



Extremile Regression

Abdelaati Daouia, Irène Gijbels, Gilles Stupfler

► To cite this version:

| Abdelaati Daouia, Irène Gijbels, Gilles Stupfler. Extremile Regression. 2021. hal-03104965

HAL Id: hal-03104965

<https://hal.science/hal-03104965>

Preprint submitted on 10 Jan 2021

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.

EXTREMILE REGRESSION

Abdelaati DAOUIA, Irène GIJBELS and Gilles STUPFLER[‡]

Toulouse School of Economics, France

KU Leuven, Belgium

ENSAI & CREST, France

Abstract

Regression extremiles define a least squares analogue of regression quantiles. They are determined by weighted expectations rather than tail probabilities. Of special interest is their intuitive meaning in terms of expected minima and maxima. Their use appears naturally in risk management where, in contrast to quantiles, they fulfill the coherency axiom and take the severity of tail losses into account. In addition, they are comonotonically additive and belong to both the families of spectral risk measures and concave distortion risk measures. This paper provides the first detailed study exploring implications of the extremile terminology in a general setting of presence of covariates. We rely on local linear (least squares) check function minimization for estimating conditional extremiles and deriving the asymptotic normality of their estimators. We also extend extremile regression far into the tails of heavy-tailed distributions. Extrapolated estimators are constructed and their asymptotic theory is developed. Some applications to real data are provided.

Keywords: Asymmetric least squares; Extremes; Heavy tails; Regression extremiles; Regression quantiles; Tail index.

*Corresponding author.

[†]Abdelaati Daouia is Maître de Conférences, Université Toulouse 1 Capitole, TSE, 21 allée de Brienne, Manufacture des Tabacs, 31000 Toulouse, France (E-mail: abdelaati.daouia@tse-fr.eu). Irène Gijbels is Professor of Statistics, KU Leuven, Statistics and Risk Section, Celestijnenlaan, 200b - box 2400, 3001 Leuven, Belgium (E-mail: Irene.Gijbels@kuleuven.be). Gilles Stupfler is Lecturer in Statistics, ENSAI & CREST, Campus de Ker Lann, 51 Rue Blaise Pascal, 35172 Bruz Cedex, France (E-mail: gilles.stupfler@ensai.fr).

1 Introduction

A basic tool in different scientific fields for analyzing the impact of a set of regressors X on the distribution of a response Y is quantile regression. For $\tau \in (0, 1)$, the conditional τ th quantile of Y given $X = x$ is defined as the minimizer

$$q_\tau(x) \in \arg \min_{\theta \in \mathbb{R}} \mathbb{E} \{ [|\tau - \mathbb{I}(Y \leq \theta)| \cdot |Y - \theta| - |\tau - \mathbb{I}(Y \leq 0)| \cdot |Y|] \mid X = x \}, \quad (1)$$

with $\mathbb{I}(\cdot)$ being the indicator function. Subtracting $|\tau - \mathbb{I}(Y \leq 0)| \cdot |Y|$ in the expectation makes the integrand well-defined and finite without assuming $\mathbb{E}(|Y| \mid X = x) < \infty$. A disadvantage of quantile regression is that quantiles only use the information on whether an observation is below or above some specific value. However, in a financial risk management context for example, not taking into account the effective magnitude of high values of losses, might not be wise. Conditional expectiles deal with this drawback, and lead to coherent and more realistic risk measures as compared to quantile-based risk measures, as evidenced by [1] and [4], among others. The conditional τ th expectile is defined as

$$e_\tau(x) = \arg \min_{\theta \in \mathbb{R}} \mathbb{E} \{ [|\tau - \mathbb{I}(Y \leq \theta)| \cdot |Y - \theta|^2 - |\tau - \mathbb{I}(Y \leq 0)| \cdot |Y|^2] \mid X = x \}, \quad (2)$$

obtained in a similar way to $q_\tau(x)$ in (1) but replacing absolute deviations by squared deviations (Newey and Powell [8]). Expectiles depend on both the tail realizations and their probability, while quantiles only depend on the frequency of tail observations. An inconvenience of expectiles is their lack of transparent interpretation, due to the absence of a closed form expression of $e_\tau(x)$ as a solution to the asymmetric least squares problem (2), for all $\tau \neq \frac{1}{2}$. The absence of an explicit expression makes the treatment of expectiles a hard mathematical problem from the perspective of extreme value theory, for instance when it comes to estimating tail risk (Daouia *et al.* [4]).

Very recently, Daouia *et al.* [3] considered an alternative class to expectiles, called *extremiles*, which defines a new least squares analogue of quantiles. A starting point for the introduction of this class was that the unconditional τ th quantile of Y , with continuous cumulative distribution function F , can alternatively be obtained from

$$q_\tau \in \arg \min_{\theta \in \mathbb{R}} \mathbb{E} \{ J_\tau(F(Y)) \cdot [|\theta - Y| - |Y|] \}, \quad (3)$$

where $J_\tau(\cdot) = K'_\tau(\cdot)$, with

$$K_\tau(t) = \begin{cases} 1 - (1 - t)^{s(\tau)} & \text{if } 0 < \tau \leq 1/2 \\ t^{r(\tau)} & \text{if } 1/2 \leq \tau < 1 \end{cases} \quad (4)$$

being a distribution function with support $[0, 1]$, and $r(\tau) = s(1 - \tau) = \log(1/2)/\log(\tau)$. See Section 2.1 in [3]. The *unconditional extremile of order τ* is then defined by substituting the absolute deviations with squared deviations, *i.e.*

$$\xi_\tau = \arg \min_{\theta \in \mathbb{R}} \mathbb{E} \left\{ J_\tau(F(Y)) \cdot [|Y - \theta|^2 - |Y|^2] \right\}. \quad (5)$$

Unlike expectiles, extremiles can be motivated via several angles and enjoy various interpretations and closed form expressions. For an overview on this issue, and the specific merits related to these interpretations and explicit expressions, see Daouia *et al.* [3]. In the presence of covariates, one can define conditional extremiles by considering a conditional version of (5). It will be evidenced in Section 2 that conditional extremiles enjoy the same advantages as unconditional extremiles. Obviously statistical inference for conditional quantities, such as conditional quantiles, expectiles and extremiles, requires specific regression tools as compared to statistical inference for their unconditional counterparts.

