

MODEL MISSPECIFICATION IN PEAKS OVER THRESHOLD ANALYSIS

BY MÁRIA SÜVEGES AND ANTHONY C. DAVISON¹

Ecole Polytechnique Fédérale de Lausanne

Classical peaks over threshold analysis is widely used for statistical modeling of sample extremes, and can be supplemented by a model for the sizes of clusters of exceedances. Under mild conditions a compound Poisson process model allows the estimation of the marginal distribution of threshold exceedances and of the mean cluster size, but requires the choice of a threshold and of a run parameter, K , that determines how exceedances are declustered. We extend a class of estimators of the reciprocal mean cluster size, known as the extremal index, establish consistency and asymptotic normality, and use the compound Poisson process to derive misspecification tests of model validity and of the choice of run parameter and threshold. Simulated examples and real data on temperatures and rainfall illustrate the ideas, both for estimating the extremal index in nonstandard situations and for assessing the validity of extremal models.

1. Introduction. When extreme-value statistics are applied to time series it is common to proceed as though the data are independent and identically distributed, although they may be nonstationary with complex covariate effects and with rare events generated by several different mechanisms. Moreover, models that are mathematically justified only as asymptotic approximations may be fitted to data for which these approximations are poor. In this paper we suggest diagnostics for failure of these models and illustrate their application.

A standard approach to modeling the upper tail of a distribution is the so-called peaks over threshold procedure [Davison and Smith (1990)], under which a threshold u is applied to data x_1, \dots, x_n from a supposedly stationary time series, leaving N positive exceedances $x_j - u$. Extrapolation beyond the tail of the data is based on a fit of the generalized Pareto distribution [Pickands (1975)]

$$(1) \quad H(y) = \begin{cases} 1 - (1 + \xi y/\sigma)^{-1/\xi}, & \xi \neq 0, \\ 1 - \exp(-y/\sigma), & \xi = 0, \end{cases}$$

to the N exceedances, treated as independent. The parameters in (1) are a scale parameter $\sigma > 0$ and a shape parameter $\xi \in \mathbb{R}$ that controls the weight of the tail

Received July 2009; revised September 2009.

¹Supported in part by the Swiss National Science Foundation and the CCES project EXTREMES (<http://www.cces.ethz.ch/projects/hazri/EXTREMES>).

Key words and phrases. Cluster, extremal index, extreme value theory, likelihood, model misspecification, Neuchâtel temperature data, Venezuelan rainfall data.

of the distribution, whose r th moment exists only if $r\xi < 1$. In many applications ξ appears to lie in the interval $(-0.5, 0.5)$, but uncertainty about its value generally leads to alarmingly wide confidence intervals for quantities of interest such as a return period T or a return level x_T ; these satisfy $\text{pr}(X > x_T) = 1/T$. If the time series is white noise, n is large and N/n is small, then exceedances of u appear as a Poisson process, and under mild conditions we may use the tail approximation $\text{pr}(X > x) \approx (N/n)\{1 - \hat{H}(x - u)\}$, where \hat{H} is the estimate of (1). A crucial preliminary to using such methods is the choice of threshold u , which is usually performed graphically using stability properties of (1): if $Y \sim H$ and $u > 0$, then conditional on $Y > u$ the exceedance $Y - u$ has distribution (1) with parameters ξ and $\sigma' = \sigma + \xi u$; and if $\xi < 1$, then the mean residual life $E(Y - u \mid Y > u) = \sigma'/(1 - \xi)$. It is standard practice to plot empirical versions of these quantities for a range of potential thresholds, and to use only values of u above which the estimates appear to be stable and the empirical mean residual life appears to be linear. See Coles [(2001), Chapter 4] or Beirlant et al. (2004) for more details, and Beirlant, Vynckier and Teugels (1996) and Sousa and Michailidis (2004) for variants of the last plot intended to stabilize it in the presence of heavy-tailed data.

This approach to tail modeling is based on a general and well-developed probabilistic theory of extremes [Leadbetter, Lindgren and Rootzén (1983), Embrechts, Klüppelberg and Mikosch (1997), Falk, Hüsler and Reiss (2004)] and is widely used: in 2008 alone the Web of Science records around 450 articles in which the terms “peaks over threshold” or “generalised Pareto distribution” appear in the abstract or title. Thus, it is important to develop simple tools for diagnosis of the failure of such models.

One source of failure is the choice of threshold. A bad choice may yield a poor tail approximation, both because the generalized Pareto distribution is inappropriate if u is too small and because independence assumptions used to fit the model are invalid: in practice, the observations, and therefore the exceedances, are almost always dependent. This dependence is often reduced by declustering the exceedances, for example, declaring that those lying closer together than a run parameter K belong to the same cluster, and fitting (1) only to the largest exceedance of each cluster. However, a poor choice of K will give a poor inference, so it is essential to check how the results depend on the choices of threshold u and run parameter K .

A key issue is thus the effect of possible model misspecification on inference. White (1994) gives conditions under which the maximum likelihood estimator derived from a misspecified model is a consistent and normally distributed estimator of the parameter that minimizes the Kullback–Leibler discrepancy between the true and the assumed models, and constructs tests for misspecification. In the context of statistics of extremes, the peaks over the threshold model may be justified by a compound Poisson process model for the exceedances of a random process above a high threshold [Hsing (1987), Hsing, Hüsler and Leadbetter (1988)]. This model, which we outline in Section 2.1, can be checked through the projection of

the two-dimensional limiting point process of exceedances onto the time axis: if the projection is misspecified, we should be wary about using peaks over thresholds.

The main contribution of this paper is to construct diagnostics for the adequacy of peaks over threshold models. Section 2 introduces inter-exceedance times truncated by the run parameter K , which we call K -gaps, and discusses the selection of an appropriate run parameter and threshold. We propose the use of an information sandwich as a diagnostic for model failure, and, as a byproduct, we extend the maximum likelihood estimator of the reciprocal mean cluster size, the extremal index, given in Süveges (2007). Section 3 uses data simulated from two autoregressive models and a Markov chain model to illustrate the application of our ideas. Section 4 applies them to real data, to elucidate nonstationarity and tuning parameter selection and to aid extremal index estimation when the basic assumptions of extreme-value theory are violated. Section 5 contains a brief discussion.

2. Theory.

2.1. *Likelihood.* We consider asymptotic models for the upper extremes of a strictly stationary random sequence X_1, \dots, X_n with marginal distribution function F . A standard approach is to consider the limiting point process of rescaled variables $\mathcal{N}_n = \sum_{i=1}^n \delta_{i/n, (X_i - b_n)/a_n}$ as $n \rightarrow \infty$, where the sequences $\{b_n\} \subset \mathbb{R}$ and $\{a_n\} \subset \mathbb{R}_+$ are chosen so that the maximum $a_n^{-1}(\max\{X_i\} - b_n)$ has a non-degenerate limiting distribution G [Resnick (1987)]. Under mild conditions, if \mathcal{N}_n converges in distribution to a point process \mathcal{N} as $n \rightarrow \infty$, this must have the representation [Hsing (1987)]

$$\mathcal{N} = \sum_{i=1}^{\infty} \sum_{j=1}^{M_i} \delta_{(S_i, X_{ij})},$$

where (S_i, X_{i1}) are the points of a nonhomogeneous Poisson point process with mean measure $|\cdot| \times \tau(\cdot)$ on $[0, 1) \times (x_L, x_R]$, $|\cdot|$ is the Lebesgue measure, x_L and x_R are the left and right endpoints of G and $\tau(x, \infty] = -\log G(x)$. The X_{ij} are such that, for all i , the variables

$$Y_{ij} = \frac{-\log G(X_{ij})}{-\log G(X_{i1})}, \quad j = 1, \dots, M_i,$$

are the points of a point process γ_i on $[1, \infty)$ with an atom at unity. The γ_i are independent of the nonhomogeneous Poisson process (S_i, X_{i1}) and of each other, and are identically distributed. Thus, the point process limit of \mathcal{N}_n is a compound Poisson process \mathcal{N} comprising independent identically distributed clusters, the i th of which has M_i exceedances Y_{i1}, \dots, Y_{iM_i} that may have different sizes but occur simultaneously.

