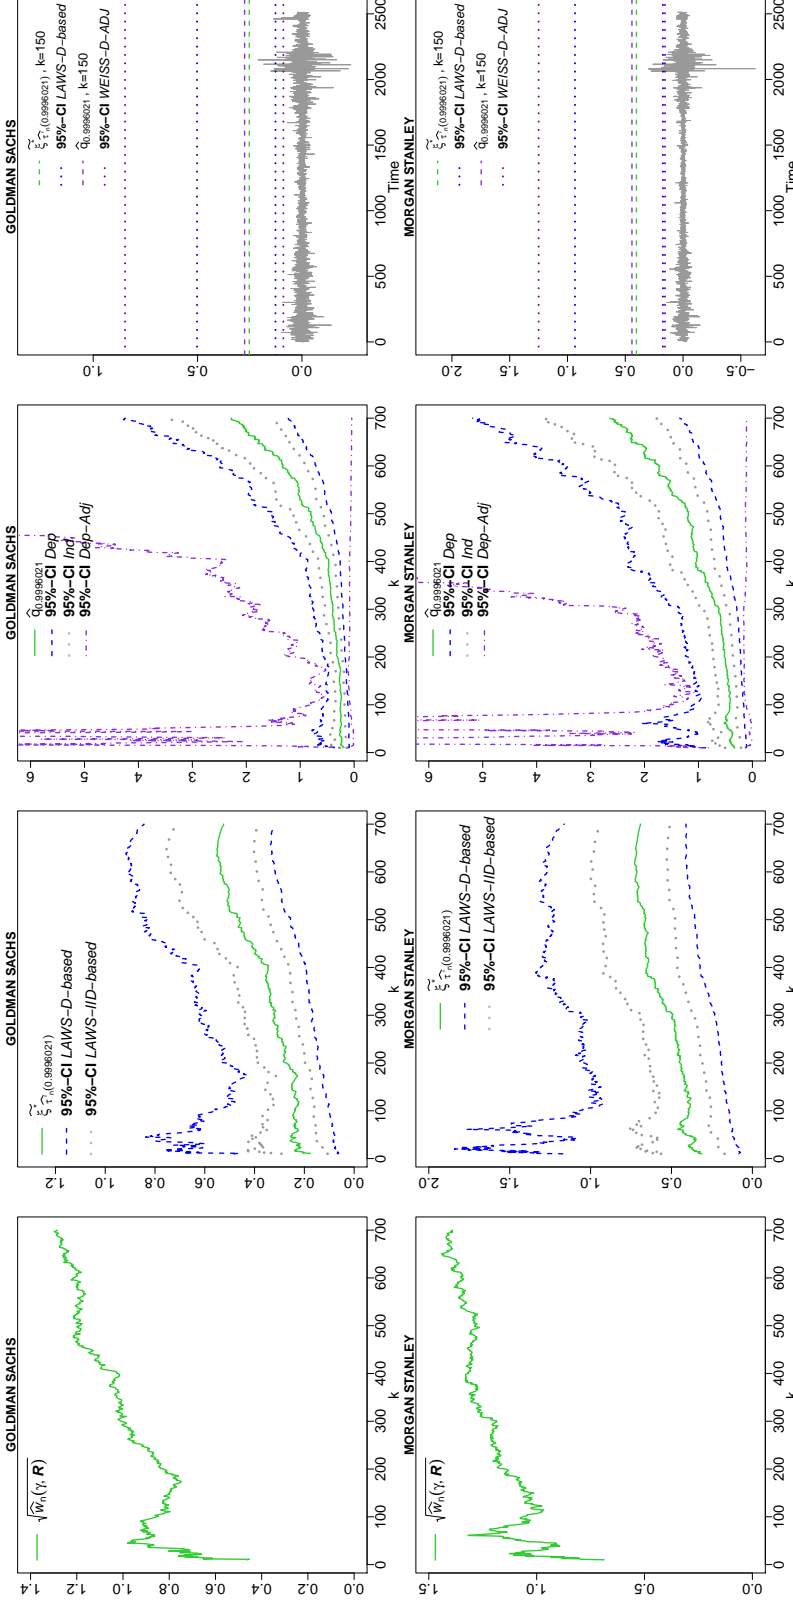


Figure 3: Goldman Sachs and Morgan Stanley loss returns data. From left to right, first panel: estimate $\widehat{\tau}'_n(0.9996021)$ (green line). Second panel: composite extrapolated LAWS estimate (green line) at level $\tau'_n = \widehat{\tau}'_n(0.9996021)$, with 95% confidence intervals based on i.i.d. theory (grey dotted line) and our LAWS-D procedure (blue dashed line). Third panel: Weissman quantile estimate (green line) at level $\alpha_n = 0.9996021$ (equivalently, composite extrapolated QB estimate at level $\tau'_n = \widehat{\tau}'_n(0.9996021)$), with 95% confidence intervals based on i.i.d. theory (grey dotted line), our LAWS-D procedure (blue dashed line) and the method of Drees (2003) (purple dashed-dotted line). Fourth panel: Negative daily log-returns data, with the composite extrapolated LAWS estimate (green dashed line) at level $\tau'_n = \widehat{\tau}'_n(0.9996021)$, compared to the standard Weissman quantile estimate $\widehat{q}_{0.9996021}$ (purple dashed line), both with 95% confidence intervals (blue dotted line and purple dotted line respectively). Top row: Goldman Sachs data, bottom row: Morgan Stanley data.




Figure 4: Left panel: Dow Jones log-returns versus S&P 500 log-returns. Right panel: Morgan Stanley loss-returns versus Goldman Sachs loss-returns.