The aim of this paper is to study conditional extremiles, *i.e.* to pursue extremile regression, in a general setting. The main contributions of this paper consist of (i) discussing probabilistic properties of regression extremiles; (ii) studying and establishing the asymptotic behaviour of their nonparametric estimators; (iii) investigating conditional extremile estimators when applied to the far tail (case $\tau = \tau_n \rightarrow 1$, as the sample size $n \rightarrow \infty$); and (iv) illustrating the practical use of noncentral conditional extremiles. We shall discuss below in Section 5 the various merits of extremile regression.

The paper is organized as follows. Section 2 presents the class of regression extremiles and their basic probabilistic properties. Section 3 deals with estimation of ordinary conditional extremiles for fixed orders τ . Extrapolated estimators of tail regression extremiles, for high orders $\tau = \tau_n \rightarrow 1$ as $n \rightarrow \infty$, are developed in Section 4 for heavy-tailed conditional distributions. Section 5 concludes. All the necessary mathematical proofs, practical implementation guidelines and simulation results are given in the supplementary file.

2 Class of regression extremiles

Let $X \in \mathbb{R}^d$ and $Y \in \mathbb{R}$ be two random variables. Denote by $F(\cdot|x)$ the cumulative distribution function of Y given $X = x$ and by $q_\tau(x) = F^{-1}(\tau|x) = \inf\{y \in \mathbb{R} | F(y|x) \geq \tau\}$ the related conditional quantile of order $\tau \in (0, 1)$. For ease of presentation, we assume throughout the paper that $F(\cdot|x)$ is continuous. The order- τ extremile of this distribution function, as introduced in (5), defines the regression τ th extremile of Y given $X = x$.

Definition 1 *Let Y given $X = x$ have a finite absolute first moment. Then, for any $\tau \in (0, 1)$, the conditional order- τ extremile of Y given $X = x$ is*

$$\xi_\tau(x) = \arg \min_{\theta \in \mathbb{R}} \mathbb{E} \{ J_\tau(F(Y|X)) \cdot [|Y - \theta|^2 - |Y|^2] | X = x \}. \quad (6)$$

Particularly useful is to look at $\xi_\tau(x)$ as the following probability-weighted moment, expected maximum or expected minimum.

Proposition 1 *Let Y given $X = x$ have a finite absolute first moment. Then, for any $\tau \in (0, 1)$, we have the following equivalent closed form expressions:*

$$\begin{aligned} \xi_\tau(x) &= \mathbb{E}[Y J_\tau(F(Y|X)) | X = x] = \int_0^1 J_\tau(t) q_t(x) dt = \int_0^1 q_t(x) dK_\tau(t), \\ \text{and } \xi_\tau(x) &= \begin{cases} \mathbb{E}[\max(Y_x^1, \dots, Y_x^r)] & \text{when } \tau = (1/2)^{1/r} \text{ with } r \in \mathbb{N} \setminus \{0\}, \\ \mathbb{E}[\min(Y_x^1, \dots, Y_x^s)] & \text{when } \tau = 1 - (1/2)^{1/s} \text{ with } s \in \mathbb{N} \setminus \{0\}, \end{cases} \end{aligned}$$

for independent observations Y_x^i drawn from the conditional distribution of Y given $X = x$.

In the central case $\tau = 1/2$, we have $r(\tau) = s(\tau) = 1$, and hence $\xi_\tau(x)$ reduces to the standard regression mean $\mathbb{E}(Y|X = x)$. The limit cases $\tau \uparrow 1$ (i.e. $r(\tau) \rightarrow \infty$) and $\tau \downarrow 0$ (i.e. $s(\tau) \rightarrow \infty$) lead to the upper and, respectively, lower endpoints of the support of $F(\cdot|x)$. Further important properties are established in the following.

Proposition 2 (i) *If Y given $X = x$ has a symmetric distribution with finite absolute first moment, then $\xi_{1-\tau}(x) = 2\mathbb{E}(Y|X = x) - \xi_\tau(x)$, for any $\tau \in (0, 1)$.*

(ii) *If $Y = m(X) + \sigma(X)\varepsilon$, where $\sigma(X) > 0$ and ε is independent of X and has a finite absolute first moment, then $\xi_\tau(x) = m(x) + \sigma(x)\xi_{\tau,\varepsilon}$, for any $\tau \in (0, 1)$, where $\xi_{\tau,\varepsilon}$ denotes the τ th extremile of ε .*

An implication of Proposition 2 is that, for symmetric conditional distributions, the lower and upper extremile curves are symmetric about the regression mean. Also, the extremile curves are parallel to each other if the distribution of the response is homogeneous. These properties hold for conditional quantiles as well.

3 Estimation method

Our approach is a local linear estimation based on the definition (6) which is of particular relevance when considering flexible regression specifications such as local polynomial approximations. We restrict our analysis here to one-dimensional covariates X ($d = 1$).

3.1 Least squares kernel smoothing

For a generic estimator $\hat{F}(\cdot|x)$ of $F(\cdot|x)$, the local linear check function minimization solves the weighted least squares problem

$$\arg \min_{(\alpha, \beta) \in \mathbb{R}^2} \sum_{i=1}^n J_\tau \left(\hat{F}(Y_i|x) \right) \{Y_i - \alpha - \beta(x - X_i)\}^2 L \left(\frac{x - X_i}{h_n} \right) \quad (7)$$

to get the estimators $\check{\alpha} = \check{\xi}_{LL,\tau}(x)$ and $\check{\beta} = \check{\xi}'_{LL,\tau}(x)$ of $\xi_\tau(x)$ and $\xi'_\tau(x)$, respectively, where $L(\cdot)$ is a kernel function and $h_n > 0$ a bandwidth sequence. Standard weighted least squares theory leads to the following explicit solution

$$\begin{pmatrix} \check{\alpha} \\ \check{\beta} \end{pmatrix} = \left(\mathbf{X}_{LL}^T \mathbf{W}_{\hat{F},L} \mathbf{X}_{LL} \right)^{-1} \mathbf{X}_{LL}^T \mathbf{W}_{\hat{F},L} \mathbf{Y},$$

where \mathbf{Y} is the column vector of dimension n containing all Y_i , $i = 1, \dots, n$, and \mathbf{X}_{LL} is the usual design matrix of the local linear fitting technique, *i.e.* the $n \times 2$ matrix with a vector of 1's as a first column, and where the second column consists of the values $x - X_i$, $i = 1, \dots, n$. Furthermore, the weight matrix in the weighted least squares problem is

$$\mathbf{W}_{\hat{F},L} = \text{diag} \left(J_\tau \left(\hat{F}(Y_i|x) \right) L \left(\frac{x - X_i}{h_n} \right) \right)_{i=1, \dots, n}.$$