This result implies convergence in distribution of maxima to the generalized extreme-value distribution, convergence of threshold exceedances to the generalized Pareto distribution, and convergence of the projection $\mathcal{N}_n^* = \sum_{i=1}^n \delta_{i/n}$ to a compound Poisson process with points at S_i and marks M_i . In the limit the inter-exceedance times follow a mixture of an exponential distribution and a point mass at zero [Ferro and Segers (2003)], and this remains true for inter-exceedance times truncated by some fixed positive value; see below. For a sequence of thresholds u_n , define the inter-exceedance times in the sequence $\{X_i\}$ by

$$T(u_n) = \min\{k \geq 1 : X_{k+1} > u_n | X_1 > u_n\},$$

and the corresponding K -gaps by

$$S^{(K)}(u_n) = \max\{T(u_n) - K, 0\}, \quad K = 0, 1, \dots$$

Then Theorem 1 of Ferro and Segers (2003) can be modified to yield a limiting distribution for the K -gaps, which Süveges (2007) gave for $K = 1$. The proof requires only small modifications of the original, by considering $\text{pr}\{\bar{F}(u_n)[T(u_n) - K] > t\}$, where $\bar{F}(u_n) = 1 - F(u_n)$. Let $\mathcal{F}_{i,j}(u_n)$ denote the σ -field generated by the events $X_r \leq u_n, r = i, \dots, j$. For any $\mathcal{A} \in \mathcal{F}_{1,k}(u_n)$ with $\text{pr}(\mathcal{A}) > 0$, $\mathcal{B} \in \mathcal{F}_{k+l,n}(u_n)$ and k, l integers such that $k = 1, \dots, n - l$, define

$$\alpha^*(n, l) = \max_k \sup_{\mathcal{A}, \mathcal{B}} |\text{pr}(\mathcal{B} | \mathcal{A}) - \text{pr}(\mathcal{B})|.$$

Then we have the following result.

THEOREM 2.1. *Suppose there exist sequences of integers $\{r_n\}$ and of thresholds $\{u_n\}$ such that as $n \rightarrow \infty$, we have $r_n \rightarrow \infty$, $r_n \bar{F}(u_n) \rightarrow \tau$ and $\text{pr}\{M_{r_n} \leq u_n\} \rightarrow e^{-\theta\tau}$ for some $\tau \in (0, \infty)$ and $\theta \in (0, 1]$. Moreover, assume that there exists a sequence $l_n = o(n)$ for which $\alpha^*(cr_n, l_n) \rightarrow 0$ as $n \rightarrow \infty$ for all $c > 0$. Then as $n \rightarrow \infty$,*

$$(2) \quad \text{pr}\{\bar{F}(u_n)S^{(K)}(u_n) > t\} \longrightarrow \theta \exp(-\theta t), \quad t > 0,$$

where the extremal index θ lies in the interval $(0, 1]$ and is the reciprocal of the mean cluster size, that is, $E(M_i) = \theta^{-1}$.

Equation (2) corresponds to a limiting mixture model for the intervals between successive exceedances: with probability θ the interval is an exponential variable with rate θ , and otherwise it is of length zero, yielding a compound Poisson process of exceedance times and a likelihood-based estimator of θ . Suppose that N observations from a stationary random sequence X_1, \dots, X_n exceed the threshold u_n , let the indices $\{j_i : X_{j_i} > u_n\}$ denote the locations of the exceedances, let $T_i = j_{i+1} - j_i$ denote the inter-exceedance times, and let $S_i^{(K)} = \max(T_i - K, 0)$

denote the i th K -gap, for $i = 1, \dots, N - 1$ and $K = 0, 1, \dots$. Assuming independence of the gaps $S_1^{(K)}, \dots, S_{N-1}^{(K)}$, the limiting distribution (2) leads to the log likelihood

$$(3) \quad \ell_K(\theta; S_i^{(K)}) = (N - 1 - N_C) \log(1 - \theta) + 2N_C \log \theta - \theta \sum_{i=1}^{N-1} \bar{F}(u_n) S_i^{(K)}$$

for θ , where $N_C = \sum_{i=1}^{N-1} I(S_i^{(K)} \neq 0)$, and to a closed-form maximum likelihood estimator $\hat{\theta}_n$, which is the smaller root of a quadratic equation. Below we modify the log likelihood (3) to allow nonstationarity to be detected by using smoothing [Fan and Gijbels (1996), Süveges (2007)].

2.2. *Model misspecification.* The point process approach can fail because the assumptions of strict stationarity and independence at extreme levels are violated, but even if they are fulfilled, the chosen threshold parameter u_n and the run parameter K may be inappropriately small, thereby leading to a poor extreme-value approximation or to dependent threshold exceedances. In order to detect such difficulties, we turn to classical work on model misspecification [White (1982)]. Under broad assumptions, the maximum likelihood estimator derived from a misspecified likelihood $\ell(\theta)$ exists as a local maximum of $\ell(\theta)$. When the true model ℓ_0 is not contained in the postulated model family, that is, there is no θ_0 such that $\ell_0 = \ell(\theta_0)$, this estimator is consistent for that parameter value θ_* within the misspecified family $\ell(\theta)$ that minimizes the Kullback–Leibler discrepancy with the true distribution. Define $J(\theta) = E_0\{\ell'(\theta, S_j)^2\}$, $I(\theta) = -E_0\{\ell''(\theta, S_j)\}$, where the prime denotes differentiation with respect to θ , and E_0 means expectation under the true model, and their empirical counterparts

$$\bar{J}_n(\theta) = (N - 1)^{-1} \sum_{j=1}^{N-1} \ell'(\theta, S_j)^2, \quad \bar{I}_n(\theta) = -(N - 1)^{-1} \sum_{j=1}^{N-1} \ell''(\theta, S_j).$$

Under regularity conditions satisfied by the limiting distribution (2) when $0 < \theta < 1$, Theorem 3.2 of White (1982) yields that, as $n \rightarrow \infty$,

$$\sqrt{n}(\hat{\theta}_n - \theta_*) \xrightarrow{d} N\{0, I(\theta_*)^{-2} J(\theta_*)\}, \quad \bar{I}_n(\hat{\theta}_n)^{-2} \bar{J}_n(\hat{\theta}_n) \xrightarrow{\text{a.s.}} I(\theta_*)^{-2} J(\theta_*),$$

where \xrightarrow{d} and $\xrightarrow{\text{a.s.}}$ denote weak and almost sure convergence, respectively. Thus, the estimator derived from (3) using an arbitrary run parameter K is consistent for the value θ_* minimizing the Kullback–Leibler divergence with the true distribution, and is asymptotically normally distributed with the sandwich variance $I(\theta_*)^{-1} J(\theta_*) I(\theta_*)^{-1}$, which can be estimated by its empirical counterpart evaluated at $\hat{\theta}_n$.

It is straightforward to verify that the above theory applies to the log likelihood (3), so that as u_n increases in such a way that $N \rightarrow \infty$, the corresponding maximum likelihood estimator is consistent and asymptotically normal.

