

DISCUSSION OF THE DANISH DATA ON LARGE FIRE INSURANCE LOSSES

SIDNEY I. RESNICK
Cornell University

ABSTRACT

Alexander McNeil's (1996) study of the Danish data on large fire insurance losses provides an excellent example of the use of extreme value theory in an important application context. We point out how several alternate statistical techniques and plotting devices can buttress McNeil's conclusions and provide flexible tools for other studies.

KEYWORDS

Heavy tails, regular variation, Hill estimator, Poisson processes, linear programming, parameter estimation weak convergence, consistency, estimation, independence, auto-correlations.

1. INTRODUCTION

McNeil's (1996) interesting study of large fire insurance losses provides an excellent case history illustrating a variety of extreme value techniques. The goal of my remarks is to show additional techniques and plotting strategies which can be employed for similar data.

Our remarks concentrate on the following:

- Diagnostics for assessing the appropriateness of heavy tailed models.
- Diagnostics for testing for independence.

It is customary in many insurance studies involving heavy tailed phenomena to assume independence without actually statistically checking this important fact so some attention is given to this issue.

2. APPROPRIATENESS OF HEAVY TAILED MODELS

Given a particular data set, there are various methods of checking that a heavy tailed model is appropriate. The methods given below (these are also reviewed in Resnick 1995, 1996; Feigin and Resnick, 1996) supplement the techniques discussed by McNeil such as mean excess plots and QQ-plots against exponential quantiles. Unlike the mean excess plot, the following methods do not depend on existence of a finite mean for the marginal distribution of the stationary time series. This is important since it is becoming clear that it is not difficult to find examples of heavy tailed data which

require infinite mean models for adequate fits. (See for example the teletraffic examples in Resnick (1995, 1996)).

For the discussion that follows, we suppose $\{X_n, n \geq 1\}$ is a stationary sequence and that

$$P[X_1 > x] = x^{-\alpha} L(x), \quad x \rightarrow \infty \tag{2.1}$$

where L is slowly varying and $\alpha > 0$. Consider the following techniques:

(1) *The Hill plot.* Let

$$X_{(1)} > X_{(2)} > \dots > X_{(n)}$$

be the order statistics of the sample X_1, \dots, X_n . We pick $k < n$ and define the Hill estimator (Hill, 1975) to be

$$H_{k,n} = \frac{1}{k} \sum_{i=1}^k \log \frac{X_{(i)}}{X_{(k+1)}}.$$

Note k is the number of upper order statistics used in the estimation. The Hill plot is the plot of

$$((k, H_{k,n}^{-1}), 1 \leq k < n)$$

and if the $\{X_n\}$ process is iid or a linear moving average or satisfies certain mixing conditions then since $H_{k,n} \xrightarrow{P} \alpha^{-1}$ as $n \rightarrow \infty, k/n \rightarrow 0$ the Hill plot should have a stable regime sitting at height roughly α . See Mason (1982), Hsing (1991), Resnick and Stărică (1995, 1996a), Rootzen et. al (1990), Rootzen (1996). In the iid case, under a second order regular variation condition, $H_{k,n}$ is asymptotically normal with asymptotic variance $1/\alpha^2$. (See de Haan and Resnick, 1996).

(2) *The smooHill Plot.* The Hill Plot often exhibits extreme volatility which makes finding a stable regime in the plot more guesswork than science and to counteract this, Resnick and Stărică (1996a) developed a smoothing technique yielding the smooHill plot: Pick an integer u (usually 2 or 3) and define

$$smooH_{k,n} = \frac{1}{(u-1)k} \sum_{j=k+1}^{uk} H_{j,n}.$$

In the iid case when a second order regular variation condition holds, the asymptotic variance of $smooH_{k,n}$ is less than that of the Hill estimator, namely:

$$\frac{1}{\alpha^2} \frac{2}{u} \left(1 - \frac{\log u}{u}\right).$$

The sensitivity of the Hill estimate to the choice of k corresponds in McNeil's work to the sensitivity of the fit of the generalized Pareto to the data to the choice of threshold. Perhaps some comparable smoothing technique would help in GPD fitting.