Clearly, the asymptotic behavior of $\hat{F}(\cdot|x)$ will be crucial to the analysis of the asymptotic and finite-sample behavior of $\check{\xi}_{LL,\tau}(x)$. Let us first discuss the properties of the latter estimator under some general high-level conditions, including the following assumptions:

(C1) The random vector (X, Y) has a joint density $f_{(X,Y)}$ which is twice continuously differentiable in its first argument and such that for each x_0 , we can write

$$\sup_{x \in U} \{f_{(X,Y)}(x, y) + |\partial_x f_{(X,Y)}(x, y)| + |\partial_{xx}^2 f_{(X,Y)}(x, y)|\} \leq h(y)$$

for some neighborhood U of x_0 , where h is a nonnegative measurable function satisfying $\int_{\mathbb{R}} (1 + |y|^{2+\delta}) h(y) dy < \infty$ for some $\delta > 0$;

(C2) The density f_X of X is continuous and positive on the interior of its support;

(C3) The kernel L is a symmetric and bounded density function with compact support.

Theorem 1 *Assume that conditions (C1)–(C3) hold, and that $\hat{F}(\cdot|x)$ is a uniformly consistent estimator of $F(\cdot|x)$ satisfying*

$$\begin{aligned} & \frac{1}{\sqrt{nh_n}} \sum_{i=1}^n J_\tau \left(\hat{F}(Y_i|x) \right) \{Y_i - \alpha - \beta(x - X_i)\} \left(\frac{x - X_i}{h_n} \right)^j L \left(\frac{x - X_i}{h_n} \right) \\ &= \frac{1}{\sqrt{nh_n}} \sum_{i=1}^n J_\tau (F(Y_i|x)) \{Y_i - \alpha - \beta(x - X_i)\} \left(\frac{x - X_i}{h_n} \right)^j L \left(\frac{x - X_i}{h_n} \right) + o_{\mathbb{P}}(1) \end{aligned} \quad (8)$$

locally uniformly in $(\alpha, \beta) \in \mathbb{R}^2$, for $j = 0, 1$. Let $h_n \rightarrow 0$ be such that $nh_n^5 \rightarrow 0$, as $n \rightarrow \infty$. Then, for any x interior to the support of X ,

$$\sqrt{nh_n} \left\{ \check{\xi}_{LL,\tau}(x) - \xi_\tau(x) \right\} \xrightarrow{d} \mathcal{N} \left(0, \frac{\|L\|_2^2}{f_X(x)} V_\tau(x) \right), \quad \text{as } n \rightarrow \infty, \quad (9)$$

where $V_\tau(x) = \mathbb{E} \left[J_\tau^2 (F(Y|x)) \{Y - \xi_\tau(x)\}^2 \mid X = x \right]$.

Assumption (8) is evidently the central condition to be checked as part of Theorem 1. It ensures that the asymptotic analysis of $\check{\xi}_{LL,\tau}(x)$ can be performed by replacing $\hat{F}(\cdot|x)$ with $F(\cdot|x)$. It should also be noted that in the higher-dimensional setting with d -dimensional covariates X , nonparametric estimators converge in general at the rate $\sqrt{nh^d}$, for a given bandwidth sequence $h = h_n \rightarrow 0$. The associated bias condition, which makes it possible to find the optimal rate of convergence of the estimator, is typically $nh^{d+4} \rightarrow c \in (0, \infty)$. This is realized for h having order $n^{-1/(d+4)}$, resulting in the optimal convergence rate $n^{2/(d+4)}$. This gets slower as d grows, and is an example of the well-known curse of dimensionality phenomenon.

The next corollary gives a simpler result for estimators $\hat{F}(\cdot|x)$ having the typical rate of convergence $n^{2/5}$. Examples of such estimators include the traditional Nadaraya-Watson estimator [7], the nearest-neighbor estimator [9] and the (improved) local linear estimator [5]. For those estimators $\hat{F}(\cdot|x)$, we derive the asymptotic normality of $\check{\xi}_{LL,\tau}(x)$ when estimating noncentral regression extremiles $\xi_\tau(x)$.

Corollary 1 *Assume that conditions (C1)–(C3) hold. Let $\hat{F}(\cdot|x)$ be an estimator of $F(\cdot|x)$ satisfying $n^{2/5} \sup_{y \in \mathbb{R}} |\hat{F}(y|x) - F(y|x)| = O_{\mathbb{P}}(1)$. Let finally $h_n \rightarrow 0$ be such that $nh_n^5 \rightarrow 0$, as $n \rightarrow \infty$. Then, with the notation of Theorem 1, for any x interior to the support of X and any $\tau \in (0, 1 - 1/\sqrt{2}] \cup [1/\sqrt{2}, 1)$, we have the convergence (9).*

The condition $\tau \in (0, 1 - 1/\sqrt{2}] \cup [1/\sqrt{2}, 1)$ should not be viewed as a restriction in practice. Indeed, by Proposition 1, regression extremiles in the right tail ($\tau \geq 1/2$) are most easily interpreted when the power $r(\tau) = \log(1/2)/\log(\tau)$ in (4) is an integer, since then $\xi_\tau(x) = \mathbb{E} \left[\max \left(Y_x^1, \dots, Y_x^{r(\tau)} \right) \right]$, for independent observations Y_x^i drawn from the conditional distribution of Y given $X = x$. In this case, the condition $\tau \in [1/\sqrt{2}, 1)$ is equivalent to $r(\tau) \geq 2$, and hence all expected maxima and corresponding extremiles are covered by this condition, except for the conditional expectation $\xi_{1/2}(x) = \mathbb{E}(Y|X = x)$ whose estimation obviously does not require extremile regression. Likewise, regression extremiles in the left tail ($\tau \leq 1/2$) are interpreted as $\xi_\tau(x) = \mathbb{E} \left[\min \left(Y_x^1, \dots, Y_x^{r(1-\tau)} \right) \right]$ when $r(1-\tau) \in \mathbb{N} \setminus \{0\}$. In this case, the condition $\tau \in (0, 1 - 1/\sqrt{2}]$ is equivalent to $r(1-\tau) \geq 2$, and so apart from $\xi_{1/2}(x) = \mathbb{E}(Y|X = x)$, all expected minima and corresponding extremiles are covered by this condition.