2.3. *Diagnostics: the information matrix test.* Tests to detect model misspecification may be based on the fact that the Fisher information for a well-specified regular model equals the variance of the score statistic, that is, $J(\theta) = I(\theta)$. If we write [White (1982)]

$$(4) \quad D(\theta) = J(\theta) - I(\theta),$$

then one possible misspecification test amounts to testing the null hypothesis $H_0 : D(\theta) = 0$ against the alternative $H_1 : D(\theta) \neq 0$. Let $d(s_j, \theta)$ denote the one-observation version of $D(\theta)$, let $D_n(\theta) = n^{-1} \sum_{j=1}^n d(s_j, \theta)$ denote the empirical counterpart of $D(\theta)$, and let

$$V(\theta) = E\{[d(S_j, \theta) + D'(\theta)I(\theta)^{-1}\ell'(\theta, S_j)]^2\}$$

and $V_n(\theta)$ denote the asymptotic variance of $D(\theta)$ and its empirical counterpart. Detailed formulae are given in the [Appendix](#). Under mild regularity conditions, [White \(1982\)](#) proves the following theorem, here given for a scalar parameter.

THEOREM 2.2. *If the assumed model $\ell(S_j, \theta)$ contains the true model for some $\theta = \theta_0$, then as $n \rightarrow \infty$, $\sqrt{n}D_n(\hat{\theta}_n) \xrightarrow{d} N(0, V(\theta_0))$, $V_n(\hat{\theta}_n) \xrightarrow{a.s.} V(\theta_0)$, and the test statistic $T(\hat{\theta}_n) = nD_n(\hat{\theta}_n)^2 V_n(\hat{\theta}_n)^{-1}$ is distributed as χ_1^2 .*

To check the quality of this chi-squared approximation, we performed simulations from the AR(1) and AR(2) processes described in Section 3, using choices of threshold and run parameter under which the models are well specified. Probability plots of the simulated $T(\hat{\theta}_n)$ showed that the χ_1^2 approximation is good for $N \geq 80$, and tends to be conservative if $N < 80$. Thus, relying on this approximation can lead to a loss of power when the number of exceedances is small, in which case it is difficult to detect misspecification anyway. Below we shall use the chi-squared quantiles without further comment.

3. Simulated examples. For a numerical assessment of the ideas in Section 2, we apply them to three processes:

AR(1): $Y_i = \phi Y_{i-1} + Z_i$ with $\phi = 0.7$ and Z_i standard Cauchy, with $K = 1$ and $\theta = 0.3$;

AR(2): $Y_i = \phi_1 Y_{i-1} + \phi_2 Y_{i-2} + Z_i$, with $\phi_1 = 0.95$, $\phi_2 = -0.89$ and Z_i Pareto with tail index 2, with $K = 5$ and $\theta = 0.25$;

Markov chain: With Gumbel margins, a symmetric logistic bivariate distribution for consecutive variables and dependence parameter $r = 2$ [[Smith \(1992\)](#)], with $\theta = 0.33$ and K unknown.

For each process we generated series of length $n = 8000$ and obtained sequences of inter-exceedance times; the top row of [Figure 1](#) shows a short sample with

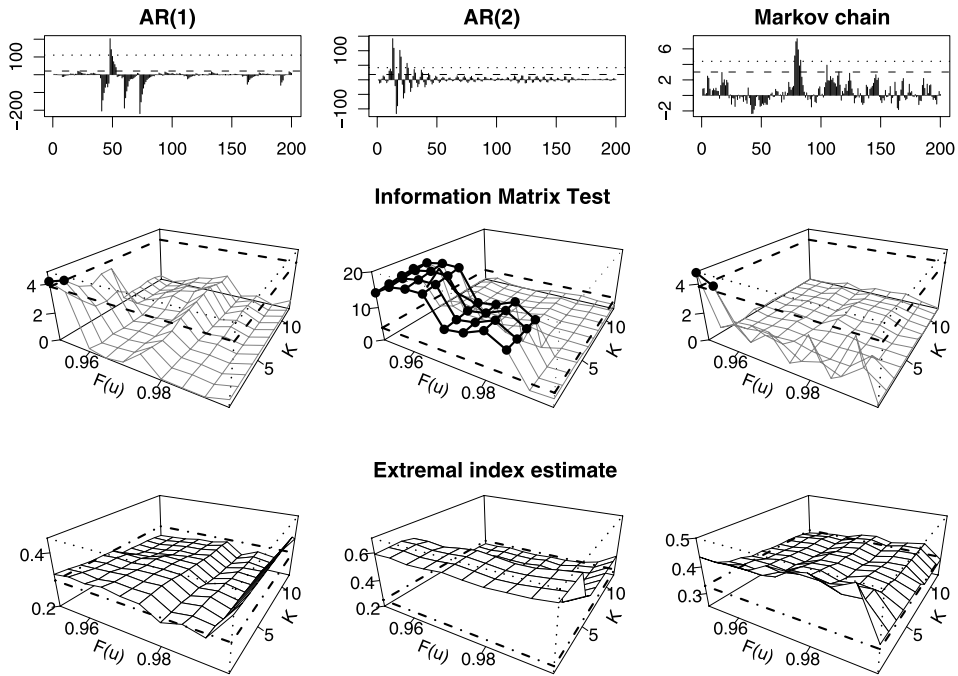


FIG. 1. Illustration of the diagnostics, based on data simulated from an AR(1) model (left column), an AR(2) model (middle column) and a Markov chain (right column). The top row gives an impression of each series, together with its 0.95 and 0.99 quantiles (dashed and dotted horizontal lines, respectively). The second row shows the information matrix test $T(\hat{\theta})$ (gray surface) and its 5% critical value $\chi_1^2(0.95) = 3.84$ (thick dashed black line around the box), with the values above 3.84 accentuated by black blobs. The third row shows the estimated extremal index $\hat{\theta}$ (black surface) as a function of the run parameter K and threshold u , with the true value of θ (thick dash-dotted black line).

a typical extreme cluster from each process. We then calculated the maximum likelihood estimates $\hat{\theta}$ for $K = 1, \dots, 12$ and thresholds corresponding to the 0.95, 0.955, \dots , 0.995 quantiles. The second row of Figure 1 shows the resulting surfaces for the information matrix statistic $T(\hat{\theta})$. The lowest panels show the estimated extremal index. For the AR(1) process, the information matrix test suggests misspecification for the combination of low thresholds with small run parameter K , but this disappears when u or K is increased. For the AR(2) process, the information matrix test indicates clear misspecification for most thresholds when $K \leq 5$, and there is then also substantial overestimation of the extremal index. The information matrix test for the Markov chain suggests that although well-specifiedness cannot be rejected for $K = 1$ and 2, inference with a larger run parameter will be more reliable. Correspondingly, the extremal index estimate is closer to the true value for larger run parameters.

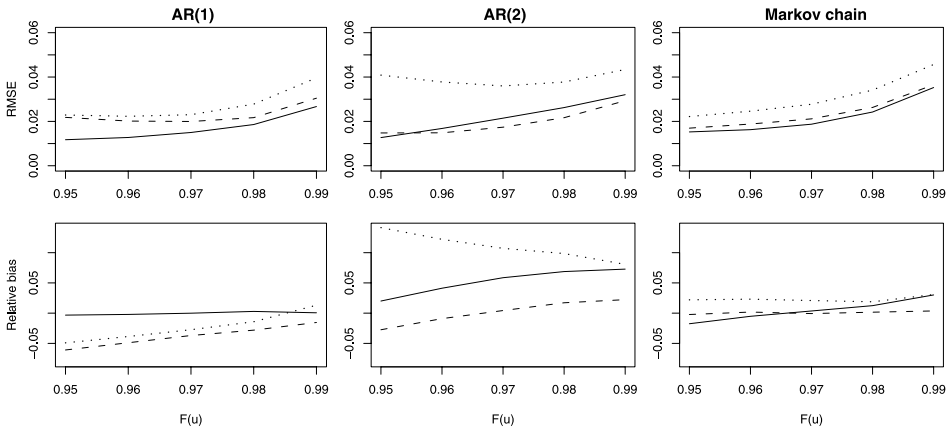


FIG. 2. Root mean squared error (top row) and relative bias (bottom row) of the K -gaps maximum likelihood (solid), the iterative least squares [Süveges (2007), dashed] and the intervals (dotted) estimators on the AR(1) process (left panels), on the AR(2) process (middle panels) and on the Markov chain (right panels). $K = 1$ for the AR(1) process, $K = 6$ for the AR(2) process, and $K = 5$ for the Markov chain. The number of observations was $n = 30,000$.