(3) *Alt plotting; Changing the scale.* As an alternative to the Hill plot, it is sometimes useful to display the information provided by the Hill or smooHill estimation as

$$\{(\theta, H_{\lfloor n^\theta \rfloor, n}^{-1}), 0 \leq \theta \leq 1\}$$

and similarly for the smooHill plot where we write $\lceil y \rceil$ for the smallest integer greater or equal to $y \geq 0$. We call such plots the *alternative Hill plot* abbreviated AltHill and the *alternative smoothed Hill plot* abbreviated AltsmooHill. The alternative display is sometimes revealing since the initial order statistics get shown more clearly and cover a bigger portion of the displayed space. However, when the data is Pareto or nearly Pareto, this alternate plotting device is less useful since in the Pareto case, the Hill estimator applied to the full data set is the maximum likelihood estimator and hence the correct answer is usually found at the right end of the Hill plot.

(4) *Dynamic and static QQ-plots.* As we did for the Hill plots, pick k upper order statistics

$$X_{(1)} > X_{(2)} > \dots > X_{(k)}$$

and neglect the rest. Plot

$$\left\{ \left(-\log\left(1 - \frac{j}{k+1}\right), \log X_{(j)} \right), 1 \leq j \leq k \right\}. \quad (2.2)$$

If the data are approximately Pareto or even if the marginal tail is only regularly varying, this should be approximately a straight line with slope $1/\alpha$. The slope of the least squares line through the points is an estimator called the QQ-estimator (Kratz and Resnick, 1996). Computing the slope we find that the QQ-estimator is given by

$$\widehat{\alpha}^{-1}_{k,n} = \frac{\frac{1}{k} \sum_{i=1}^k \left(\log\left(\frac{i}{k+1}\right) \right) \log\left(\frac{X_{(i)}}{X_{(k+1)}}\right) - \frac{1}{k} \sum_{i=1}^k \left(-\log\left(\frac{i}{k+1}\right) \right) H_{k,n}}{\frac{1}{k} \sum_{i=1}^k \left(-\log\left(\frac{i}{k+1}\right) \right)^2 - \left(\frac{1}{k} \sum_{i=1}^k \left(-\log\left(\frac{i}{k+1}\right) \right) \right)^2}. \quad (2.3)$$

There are two different plots one can make based on the QQ-estimator. There is the dynamic QQ-plot obtained from plotting $\{k, 1/\widehat{\alpha}^{-1}_{k,n}, 1 \leq k \leq n\}$ which is similar to the Hill plot. Another plot, the static QQ-plot, is obtained by choosing and fixing k , plotting the points in (3.2) and putting the least squares line through the points while computing the slope as the estimate of α^{-1} .

The QQ-estimator is consistent for the iid model if $k \rightarrow \infty$ and $k/n \rightarrow 0$ and under a second order regular variation condition and further restriction on $k(n)$, it is asymptotically normal with asymptotic variance $2/\alpha^2$. This is larger than the asymptotic variance of the Hill estimator but the volatility of the QQ-plot always seems to be less than that of the Hill estimator.

(5) *De Haan's moment estimator.* McNeil discusses the extreme value distributions (see also Resnick, 1987; de Haan, 1970; Leadbetter et al, 1983; Castillo, 1988; Embrechts et al 1997) which can be parameterized as a one parameter family

$$G_{\xi}(x) = \exp\{-(1 + \xi x)^{-\xi^{-1}}\}, \quad \xi \in \mathfrak{R}, 1 + \xi x > 0.$$