3.2 Empirical data examples

We now illustrate the usefulness of extremile regression on two real datasets about triceps skinfold variation and motorcycle insurance payouts. The first dataset ‘*dataTriceps*’, kindly sent by Keming Yu, comprises triceps skinfold measurements of 892 girls and women up to age 50, recorded in three Gambian villages during the dry season of 1989. To understand the evolution of triceps skinfold with age, Yu and Jones [12] proposed to

look at a collection of estimated quantiles as a function of age. The obtained fits via local linear check function minimization are graphed in Figure 1 (left panel). To calculate these conditional quantiles, we used the locally polynomial quantile regression function `lprq` of the R package `quantreg`, in conjunction with the optimal bandwidth h_{q_τ} chosen by the Yu and Jones [12] selection method. The competing conditional extremiles $\check{\xi}_{LL,\tau}$ in (7), obtained with the bandwidth h_{ξ_τ} from our automatic selection strategy developed in Section B of the supplementary file, are given in the same figure (left panel) in solid lines, along with some 95% pointwise asymptotic confidence intervals in dashed lines. To calculate the probability weights $J_\tau \left(\hat{F}(Y_i|x) \right)$ in $\check{\xi}_{LL,\tau}(x)$, we used in all our examples the local linear estimator $\hat{F}(\cdot|x)$. In the absence of a rule-of-thumb bandwidth selection for estimated expectiles via the check function method of Yao and Tong [11], we superimpose in the same figure (left panel) the expectile curves corresponding to h_{q_τ} (dashed lines) and those corresponding to h_{ξ_τ} (solid lines); the difference between the resulting expectile curves is negligible although h_{q_τ} and h_{ξ_τ} are appreciably different for each τ .

The messages yielded by the three regression methods are broadly similar, indicating particularly that adulthood corresponds to a much greater variability in triceps skinfold compared to childhood. Still, expectiles beyond the regression mean exhibit less evidence of the obvious variation and over-dispersion of the triceps skinfold as age increases: the widening of extreme expectiles seems to be rather “narrow”. By contrast, there is a distinct tendency for the noncentral extremiles and quantiles to be more spread, suggesting better capability of fitting both location and sparseness in data points. That said, extremile regression seems to be beneficial at least in producing smoother and more pleasing fits of conditional location and spread beyond the regression mean. Of course, the quantile curves can be smoothed by resorting to local linear double-kernel smoothing, but this is unnecessary for extremiles. Moreover, we are not aware of any ready-made procedure for constructing asymptotic confidence intervals of conditional quantiles and expectiles based directly on the limit distributions of their local linear estimators.

The advantages of extremile regression at the tails become even more pronounced when considering heavy-tailed scenarios as is the case in most actuarial and financial applica-

tions. The second dataset ‘*dataOhlsson*’, available from the R package `insuranceData`, contains 670 motorcycle-related claims recorded from 1994 to 1998 by the Swedish insurer Wasa. The scatterplot and local linear fits are given in Figure 1 (right panel). Here, tail regression extremiles show more alertness and reactivity to unexpected high losses than their expectile counterparts. They also exhibit better smoothness and stability than their quantile competitors and do not show any crossing, unlike the unpleasant quantile crossings that are incompatible with what occurs at the population level.

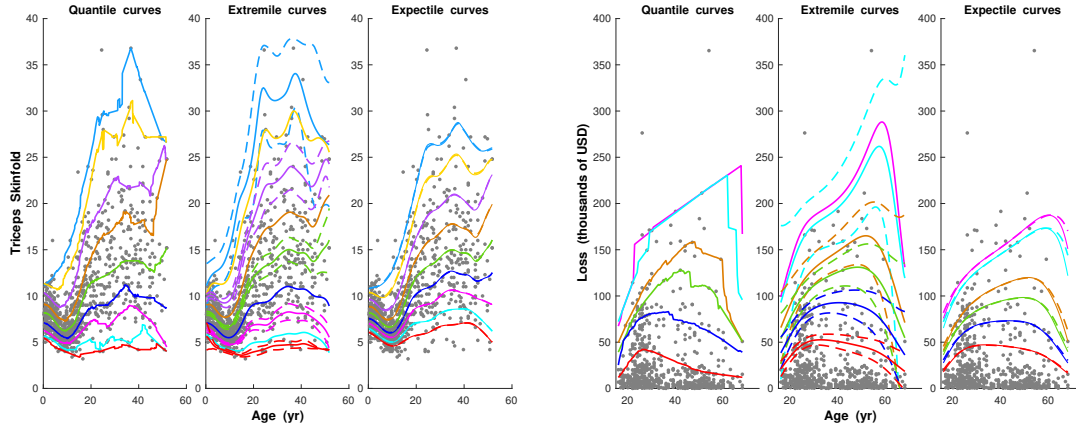


Figure 1: *Left panel: dataTriceps, with smoothed 1%, 3%, 10%, 25%, 50%, 75%, 90%, 97% and 99% quantile (left), extremile (middle) and expectile curves (right) in solid lines, and 95% pointwise asymptotic confidence intervals for $\xi_{.01}$, $\xi_{.1}$, $\xi_{.5}$, $\xi_{.9}$, $\xi_{.99}$ in dashed lines. Right panel: dataOhlsson, with smoothed 75%, 90%, 95%, 97%, 99% and 99.2% regression curves (solid), and confidence intervals for $\xi_{.75}$, $\xi_{.9}$, $\xi_{.95}$, $\xi_{.97}$ and $\xi_{.99}$ (dashed).*