To assess the quality of the extremal index estimator based on (3), a simulation study was performed with 1000 replications of each of these processes, using $K = 1$ for the AR(1) process, $K = 6$ for the AR(2), and $K = 5$ for the Markov chain as suggested by the misspecification tests. We simulated processes of lengths $n = 2000$ and $n = 30,000$, and used thresholds corresponding to the 0.95, 0.96, 0.97, 0.98 and 0.99 quantiles. The median relative bias and the root mean squared error for the case $n = 30,000$ are shown in Figure 2. The plots confirm that if a suitable run parameter is chosen, then the maximum likelihood estimator has generally lower bias and root mean squared error than the most commonly used competitor, the intervals estimator, and another good estimator based on an iterative weighted least squares fit to the longest inter-exceedance times [Süveges (2007)].

In order to explore the behavior of the misspecification tests in the case of long-range dependence, we simulated fractionally differenced ARIMA(1, 0, d) processes of length $n = 8000$, with Gaussian white noise innovations, autoregressive parameter 0.5 and difference parameter $d = 0, 0.3$ and 0.49 . The sequence with $d = 0$ corresponds to a stationary normal AR(1) process, and therefore has extremal index $\theta = 1$. For the other cases, no theoretical calculations are known to us concerning the existence of the extremal index. The autocorrelation functions of the data, shown in the second row of Figure 3, show long memory when $d > 0$. The test statistic, plotted in the third row of Figure 3, shows a lengthening dependence range: misspecification is indicated at thresholds up to $u = F^{-1}(0.985)$ and $K < 8$ for $d = 0.3$, and at all thresholds and $K < 8$ for $d = 0.49$. The absence of misspecification at very high thresholds for $d = 0.3$ may be due to the effect of increasing the threshold while keeping the length of the series fixed.

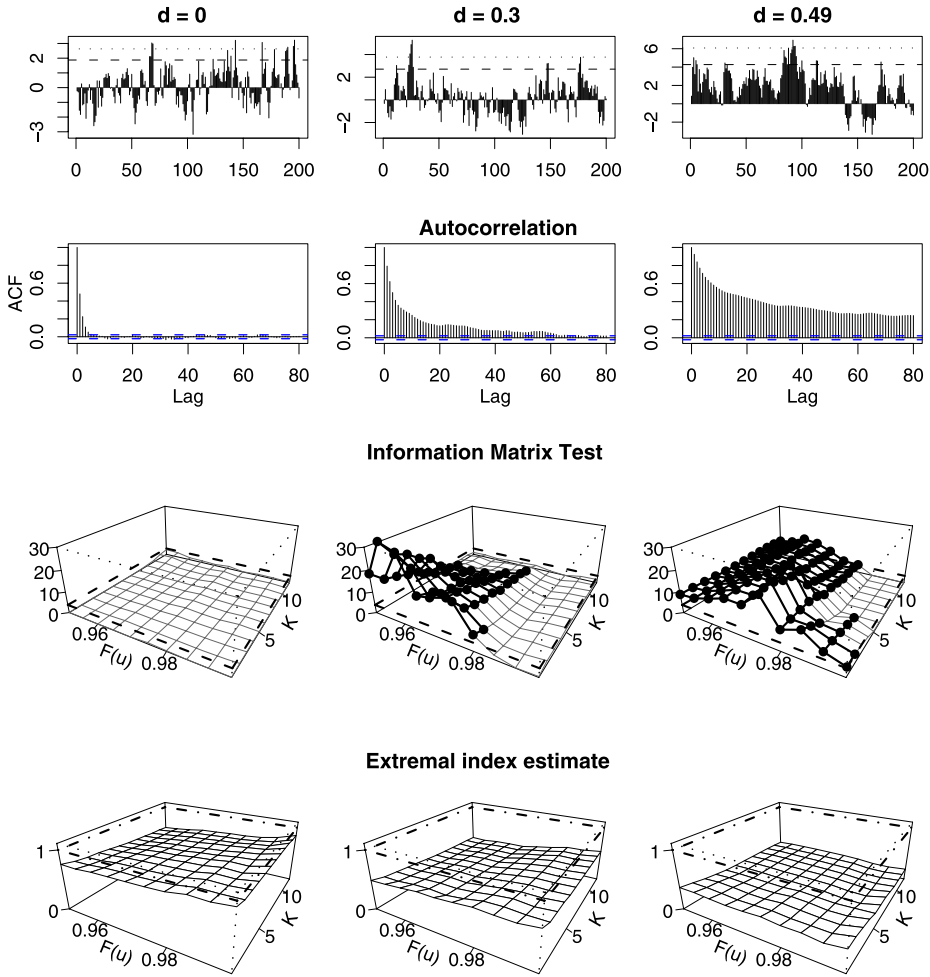


FIG. 3. The behavior of the information matrix test as a function of the dependence range of the process. The top row shows three ARIMA(1, 0, d) sequences, with $d = 0$ (left column), $d = 0.3$ (middle column) and $d = 0.49$ (right column), and the second row their autocorrelation functions. The third row contains the information matrix test $T(\hat{\theta})$ (gray surface), with the thick dashed black lines indicating the critical value for the information matrix test $\chi_1^2(0.95) = 3.84$, and the black blobs indicating $T(\hat{\theta}) > 3.84$. The fourth row presents the extremal index estimate, with the thick dash-dotted lines representing the theoretical value $\theta = 1$ for the case $d = 0$.

4. Data examples.

4.1. *Neuchâtel daily minimum summer temperatures.* For a first application to real data we use daily minimum summer temperatures from Neuchâtel from 1 January 1901 to 31 May 2006. Climatic change raises the question whether changes in temperature extremes can be summarized simply by a smooth variation of the

mean and the variance of the entire temperature distribution, or whether there are additional changes in the extremes. The daily summer minimum temperatures at Neuchâtel show a strong trend in the averages, and we investigate whether this is accompanied by a change in the clustering of the extremes. The data have been carefully homogenized, so such changes should not be due to changes in instrument siting or type, or urban effects.

We stationarized the raw data by first centering and scaling by the annual median and median absolute deviation (MAD) cycle, and then de-trending by using a ten-year moving median and MAD. Below we treat the resulting standardized temperature anomalies for the months June–August in successive years as a continuous time series. Figure 4, which shows the deseasonalized series before de-trending, shows a strong irregular variation in the mean. The presence of trend also motivates a careful misspecification analysis, not only for appropriate estimation of the extremal index, but for the assumption of stationarity.

Initially assuming stationarity of the anomalies in the 1901–2006 period, we applied the threshold selection procedures mentioned in Section 1 and described in more detail by Coles [(2001), Section 4.3] to the entire sequence. There seems to be stability above the 0.98 quantile, and generalized Pareto analysis of the complete sequence showed acceptable diagnostics. However, Figure 5, which shows these plots when applied separately to three 41-year-long periods centered on the years 1925, 1955 and 1985, casts some doubt on the overall results.

We therefore checked model misspecification as a function of time, by centering 41-year long windows successively at 15 July of each year, and calculating the information matrix test defined by equation (4), for every combination of threshold u and run parameter K . The calculations thus gave 106 sets of $T(\hat{\theta})$ values, with extremal index estimates for all (u, K) pairs for the sequence of anomalies.