When $\xi = 0$, we interpret G_0 as the Gumbel distribution

$$G_0(x) = \exp\{-e^{-x}\}, \quad x \in \mathfrak{R}.$$

A distribution whose sample maxima when properly centered and scaled converges in distribution to G_ξ is said to be in the *domain of attraction* of G_ξ which in McNeil's notation is written $F \in MDA(G_\xi)$. If $\xi > 0$ and $F \in MDA(G_\xi)$ then $1 - F \in RV_{-1/\xi}$. De Haan's moment estimator $\hat{\xi}_{k,n}$ (Dekker's, Einmahl, de Haan, 1989; de Haan, 1991; Dekkers and de Haan, 1991; Resnick and Starica, 1996b) is designed to estimate $\xi = 1/\alpha$. Note that $\hat{\xi}_{k,n}$, like the Hill estimator, is based on the k -largest order statistics. Since most common densities such as the exponential, normal, gamma and Weibull densities and many others are in the $MDA(G_0)$, the domain of attraction of the Gumbel distribution, this provides another method of deciding when a distribution is heavy tailed or not. If $\hat{\xi}_{k,n}$ is negative or very close to zero, there is considerable doubt that heavy tailed analysis should be applied and the moment estimator is usually much more reliable in these circumstances than the Hill estimator. In particular, when $\xi = 0$, the Hill estimator is not usually informative and the moment estimator does a much better job of identifying exponentially bounded tails. Smoothed versions of the moment estimator can also be devised (Resnick and Starica, 1996b) which overcome volatility in the plot of $\{k, \hat{\xi}_{k,n}, 1 \leq k \leq n\}$.




FIGURE 2.1: Tsplo and QQ plot of Danish data.




FIGURE 2.2: QQ plot of Danish.all data and parameter estimate.

Hill and Dynamic QQ




FIGURE 2.3: Hill and QQ-plot of Danish data.

Figure 2.1 gives a time series plot of the 2156 Danish data consisting of losses over one million Danish Krone (DKK) and the right hand plot is the QQ-plot (2.2) of this data yielding a remarkably straight plot. Figure 2.2 gives the QQ-plot of all of the 2492 losses recorded in the data set labeled *danish.all* and shows why McNeil was statistically wise to drop losses below one million DKK. (In the left hand plot the data is scaled to have a range of (0.3134041, 263.2503660) and the dots below height 0 represent the 325 values which are less than 1 in the scaled data.) The right hand plot in Figure 2.2 puts a line through the QQ-plot of the losses above one million and yields an estimate of $\alpha = 1.386$. Using only the largest 1500 order statistics and then estimating α from the slope of the LS line produces an estimate of $\alpha = 1.4$.

We next attempted to estimate α by means of the Hill plot. Figure 2.3 shows a Hill plot side by side with the dynamic QQ-plot. Because the plot in the right side of Figure 2.1 is so straight, we tend to trust the Hill plot near the right end of the plot. This is because the straight plot in Figure 2.1 indicates the underlying distribution is close to Pareto and for the Pareto distribution the maximum likelihood estimator of the shape parameter is the Hill estimator calculated using all the data. This analysis is confirmed by the excellent fit achieved by McNeil using a GPD with $\xi = 0.684$ or $\alpha = 1.46$ corresponding to losses exceeding a threshold of 20 million DKK. Such a GPD is a shifted Pareto.

On the other hand, examining the *altHill* and *altsmooHill* plots in Figure 2.4 makes it seem unlikely that α could be as large as 2.01 which is what is given in McNeil's Figure 7. This corresponds to a $\xi = 0.497$. Our methods indicate a likely value of $\alpha = 1.45$.

In Figure 2.5 we present four views of the moment estimator $\hat{\xi}_{k,n}$ of $\xi = 1/\alpha$. The upper right graph and the lower two graphs are in *alt* scale where k , $1 \leq k \leq n$ is replaced by $\lceil n^\theta \rceil$, $0 \leq \theta \leq 1$. Interestingly, we see here and in the four views of the Hill plot, that when the data are very close to Pareto, the *alt* scale is not advantageous.

When the data is close to Pareto, the reliable part of the graph is toward the end and this is the part of the graph under emphasized by the alt scale. The situation is very different for something like stable data (Resnick, 1995) where the traditional Hill plot is incapable of identifying the correct value of α but the alt plot does a superior job.