Note that Proposition 1 provides a straightforward interpretation of the regression extremile $\xi_\tau(x)$ by making use of the asymmetry level forms $\tau = (1/2)^{1/r(\tau)}$, for $\tau \geq 1/2$, and $\tau = 1 - (1/2)^{1/r(1-\tau)}$ for $\tau \leq 1/2$. Intuitively, for example in the case of motorcycle insurance claims, in the right tail, $\xi_\tau(x) \equiv \mathbb{E}[\max(Y_x^1, \dots, Y_x^{r(\tau)})]$ gives the expected maximum claim amount among $r(\tau)$ potential claimants aged x years, with $r(.97) \approx 22.75$, $r(.99) \approx 68.96$, $r(.992) \approx 86.29$, $r(.993) \approx 98.67$, and $r(.994) \approx 115.17$. Interestingly, the regression quantile of the same order τ has the “dual” intuitive meaning as $q_\tau(x) \equiv \text{Median}[\max(Y_x^1, \dots, Y_x^{r(\tau)})]$. For a general level τ , we still keep the intuitive

meaning of $\xi_\tau(x)$ as an expected maximum on the right tail ($\frac{1}{2} \leq \tau < 1$) in the sense that

$$\mathbb{E} \left[\max(Y_x^1, \dots, Y_x^{\lfloor r(\tau) \rfloor}) \right] \leq \xi_\tau(x) \leq \mathbb{E} \left[\max(Y_x^1, \dots, Y_x^{\lfloor r(\tau) \rfloor + 1}) \right],$$

where $\lfloor \cdot \rfloor$ denotes the floor function.

4 Extremal regression

In this section, we focus on *extremal regression* of a response variable $Y \in \mathbb{R}$ given a vector of covariates $X \in \mathbb{R}^d$. This translates into considering the order $\tau = \tau'_n \rightarrow 1$ or $\tau'_n \rightarrow 0$ as the sample size n goes to infinity. To ease the presentation, we restrict our extreme-value analysis to the case $\tau \rightarrow 1$. Similar considerations evidently apply to the left tail $\tau \rightarrow 0$.

4.1 Model assumption

We assume for the sake of simplicity that the response Y given $X = x$ is positive and $\mathbb{E}(Y|X = x) < \infty$. We focus on the challenging domain of attraction of heavy-tailed conditional distributions that better describe the tail structure and sparseness of the data in most applications in financial and natural sciences [2, 6, 10]. More precisely, we assume that the conditional tail quantile function $t \mapsto q_{1-t^{-1}}(x)$ is second-order regularly varying:

$$(E) \quad \forall y > 0, \quad \lim_{t \rightarrow \infty} \frac{1}{A(t|x)} \left(\frac{q_{1-(ty)^{-1}}(x)}{q_{1-t^{-1}}(x)} - y^{\gamma(x)} \right) = y^{\gamma(x)} \frac{y^{\rho(x)} - 1}{\rho(x)}$$

for some parameters $0 < \gamma(x) < 1$, $\rho(x) \leq 0$ and an auxiliary function $A(\cdot|x)$ having constant sign, with $A(t|x) \rightarrow 0$ as $t \rightarrow \infty$. We use throughout the convention that $(y^b - 1)/b = \log y$ for $b = 0$, so that the right-hand side reads $y^{\gamma(x)} \log y$ if the second-order parameter $\rho(x)$ is zero. The index $\gamma(x) > 0$ tunes the tail heaviness of the conditional distribution of Y given $X = x$, with higher positive values indicating heavier conditional tails. The assumption $\gamma(x) < 1$ is tailored to our requirement that $\mathbb{E}(Y|X = x) < \infty$.

4.2 Estimation procedure and main results

Here we consider the estimation of $\xi_\tau(x)$ when $\tau = \tau'_n \uparrow 1$ at an arbitrary rate as $n \rightarrow \infty$. Under assumption (E), we have by Proposition 3 of [3], applied to the conditional

distribution of Y given $X = x$, that $\xi_{\tau'_n}(x) \sim q_{\tau'_n}(x) \mathcal{G}(\gamma(x))$ as $n \rightarrow \infty$, where $\mathcal{G}(s) := \Gamma(1-s)\{\log 2\}^s$ and Γ is the Gamma function. This motivates the estimator

$$\hat{\xi}_{\tau'_n}^*(x) := \hat{q}_{\tau'_n}^*(x) \mathcal{G}(\hat{\gamma}(x)) \quad (10)$$

obtained by substituting in suitable estimators $\hat{q}_{\tau'_n}^*(x)$ of $q_{\tau'_n}(x)$ and $\hat{\gamma}(x)$ of $\gamma(x)$. Non-parametric local estimates of the tail quantities $q_{\tau'_n}(x)$ and $\gamma(x)$ have been proposed in the last decade by [2, 6, 10], among others. Prominent among these contributions is the Weissman quantile-type estimator

$$\hat{q}_{\tau'_n}^*(x) \equiv \hat{q}_{\tau'_n, \tau_n}^*(x) := \left(\frac{1 - \tau'_n}{1 - \tau_n} \right)^{-\hat{\gamma}(x)} \hat{q}_{\tau_n}(x), \quad (11)$$

where $\hat{\gamma}(x)$ and $\hat{q}_{\tau_n}(x)$ are consistent estimators of $\gamma(x)$ and $q_{\tau_n}(x)$, with $\tau_n < \tau'_n$ being a tuning sequence to be selected jointly with h_n . Combining (10) and (11), we arrive at

$$\hat{\xi}_{\tau'_n}^*(x) \equiv \hat{\xi}_{\tau'_n, \tau_n}^*(x) = \left(\frac{1 - \tau'_n}{1 - \tau_n} \right)^{-\hat{\gamma}(x)} \hat{q}_{\tau_n}(x) \mathcal{G}(\hat{\gamma}(x)). \quad (12)$$