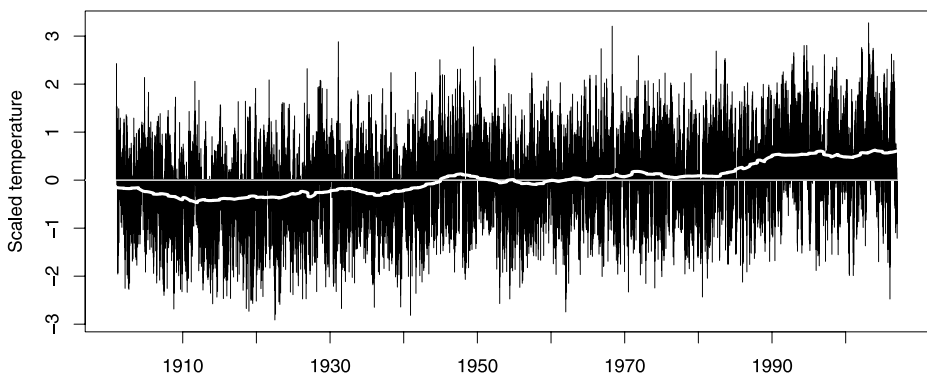


FIG. 4. *The deseasonalized Neuchâtel daily minimum temperatures (black) with their trend, estimated by a cubic spline-smoothed 10-year moving median (heavy white line). The horizontal line is intended to help appreciate the trend.*

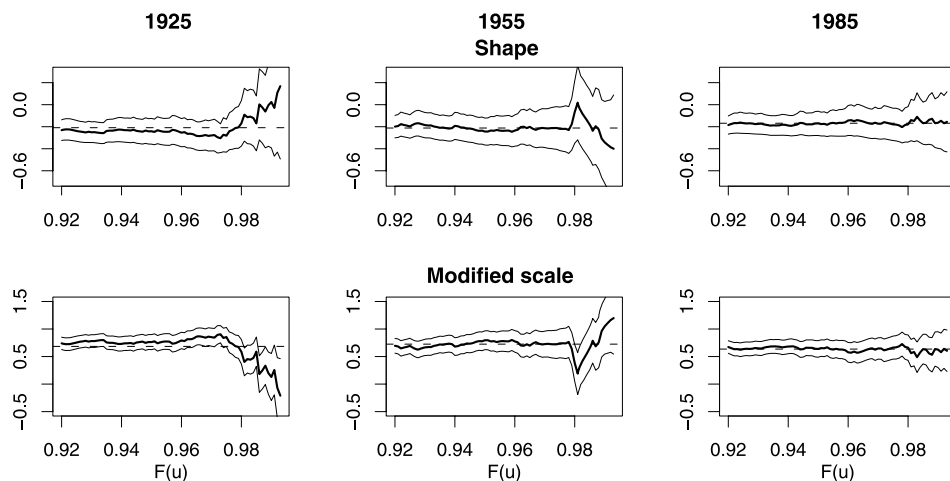


FIG. 5. Classical threshold selection plots for three 41-year windows centered on 1925, 1955 and 1985 for the Neuchâtel daily minimum temperature anomalies, showing the parameter estimate (bold) and pointwise 95% confidence limits (solid) as functions of the threshold. The dashed lines show the average estimates for the different thresholds.

The three main potential sources of misspecification here are threshold selection, the choice of run parameter and possible nonstationarity, so the test statistics proposed above must depend on these. Figure 6, which presents the surfaces of $T(\hat{\theta})$ for twelve different years, suggests misspecification in the period 1935–1970 for thresholds around the 0.97- and 0.98-quantile for all run parameters. Smaller instabilities were also found, mostly between 1985–2000, though these rarely exceeded the critical χ_1^2 -quantile. These two periods roughly coincide with the strongest nonstationarity in the summer mean temperatures: see Figure 4, in which the most marked periods of change in the 10-year median are in the 1940s and in the 1980s. The threshold suggested by classical selection methods, the 0.98-quantile, should be avoided: the ridge indicating misspecification is located at this threshold.

Although our motivation for testing is different from the usual one leading to the definition of the false discovery rate (FDR), the multiple testing setup and the dependence between information matrix tests applied successively in sliding windows should be taken into account when we want to justify the existence of a misspecification region. Our main interest does not lie in the discovery of regions where the null hypothesis of well-specifiedness is rejected at a given level of FDR [Benjamini and Hochberg (1995)], but in finding those where it cannot be rejected. Nevertheless, we tested the significance of the departure from the null hypothesis by the procedure proposed in Benjamini and Yekutieli (2001) separately for each u, K pair, and found that the misspecification is significant for all K and for $F(u) = 0.98$ between roughly 1940–1965. The information matrix test was

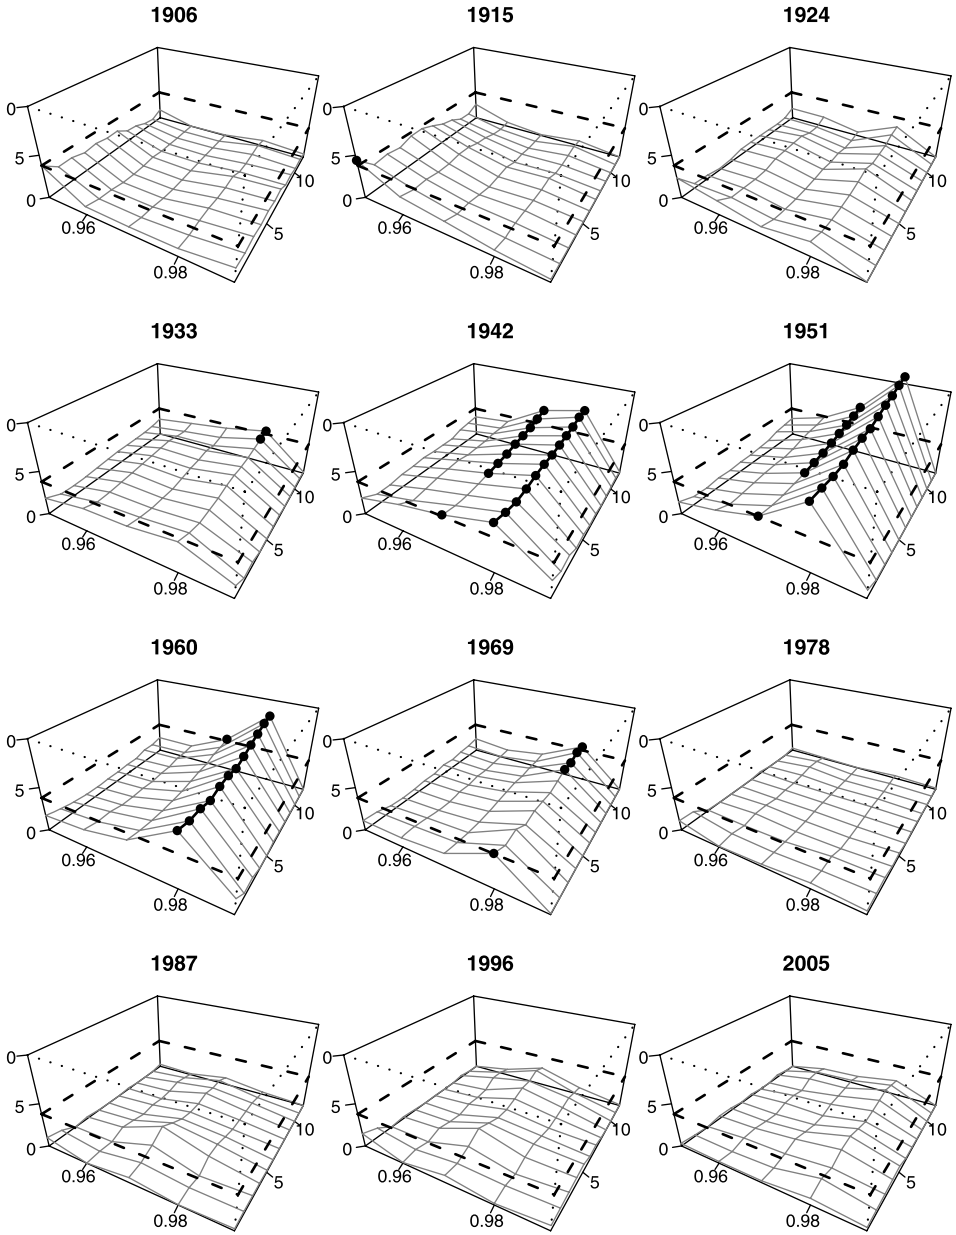


FIG. 6. Misspecification as a function of time for the Neuchatel summer daily minimum temperature anomalies. Horizontal foreground axis: threshold u as $F(u)$; horizontal left axis: run parameter K ; vertical axis: $T(\hat{\theta})$. The thick dashed lines around the box correspond to the critical 0.95-quantile of the χ^2_1 distribution, the black blobs emphasize the parameter combinations where $T(\hat{\theta}) \geq \chi^2_1(0.95)$. Years are indicated above the plots.