Based on an amalgam of the QQ, Hill and moment plots, we settle on an estimate of $\alpha = 1.4$ or $\xi = .71$.




FIGURE 2.4 Hill and smooHill plots for Danish data

3. TESTING FOR INDEPENDENCE

We outline several tests for independence which can help reassure the analyst that an iid model is adequate and that it is not necessary to try to fit a stationary time series with dependencies to the data. Some of our tests are motivated by our experience trying to fit autoregressive processes to heavy tailed data.

Here is a survey of several methods which can be used to test independence. Some of these are based on asymptotic methods using heavy tailed analysis and the rest are standard time series tests of homogeneity.

(1) *Method based on sample acf.* An exploratory, informal method for testing for independence can be based on the sample autocorrelation function $\hat{\rho}(h)$ where for h any positive integer

$$\hat{\rho}(h) = \frac{\sum_{t=1}^{n-h} (X_t - \bar{X})(X_{t+h} - \bar{X})}{\sum_{t=1}^n (X_t - \bar{X})^2}.$$

In many studies of heavy tailed data, the centering by the sample mean is omitted since if mathematical expectation does not exist, there is no advantage or sense to centering by the sample mean. However, since our chosen value of $\alpha = 1.4$ implies $E|X_1| < \infty$, we have decided to include the centering. From Davis and Resnick (1985a), if $\{X_i\}$ are iid with regularly varying tail probabilities, then

$$\lim_{n \rightarrow \infty} \hat{\rho}(h) = \begin{cases} 1, & \text{if } h = 0, \\ 0, & \text{if } h \neq 0. \end{cases}$$

Thus, if upon graphing $\hat{\rho}(h)$, $h = 0, \dots, n - h$ we get only small values for $h \neq 0$ there is no evidence against independence. The limit distribution of $\hat{\rho}(h)$, $h = 1, \dots, q$ is known (Davis and Resnick, 1985b, 1986 Corollary 1) but it is somewhat difficult to work with and the percentiles must be calculated by simulation. It is important to realize that the 95% confidence bands drawn by a typical statistics package like Splus are drawn using Bartlett's formula (Brockwell and Davis, 1991) on the assumption that the data is Gaussian or at least has finite fourth moment. This assumption is totally inappropriate for heavy tailed data and the confidence band must be drawn taking into account the heavy tailed limit distribution for $\hat{\rho}(h)$, $h = 1, \dots, l$.




FIGURE 2.5: Moment estimator plots for Danish data.

We discuss implementation of the acf based procedure when $1 < \alpha < 2$ since in the case of the Danish loss data we have settled on an estimate of $\alpha = 1.4$. Suppose $\{Y_1, \dots, Y_n\}$ are iid non-negative random variables satisfying

$$P[Y_1 > x] \sim x^{-\alpha} L(x), \quad x \rightarrow \infty, 1 < \alpha < 2$$

where L is slowly varying. From Corollary 1, page 553 of Davis and Resnick (1986), if we set $\hat{\rho}_Y(h)$ to be the lag h sample acf for Y_1, \dots, Y_n , then we have

$$\lim_{n \rightarrow \infty} P[\hat{b}_n^{-1} b_n^2 \hat{\rho}_Y(h) \leq x] = P[U_h / V_0 \leq x]$$

where U_h is a one sided stable random variable with index $\alpha = 1.4$ and V_0 is a positive stable random variable with index $\alpha/2 = 0.7$ and b_n is the solution to

$$P[Y_1 > x] = 1/n$$

and \hat{b}_n is the solution to

$$P[Y_1 Y_2 > x] = 1/n.$$

Thus an approximate symmetric 95% confidence window for the sample correlations of the Y 's would be placed at $\pm l \hat{b}_n / b_n^2$ where l satisfies

$$P[|U_h / V_0| \leq l] = .95.$$

We estimate the 95%-quantile of $|U_h/V_0|$ by simulation and if we assume the distribution of Y_i 's is Pareto from some point on, we find

$$l \frac{\hat{b}_n}{b_n^2} = l \left(\frac{n}{\log n} \right)^{-1/\alpha}.$$

The assumption of a Pareto distribution seems mild in view of Figure 2.2 and the good fit found by McNeil of the GPD with positive shape parameter.