In Theorem A.1 in the Supplementary Material document, we establish the asymptotic distribution of $\hat{\xi}_{\tau'_n}^*(x)$ in its general form (12), for generic estimators $\hat{\gamma}(x)$ and $\hat{q}_{\tau_n}(x)$, under standard assumptions in the literature of conditional extremes. Here, we specialize the discussion to well-specified estimators $\hat{q}_{\tau_n}(x)$ and $\hat{\gamma}(x)$ in the generic form (12) of $\hat{\xi}_{\tau'_n}^*(x)$. We consider the Nadaraya-Watson type estimator $\hat{q}_{\tau_n}(x) \equiv \hat{F}_{\text{NW}}^{-1}(\tau_n|x)$, where

$$\hat{F}_{\text{NW}}(y|x) := \sum_{i=1}^n \mathbb{I}(Y_i \leq y) L\left(\frac{x - X_i}{h_n}\right) \bigg/ \sum_{i=1}^n L\left(\frac{x - X_i}{h_n}\right). \quad (13)$$

As for the choice of the conditional tail index estimator $\hat{\gamma}(x)$, we will use in the sequel the notation $\alpha_n := 1 - \tau_n$ and $p_n := 1 - \tau'_n$, and consider the kernel estimator of [2]:

$$\hat{\gamma}(x) = \sum_{j=1}^J [\log \hat{q}_{1-t_j\alpha_n}(x) - \log \hat{q}_{1-\alpha_n}(x)] \bigg/ \sum_{j=1}^J \log(1/t_j), \quad (14)$$

where $(1 = t_1 > t_2 > \dots > t_J > 0)$ is a decreasing list of J weights. Note that, unlike [2], we do not assume differentiability of the conditional distribution function, and therefore the distribution of Y given X is allowed to have atoms. The asymptotic normality of the corresponding regression extremile estimator

$$\hat{\xi}_{1-p_n}^*(x) := \left(\frac{\alpha_n}{p_n} \right)^{\hat{\gamma}(x)} \hat{q}_{1-\alpha_n}(x) \mathcal{G}(\hat{\gamma}(x))$$

follows under the following additional regularity conditions:

(K1) The functions $1/\gamma$ and f_X are Lipschitz continuous and $x \mapsto \log \ell(y|x)/\log y$, where

$$\ell(y|x) := y^{1/\gamma(x)}[1 - F(y|x)], \text{ satisfies, for a norm } \|\cdot\| \text{ on } \mathbb{R}^d,$$

$$\exists c > 0, \exists y_0 > 1, \sup_{y \geq y_0} \left| \frac{\log \ell(y|x)}{\log y} - \frac{\log \ell(y|x')}{\log y} \right| = \sup_{y \geq y_0} \frac{1}{\log y} \left| \log \frac{\ell(y|x)}{\ell(y|x')} \right| \leq c \|x - x'\|.$$

(K2) The kernel L is a bounded density with support included in the unit ball of \mathbb{R}^d .

Theorem 2 Suppose (E) and (K1)–(K2) hold. Let $x \in \mathbb{R}^d$ be such that $f_X(x) > 0$. Assume further that $\rho(x) < 0$ and, as $n \rightarrow \infty$,

1. $\alpha_n \rightarrow 0$, $nh_n^d \alpha_n \rightarrow \infty$ and $nh_n^d p_n \rightarrow c < \infty$;
2. $\log(\alpha_n/p_n)/\sqrt{nh_n^d \alpha_n} \rightarrow 0$, $nh_n^{d+2} \alpha_n \log^2 \alpha_n \rightarrow 0$ and $\sqrt{nh_n^d \alpha_n} A((1 - \alpha_n)^{-1}|x) \rightarrow 0$.

Then we have

$$\frac{\sqrt{nh_n^d \alpha_n}}{\log(\alpha_n/p_n)} \left(\frac{\hat{\xi}_{1-p_n}^*(x)}{\xi_{1-p_n}(x)} - 1 \right) \xrightarrow{d} \mathcal{N} \left(0, \frac{\|L\|_2^2}{f_X(x)} V_J \gamma^2(x) \right) \quad \text{as } n \rightarrow \infty,$$

where

$$V_J = \left(\sum_{j=1}^J \frac{2(J-j)+1}{t_j} - J^2 \right) \bigg/ \left(\sum_{j=1}^J \log(1/t_j) \right)^2.$$

When choosing, for instance, the harmonic sequence $t_j = 1/j$, the variance of the limiting distribution is minimal for $J = 9$ with $V_9 \approx 1.25$ (see [2]). We shall discuss below concrete applications where $(t_j = 1/j)_{1 \leq j \leq 9}$ are employed with the discretized tuning parameter $\alpha_n = k/n$ for a sequence of integers $k \in [1, n)$. A data-driven method for selecting both k and the bandwidth h_n in practice is described in Supplement C.

4.3 Insurance payouts

This section returns to our motivating data set ‘*dataOhlsson*’ and explores estimation and inference for extreme risk associated with motorcycle insurance claims. It can be seen in Figure 1 that this data ($n = 670$) features heavy tails and data sparsity in the tail areas. Figure 2 (top left) plots the tail index estimates $\hat{\gamma}(x)$ versus k , for the empirical

quartiles x of X . The plot shows stability over the region $k \in [50, 90]$, which suggests to pick out the pointwise estimates $\hat{\gamma}(x)$ over this interval. The top right panel plots the final estimates $\hat{\gamma}(x)$ versus x obtained via our data-driven device (red curve), along with the estimates using $k \in \{50, 70, 90\}$. It is remarkable that the automatic selection points towards similar results as these k values from the stable region. The four estimated curves indicate tail indices $\hat{\gamma}(x) \in [0.25, 0.65]$ which, as expected, reflect a strong conditional tail heaviness. Hereafter in this extremal regression study, we focus mainly on the estimates obtained for $x \leq 55$ to mitigate data scarcity and boundary effects beyond this range.