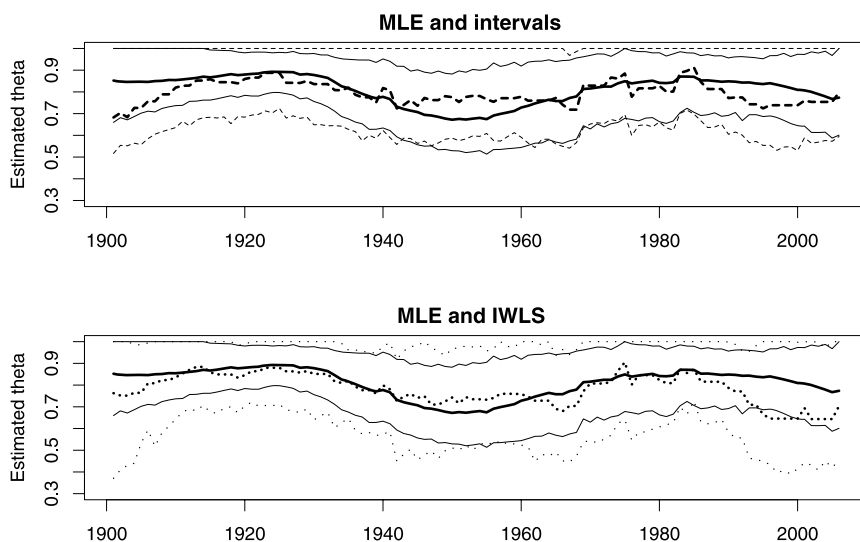


FIG. 7. Comparison of the maximum likelihood estimator using $K = 4$ and $F(u) = 0.99$ (heavy solid in both panels) to the intervals (top panel, thick dashed) and to the iterative weighted least squares (bottom panel, thick dotted) estimators. The 95% confidence intervals are shown by thin versions of the lines.

closest to zero at the point $F(u) = 0.99$, $K = 4$ over the whole century, so we chose these for smooth estimation of the extremal index. Figure 7 shows the resulting locally constant weighted K -gaps estimates, compared to intervals and to iterative weighted least squares estimates based on 41-year long sliding windows. The confidence intervals for the K -gaps estimator are based on asymptotic normality. Nonparametric bootstrap intervals were also calculated, but showed only slight differences mostly in the middle of the century, where the bootstrap interval was slightly wider. The value of θ dips in the 1950s and in recent years, but overall any evidence for changes in the clustering of summer minimum daily temperatures seems to be weak.

The information matrix test suggests the existence of fluctuations in the time point process of extreme anomalies. Using the ten-year moving window to de-trend the series, and a 41-year window for the information matrix tests and the estimation of the extremal index, only the combination of a high threshold and a relatively high run parameter seem to yield a well-specified model. The period where the models are misspecified roughly coincides with a local peak in the 10-year moving median of the data set. Changes in the median temperatures may be accompanied by changes in clustering characteristics, but perhaps using a 10-year moving window for de-trending is insufficient to remove mean fluctuations, so the anomalies display traces of residual nonstationarity that then appear in the 41-year moving windows. Our investigation thus emphasizes the importance of an appropriate treatment of long-term trends. Many studies of climate extremes

use varying thresholds based on local quantile estimation or on an assumption of a trend of simple parametric form, the first of which corresponds to our ad hoc selection of window length for de-trending; see, for example, [Kharin and Zwiers \(2005\)](#), [Nogaj et al. \(2006\)](#) or [Brown, Caesar and Ferro \(2008\)](#). Our results indicate that model quality is highly sensitive to such choices, so it is necessary to check whether the models are well specified, in order to avoid biased estimates with underestimated variances. Climatological studies commonly directly compare periods of a few decades, which are assumed stationary, but this too should be performed with care. In our study, the time-scale of the fluctuations found at extreme levels is shorter than a few decades on thresholds $u < F^{-1}(0.99)$ in the period 1940–1965. As this is a period where the global mean temperature based on observational data has a marked local peak, this might arise at other stations also, and other climate variables may also show instability on such time-scales.

4.2. Daily rainfall in Venezuela. The rainfall data recorded daily between 1 January 1961 and 31 December 1999 at Maiquetia airport in Venezuela provide a striking example of the difficulties of using simple extreme-value methods for risk estimation. Prior to December 1999, the annual maxima of the data set were fitted with a Gumbel model, with no diagnostics indicating a bad fit. However, after an unusually wet fortnight in December 1999, extensive destruction and around 30,000 deaths [[Larsen et al. \(2001\)](#)] were caused by three consecutive daily precipitation totals of 120, 410.4 and 290 mm, the largest of which was almost three times greater than the maximum of the preceding 40 years. The return period for the peak value of 410.4 mm under simple models can be expressed in thousands or even millions of years. Why do classical extreme-value methods fail so catastrophically, and could more sophisticated methods have given a different return period estimate for such an event?

[Coles and Pericchi \(2003\)](#) apply a Bayesian approach to the point process representation [[Smith \(1989\)](#)], which is essentially equivalent to fitting the generalized Pareto distribution. They use a threshold corresponding approximately to the 0.96-quantile and including all exceedances, and argue in favor of partitioning the sequence into two seasons, the one of interest being from mid-November to April. With these refinements, they obtain a predictive return period of approximately 150 years for 410.4 mm. However, the classical threshold selection plots for these months, shown in the left three panels of [Figure 8](#), indicate trouble with the model: parameter stability is compatible with the confidence intervals only for thresholds much higher than the 0.96-quantile. Ignoring the tendency of extremes to cluster may also have implications for the estimates and their variance, because the independence assumption is violated. Moreover, which estimate should we choose if different methods give very different answers?

Following [Coles and Pericchi \(2003\)](#) and backed by meteorological information, we took only the months December–April, initially excluding December 1999, and calculated our diagnostics for $u \in [F^{-1}(0.95), F^{-1}(0.995)]$ and

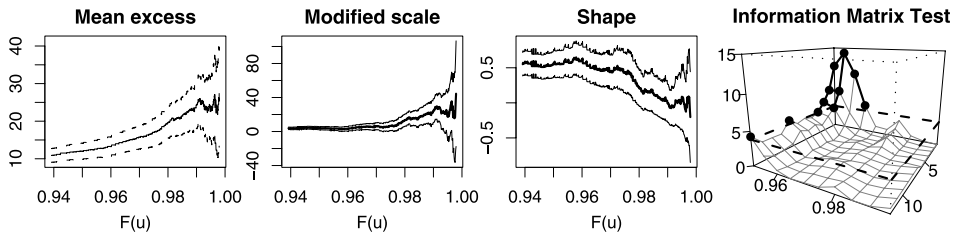


FIG. 8. Classical threshold selection plots for the Venezuelan daily rainfall data for the months December–April, between January 1961 and April 1999. The three panels on the left show the mean excess plot and the modified scale and the shape parameters of GPD fits as a function of threshold. The rightmost panel is the statistic $T(\hat{\theta})$ as a function of the run parameter K and the threshold u on the probability scale $F(u)$.