Figure 3.1 presents this technique applied to the Danish loss data. No spike is protruding from the band and hence this acf based technique does not provide any evidence against the assumption of independence.




FIGURE 3.1 95% confidence band for the acf of the Danish loss data

(2) *Tests based on asymptotic theory.* Estimators of autoregressive coefficients for heavy tailed time series can be used to fashion tests for independence against autoregressive alternatives. If the autoregression is described as

$$X_t = \sum_{i=1}^p \phi_i X_{t-i} + Z_t, \quad t = 0, 1, \dots$$

where $\{Z_t\}$ are iid heavy tailed residuals, then we test if

$$\phi_1 = \dots = \phi_p = 0,$$

that is independence, by rejecting when the maximal estimated coefficient

$$\bigvee_{i=1}^p |\hat{\phi}_i(n)|$$

is too large. This procedure has been implemented by Feigin, Resnick and Stărica (1996) based on *linear programming* (LP) estimators under the assumption that the iid heavy tailed residuals $\{Z_t\}$ are non-negative. See also Feigin and Resnick (1993).

It would not be possible to fix the size of the LP test if the limit distribution of the LP estimator did not considerably simplify. Fortunately it does under the null hypothesis of independence and we then have

$$b_n(\hat{\phi}_1(n), \dots, \hat{\phi}_p(n)) \Rightarrow \mathbf{L} \equiv (V_1^{-1}, \dots, V_p^{-1})$$

where for $x_i \geq 0; i = 1, \dots, p$ we have that

$$P[V_i \leq x_i, i = 1, \dots, p] = \exp\left\{-\int_{(y_1, \dots, y_p) \in [0, \infty)^p} \left(\bigwedge_{l=1}^p y_l x_l\right)^{-\alpha} F(dy_1) \dots F(dy_p)\right\}. \quad (3.2)$$

This means that if we want a 0.05 level rejection region, we should reject when $\bigvee_{i=1}^p |\hat{\phi}_i(n)| > K(.05)$ where $K(.05)$ is defined by

$$P\left[\bigvee_{i=1}^p |\hat{\phi}_i(n)| > K(.05)\right] = .05$$

and to find an approximate value of $K(.05)$ we write

$$P\left[\bigvee_{i=1}^p |\hat{\phi}_i(n)| > K(.05)\right] \approx P\left[\bigvee_{i=1}^p L_i > b_n K(.05)\right] \leq pP[L_1 > b_n K(.05)] = pe^{-c(b_n K(.05))^\alpha}, \quad (3.3)$$

where $c = E(Z_1^{-\alpha})$. This yields

$$K(.05) \approx \frac{\left(\frac{-\log(.05/p)}{c}\right)^{1/\alpha}}{b_n} = \frac{\left(\frac{\log(20p)}{c}\right)^{1/\alpha}}{b_n}.$$

We need to estimate α , c and b_n . One way to do this is to use the QQ-plot (Feigin, Resnick and Stărică, 1996; Kratz and Resnick, 1996) which yields both \hat{b}_n (as the

intercept of the fitted line) and $\hat{\alpha}$ (as the reciprocal of the slope of the fitted line) and then we can get

$$\hat{c} = n^{-1} \sum_{i=1}^n X_i^{-\hat{\alpha}}.$$

The asymptotic test is implemented and shown in Figure 3.2. None of the estimated coefficient values extend above the bar representing $K(.05)$ so this method provides no evidence against the hypothesis of independence.