To estimate conditional extremiles $\xi_{\tau'_n}(x)$ at extreme levels, Supplement D gives Monte Carlo evidence that the extrapolated estimates $\hat{\xi}_{\tau'_n}^*(x)$ in (12) are more efficient relative to the direct estimates $\check{\xi}_{LL, \tau'_n}(x)$ from ordinary local linear regression in (7). For $\tau'_n = .99$, the middle panel of Figure 2 plots $\hat{\xi}_{\tau'_n}^*(x)$ versus k (left), for the empirical quartiles x of X , and plots the final estimates $x \mapsto \hat{\xi}_{\tau'_n}^*(x)$ (right), obtained by using the automatic selection and three values of k (selected in the stable region $[5, 25]$ of the plots shown on the left-hand side). The data-driven method affords a smoother and more stable estimated curve (in red). This extrapolated curve $x \mapsto \hat{\xi}_{\tau'_n}^*(x)$ is superimposed in the bottom panel (left) with the curve $x \mapsto \check{\xi}_{LL, \tau'_n}(x)$ of the local linear estimator (in solid blue), along with their corresponding asymptotic pointwise 95% confidence intervals (in dashed blue for $\check{\xi}_{LL, \tau'_n}(x)$ and dashed red for $\hat{\xi}_{\tau'_n}^*(x)$). There is a substantial difference between $\hat{\xi}_{\tau'_n}^*(x)$ and $\check{\xi}_{LL, \tau'_n}(x)$, which may exceed 80,000 USD for claimants' ages x outside the interval $[30, 40]$. The extrapolated quantile (Value at Risk) estimator $\hat{q}_{\tau'_n}^*(x)$, described in (11), is also graphed in the same figure (in dashed green). It lies below the extremile competitor (in solid red), close and sometimes beyond its lower confidence bound (in dashed red), with a gap that may exceed 54,000 USD. The extrapolated estimate $\hat{\xi}_{\tau'_n}^*(x)$ shows much greater reactivity to the shape of the conditional tail, as can be observed on the right end of the plot where $\hat{\xi}_{\tau'_n}^*(x)$ visibly takes into account the few very large claims incurred by older customers to produce a more prudent measure of extreme risk relative to these older claimants.

The bottom right panel plots the extrapolated $\hat{\xi}_{\tau'_n}^*(x)$ and ordinary $\check{\xi}_{LL, \tau'_n}(x)$ estimators of the tail extremile $\xi_{\tau'_n}(x) \equiv \mathbb{E}[\max(Y_x^1, \dots, Y_x^{r(\tau'_n)})]$, in solid and dashed lines

respectively, for $\tau'_n = .992$ [*i.e.* $r(\tau'_n) \approx 86$], $\tau'_n = .995$ [*i.e.* $r(\tau'_n) \approx 138$] and $\tau'_n = .998$ [*i.e.* $r(\tau'_n) \approx 346$]. With the increase of the security level τ'_n , in contrast to the non-extrapolated estimator $\check{\xi}_{LL, \tau'_n}(x)$, the refined version $\hat{\xi}_{\tau'_n}^*(x)$ becomes clearly more alert to the claims' severity and better captures the magnitude of the two most extreme losses.

5 Concluding remarks

The use of regression extremiles appears naturally in the context of risk handling, where their interpretability is straightforward and they are fully operational in practice. Beyond their remarkable merits from the point of view of the axiomatic theory of risk measures, why should statisticians and practitioners care about extremile regression? A first distinctive property of extremile regression is that, in contrast to its quantile and expectile competitors, the local linear estimators have an explicit form that is straightforward to compute, without recourse to any approximation algorithm. A second unexpected result is that the asymptotic variance of these estimators is not merely an adaptation to the conditional setup of the asymptotic variance arising in the unconditional case from [3]. In particular, this makes inference on regression extremiles much easier than inference on regression quantiles and expectiles. We are not aware of any ready-made procedure for constructing asymptotic confidence intervals of both conditional quantiles and expectiles based on the limit distributions of their local linear estimators. A further distinctive advantage of using local linear extremiles is that they suggest better capability of fitting both location and spread of data points beyond the regression mean, especially for heavy-tailed distributions. They provide much smoother and more stable fits than their quantile counterparts and do not suffer from crossings, as illustrated through both the concrete applications on triceps skinfold variation and motorcycle insurance payouts. In the class of light-tailed conditional distributions, population extremiles and quantiles are equivalent in the tail. In this class, the merits of local linear extremile estimators lead us then to favor their use over quantile estimators. It should finally be pointed out that we restrict our local linear kernel smoothing analysis to one-dimensional covariates. Extensions of our Theorem 1 and Corollary 1 to multiple covariates are obviously of interest as well.

Acknowledgments

The authors are grateful to the Editor, an Associate Editor, and three reviewers for their very valuable comments, which led to a much improved presentation of the results. The research of A. Daouia and G. Stupfler is supported by the French National Research Agency under the grant ANR-19-CE40-0013/ExtremReg project. A. Daouia also acknowledges funding from the ANR under grant ANR-17-EURE-0010 (Investissements d’Avenir program). I. Gijbels gratefully acknowledges support from the Research Fund KU Leuven (projects GOA/12/014 and C16/20/002), and from Research Grant FWO G0D6619N (Flemish Science Foundation). G. Stupfler also acknowledges support from an AXA Research Fund Award on “Mitigating risk in the wake of the COVID-19 pandemic”.

References

- [1] Bellini, F., Klar, B., Müller, A. and Gianin, E.R. (2014). Generalized quantiles as risk measures. *Insurance Math. Econom.*, **54**, 41–48.
- [2] Daouia, A., Gardes, L., Girard, S. and Lekina, A. (2011). Kernel estimators of extreme level curves. *TEST*, **20**, 311–333.
- [3] Daouia, A., Gijbels, I. and Stupfler, G. (2019). Extremiles: A new perspective on asymmetric least squares. *J. Amer. Statist. Assoc.*, **114**, 1366–1381.
- [4] Daouia, A., Girard, S. and Stupfler, G. (2018). Estimation of tail risk based on extreme expectiles. *J. R. Stat. Soc. Ser. B*, **80**, 263–292.
- [5] Daouia, A. and Park, B.U. (2013). On Projection-Type Estimators of Multivariate Isotonic Functions. *Scand. J. Stat.*, **40**, 363–386.
- [6] Goegebeur, Y., Guillou, A. and Osmann, M. (2017). A local moment type estimator for an extreme quantile in regression with random covariates. *Comm. Statist. Theory Methods*, **46**, 319–343.
- [7] Horváth, L. and Yandell, B.S. (1988). Asymptotics of conditional empirical processes. *J. Multivariate Anal.*, **26**, 184–206.
- [8] Newey, W.K. and Powell, J.L. (1987). Asymmetric least squares estimation and testing. *Econometrica*, **55**, 819–847.
- [9] Stute, W. (1986). Conditional empirical processes. *Ann. Statist.*, **14**, 638–647.