$K = 1, \dots, 12$. The graph of the statistic $T(\hat{\theta})$ in the rightmost panel of Figure 8 displays three clear features: a region of misspecification at thresholds below the 0.96-quantile; a ridge along the 0.98-quantile with relatively higher values than the surrounding region, corresponding to the most marked instability interval in the classical threshold selection plots; and higher test statistic values for $K = 1, 2$, implying misspecification for combinations with $F(u) < 0.97$, $K \leq 2$. The best regions appear to be either $F(u) \approx 0.97$, $K = 3$ or $F(u) \approx 0.99$, $K = 3$, but there is no further indication which is preferable, and because of the ridge, the existence of a contiguous area of well-specifiedness is doubtful.

The generalized Pareto model (1) was fitted with $F(u) = 0.97$, $K = 3$ and $F(u) = 0.99$, $K = 3$. The resulting parameter estimates and standard errors are very different: $\hat{\xi} = 0.27$ (0.14), $\hat{\sigma} = 14.8$ (2.4) for 0.97 quantile, and $\hat{\xi} = -0.03$ (0.14), $\hat{\sigma} = 26.6$ (5.3) for the 0.99 quantile. Corresponding diagnostic plots are shown in the upper row of Figure 9, the left two plots referring to the lower threshold, and the right two to the higher. Neither model seems poor, but they give very different return periods for the value 410.4 mm: approximately 600 years for the model with threshold $F(u) = 0.97$, and an unreasonable 30 million years for the other. The catastrophe seems to be compatible only with the lower threshold model, despite the fact that in extreme-value statistics, models above higher thresholds are generally considered to be closer to the limiting distribution. Inclusion of December 1999 does not change the misspecification tests much, but changes the estimates to $\hat{\xi} = 0.5$ (0.15), $\hat{\sigma} = 12.6$ (2.2) for the lower threshold and $\hat{\xi} = 0.27$ (0.15), $\hat{\sigma} = 25.1$ (5.2) for the higher one. The corresponding return periods become 65 and 300 years, respectively. A look at the diagnostic plots in the bottom row of Figure 9 confirms that the former model admits the catastrophe quite smoothly, whereas it remains an outlier in the latter.

One explanation of this apparent paradox might be that the underlying process is a mixture. Rainfall is generated by different atmospheric processes, such as convective storms, cold fronts or orographic winds. If so in this case, the observed

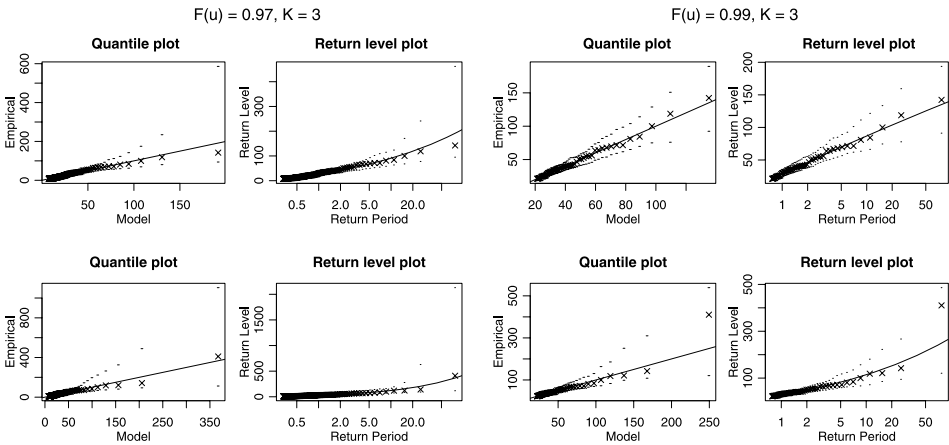


FIG. 9. Diagnostic and return level plots for the generalized Pareto fits to the Venezuelan rainfall data. The fits excluding December 1999 are shown in the top row. The left pair of plots presents the quantile-quantile and return level plots using the 0.97-quantile as the threshold, the right pair those using the 0.99-quantile. The run parameter was $K = 3$ for both. The same plots for model fits including December 1999, with the same thresholds and run parameters, are presented in the bottom row. Each panel shows the ordered data (\times) with a line representing the fitted model and a pointwise 95% envelope.

extreme process corresponds to the extremes of a mixture distribution $\sum_{i=1}^m p_i F_i$ with the component distribution functions F_i appearing with probabilities p_i , where the F_i could have different extreme-value limits: with $m = 2$, for example, F_1 might correspond to the short-tailed case $\xi = -1/2$ and F_2 to the long-tailed case $\xi = 1$. Such a mixture can show unstable behavior like that of Figure 8, which hints at the presence of at least two components: a more frequent light-tailed one with relatively high location and scale parameters, dominating the levels around the 0.99-quantile, and a rarer heavy-tailed one concentrated at lower levels and having a smaller scale parameter, but generating extremely large observations occasionally. A more sophisticated model for the clustering of extremes also suggests a mixture character, but will be reported elsewhere. This failure of simple extreme-value techniques is a warning to beware of oversimplification, and suggests that an approach linking atmospheric physics and statistical methods would provide better risk estimates.

5. Discussion. Inference about the extremal behavior of a process involves assumptions such as asymptotic independence at extreme levels and stationarity, and also entails the selection of auxiliary quantities such as threshold and run parameters in order to apply asymptotic models with finite samples. Careful investigation of possible model misspecification is therefore essential.

In this paper we have applied standard methods of detecting misspecification to the point mass-exponential mixture model (2) for the inter-exceedance times.

These tests assist in the selection of the threshold and the run parameter K and thus help to provide better estimates of both the extremal index and of the generalized Pareto distribution. Failure of the model (2) indicates failure of a more general limit, and consequently of the GPD approximation (1). Analysis of the Venezuelan rainfall data shows that misspecification tests can provide a valuable supplement to classical threshold selection procedures, can lead to improved models and better variance estimates, and may yield further insight into the structure of the data.

We have also described a maximum likelihood estimator for the extremal index, based on the point mass-exponential model (2) and on the existence of a selection procedure for K . The maximum likelihood estimator is consistent and asymptotically normal under an appropriate choice of K , and shows small asymptotic bias and root mean square error compared to the best competing estimators. The joint application of the misspecification tests and the smoothed maximum likelihood estimator proved the good properties of the procedure as an efficient method to detect violations of underlying assumptions such as nonstationarity or indicate other model problems like mixture character that cannot be disregarded using finite thresholds. It can be therefore a useful aid to fine-tuning parameters of extreme-value models or investigating their limitations.

One natural question is whether the assessment of misspecification for extremal models might better be based on (1). The difficulty with this is that the r th moment of the score statistic for ξ exists only if $r\xi > -1$, so the maximum likelihood estimators of ξ and σ are regular only if $\xi > -1/2$ [Smith (1985)] and the observed information has finite variance only if $\xi > -1/4$, and information quantities for (1) have poor finite-sample properties. The distribution (2), on the other hand, is regular for $0 < \theta < 1$, and so has no such disadvantages.

APPENDIX: FORMULAE FOR THE INFORMATION MATRIX TEST

Assume that X_1, \dots, X_n satisfy the necessary conditions for Theorem 2.1. For a threshold u_n , suppose that $N < n$ observations exceed the threshold u_n . Let the indices $\{j_i : X_{j_i} > u_n\}$ denote the times of the exceedances, and let $c_i^{(K)} = \overline{F}(u_n)s_i^{(K)} = \overline{F}(u_n) \max(j_{i+1} - j_i - K, 0)$ denote the i th observed K -gap $s_i^{(K)}$ normalized by the tail probability $\overline{F}(u_n)$ for $i = 1, \dots, N - 1$ and $K = 0, 1, \dots$. Then, denoting derivatives with respect to θ by a prime, it follows from the likelihood (3) that

$$\begin{aligned} \ell'_K(\theta; c_i^{(K)}) &= -\frac{I(c_i^{(K)} = 0)}{(1 - \theta)} + \frac{2I(c_i^{(K)} > 0)}{\theta} - c_i^{(K)}, \\ \bar{J}_n(\theta) &= (N - 1)^{-1} \sum_{j=1}^{N-1} \left\{ \frac{I(c_j^{(K)} = 0)}{(1 - \theta)^2} + \frac{4I(c_j^{(K)} > 0)}{\theta^2} + c_j^{(K)} - \frac{4c_j^{(K)}}{\theta} \right\}, \\ \bar{I}_n(\theta) &= (N - 1)^{-1} \sum_{j=1}^{N-1} \left\{ \frac{I(c_j^{(K)} = 0)}{(1 - \theta)^2} + \frac{2I(c_j^{(K)} > 0)}{\theta^2} \right\}, \end{aligned}$$

where $I(A)$ is the indicator function of the set A . Then we can derive the one-observation version, the sample mean of the difference D between the variance of the score and the Fisher information and the sample variance of the latter as

$$\begin{aligned}d(\theta; c_i^{(K)}) &= \frac{2I(c_i^{(K)} > 0)}{\theta^2} + c_j^{(K)} - \frac{4c_j^{(K)}}{\theta}, \\D_n(\theta) &= \bar{J}_n(\theta) - \bar{I}_n(\theta), \\D'_n(\theta) &= (N-1)^{-1} \sum_{j=1}^{N-1} \left\{ -\frac{4I(c_j^{(K)} > 0)}{\theta^3} + \frac{4c_j^{(K)}}{\theta^2} \right\}, \\V_n(\theta) &= (N-1)^{-1} \sum_{j=1}^{N-1} \{d(c_j^{(K)}) - D'_n \bar{I}_n(\theta)^{-1} \ell'_K(\theta; c_i^{(K)})\},\end{aligned}$$

and from there, the information matrix test statistic is obtained by substituting the appropriate quantities and the estimated value of $\hat{\theta}_n^{(K)}$ as

$$T(\hat{\theta}_n^{(K)}) = nD_n(\hat{\theta}_n^{(K)})^2 V_n(\hat{\theta}_n^{(K)})^{-1}.$$

Acknowledgments. We gratefully acknowledge helpful comments of Philippe Naveau, Jonathan Tawn, two referees, an associate editor and the editor.

REFERENCES

- BEIRLANT, J., GOEGEBEUR, Y., SEGERS, J. and TEUGELS, J. (2004). *Statistics of Extremes*. Wiley, Chichester. [MR2108013](#)
- BEIRLANT, J., VYNCKIER, P. and TEUGELS, J. L. (1996). Excess functions and estimation of the extreme-value index. *Bernoulli* **2** 293–318. [MR1440271](#)
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the False Discovery Rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. [MR1325392](#)
- BENJAMINI, Y. and YEKUTIELI, D. (2001). The control of the False Discovery Rate in multiple testing under dependency. *Ann. Statist.* **29** 1165–1188. [MR1869245](#)
- BROWN, S. J., CAESAR, J. and FERRO, C. A. T. (2008). Global changes in extreme daily temperature since 1950. *J. Geophys. Res.* **113** D05115. doi: [10.1029/2006JD008091](#).
- COLES, S. G. (2001). *An Introduction to Statistical Modeling of Extreme Values*. Springer, London. [MR1932132](#)
- COLES, S. G. and PERICCHI, L. (2003). Anticipating catastrophes through extreme-value modelling. *Appl. Statist.* **52** 405–416. [MR2012566](#)
- DAVISON, A. C. and SMITH, R. L. (1990). Models for exceedances over high thresholds (with discussion). *J. Roy. Statist. Soc. Ser. B* **52** 393–442. [MR1086795](#)
- EMBRECHTS, P., KLÜPPELBERG, C. and MIKOSCH, T. (1997). *Modeling Extremal Events for Insurance and Finance*. Springer, Berlin. [MR1458613](#)
- FALK, M., HÜSLER, J. and REISS, D. (2004). *Laws of Small Numbers: Extremes and Rare Events*. Birkhäuser, Basel. [MR2104478](#)

- FAN, J. and GIJBELS, I. (1996). *Local Polynomial Modelling and Its Applications*. Chapman & Hall, London. [MR1383587](#)
- FERRO, C. A. T. and SEGERS, J. (2003). Inference for clusters of extreme values. *J. Roy. Statist. Soc. Ser. B* **65** 545–556. [MR1983763](#)
- HSING, T. (1987). On the characterization of certain point processes. *Stochastic Process. Appl.* **26** 297–316. [MR0923111](#)
- HSING, T., HÜSLER, J. and LEADBETTER, M. R. (1988). On the exceedance point process for a stationary sequence. *Probab. Theory Related Fields* **78** 97–112. [MR0940870](#)
- KHARIN, V. V. and ZWIERS, F. W. (2005). Estimating extremes in transient climate change simulations. *Journal of Climate* **18** 1156–1173.
- LARSEN, M. C., WIECZOREK, G. F., EATON, L. S., MORGAN, B. A. and TORRES-SIERRA, H. (2001). Venezuelan debris flow and flash flood disaster of 1999 studied. *EOS, Transactions of the American Geophysical Union* **47** 572.
- LEADBETTER, M. R., LINDGREN, G. and ROOTZÉN, H. (1983). *Extremes and Related Properties of Random Sequences and Processes*. Springer, New York. [MR0691492](#)
- NOGAJ, M., YIOU, P., PAREY, S., MALEK, F. and NAVEAU, P. (2006). Amplitude and frequency of temperature extremes over the North Atlantic region. *Geophys. Res. Lett.* **33** L10801. doi: [10.1029/2003GL019019](#).
- PICKANDS, J. (1975). Statistical inference using extreme order statistics. *Ann. Statist.* **3** 119–131. [MR0423667](#)
- RESNICK, S. I. (1987). *Extreme Values, Regular Variation and Point Processes*. Springer, New York. [MR0900810](#)
- SMITH, R. L. (1985). Maximum likelihood estimation in a class of non-regular cases. *Biometrika* **72** 67–90. [MR0790201](#)
- SMITH, R. L. (1989). Extreme value analysis of environmental time series: An application to trend detection in ground-level ozone. *Statist. Sci.* **4** 367–377. [MR1041763](#)
- SMITH, R. L. (1992). The extremal index for a Markov chain. *J. Appl. Probab.* **29** 37–45. [MR1147765](#)
- SOUSA, B. and MICHAILIDIS, G. (2004). A diagnostic plot for estimating the tail index of a distribution. *J. Comput. Graph. Statist.* **13** 974–995. [MR2109061](#)
- SÜVEGES, M. (2007). Likelihood estimation of the extremal index. *Extremes* **10** 41–55. [MR2407640](#)
- WHITE, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica* **50** 1–25. [MR0640163](#)
- WHITE, H. (1994). *Estimation, Inference and Specification Analysis*. Cambridge Univ. Press, Cambridge. [MR1292251](#)

ECOLE POLYTECHNIQUE FÉDÉRALE
DE LAUSANNE
EPFL-FSB-IMA-STAT
STATION 8, 1015 LAUSANNE
SWITZERLAND
E-MAIL: Maria.Suveges@epfl.ch
Anthony.Davison@epfl.ch