- [10] Stupfler, G. (2013). A moment estimator for the conditional extreme-value index. *Electron. J. Stat.*, **7**, 2298–2343.
- [11] Yao, Q. and Tong, H. (1996). Asymmetric least squares regression and estimation: A nonparametric approach. *J. Nonparametr. Stat.*, **6**, 273–292.
- [12] Yu, K. and Jones, M.C. (1998). Local linear quantile regression. *J. Amer. Statist. Assoc.*, **93**, 228–237.

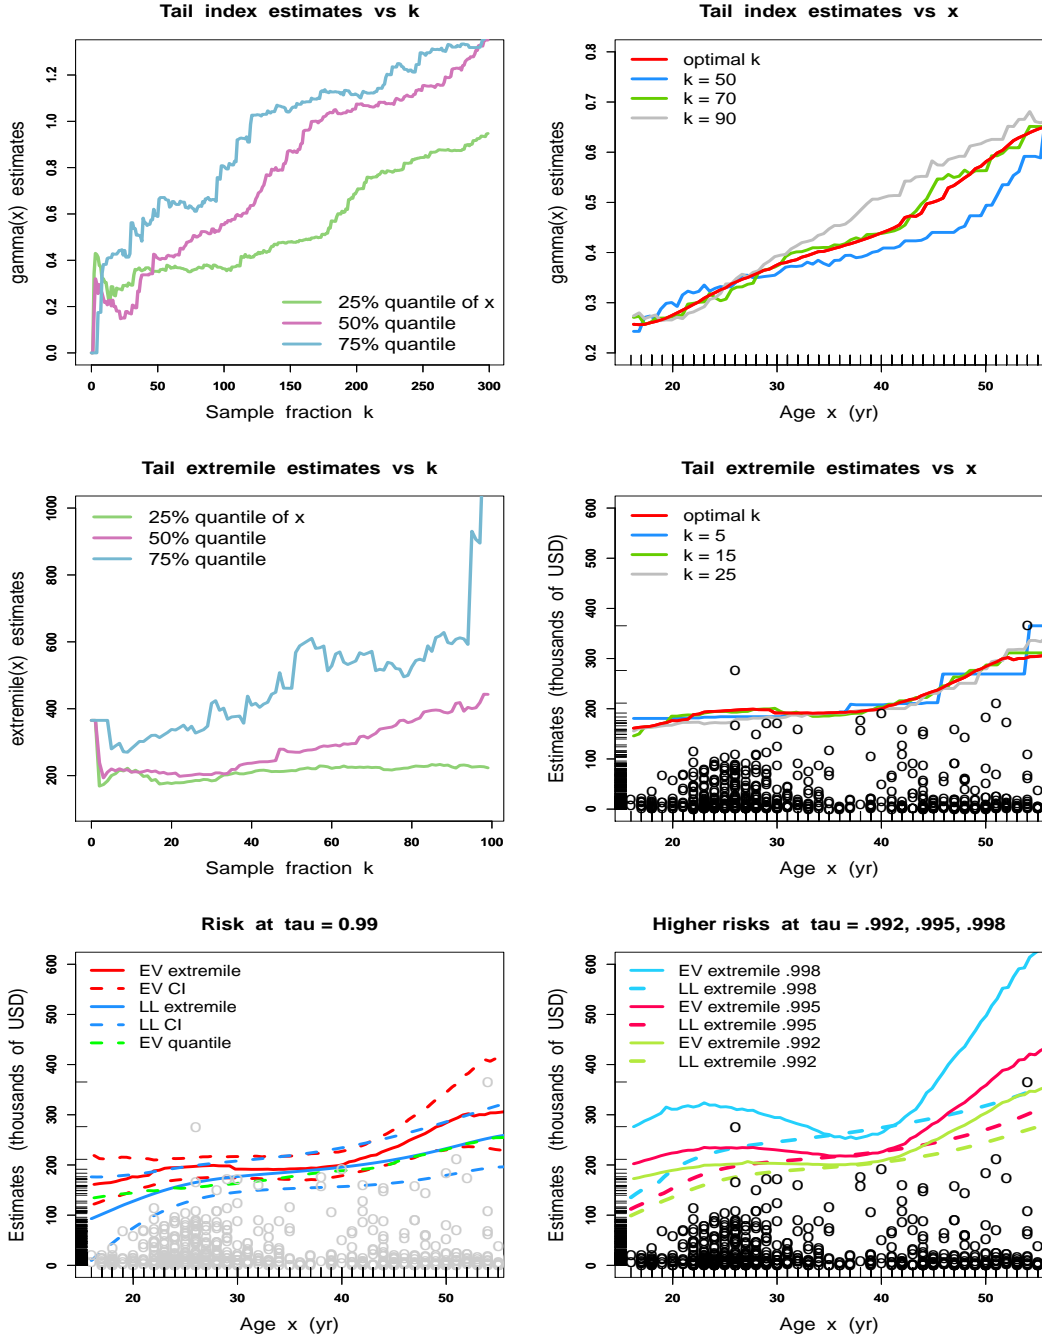


Figure 2: *Top left: plots of $\hat{\gamma}(x)$ versus k , for the quantiles of X ; Top right: estimates $x \mapsto \hat{\gamma}(x)$ obtained via the automatic selection (red) and three selected k values; Middle left: plots of $\hat{\xi}_{.99}(x)$ versus k ; Middle right: final estimates $x \mapsto \hat{\xi}_{.99}(x)$; Bottom left: Estimates $\hat{\xi}_{.99}(x)$ and $\check{\xi}_{LL,.99}(x)$ (solid red and blue), corresponding 95% confidence intervals (dashed red and blue), and $\hat{q}_{.99}^*(x)$ (dashed green); Bottom right: The estimators $\hat{\xi}_{\tau'_n}^*(x)$ (solid lines) and $\check{\xi}_{LL,\tau'_n}(x)$ (dashed lines), for $\tau'_n = .992$ (green), $\tau'_n = .995$ (red), and $\tau'_n = .998$ (blue).*