FIGURE 3.2 Asymptotic test for independence for the Danish loss data

(3) *Standard tests of randomness.* There are several standard time series tests of randomness (Brockwell and Davis, 1991, Section 9.4) which are non-parametric and can be employed in the present context. We give some examples below. We use the notation

$$\chi_n \sim AN(\mu_n, \sigma_n^2)$$

as shorthand to mean that

$$(\chi_n - \mu_n) / \sigma_n \Rightarrow N(0, 1).$$

- (1) Turning point test. If T is the number of turning points among X_1, \dots, X_n then under the null hypothesis that the random variables are iid we have

$$T \sim AN(2(n - 2) / 3, (16n - 29) / 90)$$

and this can be used as the basis of a test.

- (2) Difference-sign test. Let S be the number of $i = 2, \dots, n$ such that $X_i - X_{i-1}$ is positive. Under the null hypothesis that the random variables X_1, \dots, X_n are iid we have

$$S \sim AN\left(\frac{1}{2}(n - 1), (n + 1) / 12\right).$$

- (3) Rank test. Let P be the number of pairs (i, j) such that $X_j > X_i$ for $j > i$ and $i = 1, \dots, n - 1$. Under the null hypothesis that the random variables X_1, \dots, X_n are iid we have

$$P \sim AN\left(\frac{1}{4}n(n-1), n(n-1)(2n+5)/8\right).$$

We would reject the iid hypothesis at the 0.05 level if any of these standardized variables had an absolute value greater than 1.96. All of these tests are implemented in the Brockwell and Davis (1991) package ITSM. Data can easily be imported into their program and tested within the package for randomness.

We carried out these tests on the Danish loss data using ITSM and achieved the following results:

Turning points	1409	AN (1436.00, 19.57 ²)
Difference-sign	1079	AN (1077.50, 13.41 ²)
Rank test	1055894	AN (1161545, 50071.90 ²).

The rank test rejects the hypothesis of independence at the 5% level. The turning points and difference-sign tests fail to reject.

(4) *Stability testing on subsets of the data.* An informal but useful technique is to take a statistic, such as the sample acf, and compute it relative to different subsets of the sample. If the data is iid, the values of the statistic should be similar across different subsets.

For the sample acf, if the graphs of $\hat{\rho}_H(h)$, $h = 1, \dots, q$ look different for different subsets, then one should be skeptical of the correctness of the iid assumption. Often it is enough to split the sample into halves or thirds to generate some skepticism. One could make acf subset plots for the Danish data but since the acf values are not significantly different from 0, there seems little point to pursuing this diagnostic in this case.

(5) *Permutation test for independence.* Another approach to testing for independence in time series analysis is based on permutation tests. Here we can use any desired statistic that is designed to measure some form of dependence between successive data. This statistic might be a maximum autocorrelation or partial autocorrelation, or it may be a maximal autoregressive coefficient estimated by the linear programming paradigm.

The permutation test is based on comparing the observed value of the statistic with the permutation distribution of that statistic — that is with the distribution of values of the statistic under all the possible permutations of the time series data. If there is no dependence structure in the data, then the observed value should be a typical value for this reference permutation distribution. If there is some dependence of the type to which the statistic is sensitive, then the observed value should be extreme with respect to this reference distribution.

This approach allows one to perform tests without relying on the asymptotic theory for the particular statistic. As we have seen earlier, the asymptotic distribution for

$$\bigvee_{i=1}^p |\hat{\phi}_i(n)|$$

involves various parameters that have to be estimated. Moreover, the fact that we are not sure of the rate of convergence to the asymptotic distribution, also suggests the precautionary tactic of using a permutation test.

In the implementation we use below, we approximate the *p-value* of the actually observed statistic. This is achieved by generating 99 permutations of the time series, computing the statistic for each one, and counting the number (*C*) of these that are greater than or equal to the actually observed statistic. The *p-value* is approximated by $(1+C)\%$. The statistics considered are the maximum absolute autocorrelation (*macf*), the maximum absolute partial autocorrelation (*mpacf*), and the maximum absolute linear programming coefficient estimate (*mphi*). In each case, one must specify the value of *p*, the order over which the maximum is taken.

For the Danish loss data, we took the order to be 10 and ran the tests yielding the following *p-values*:

maximum autocorrelation	0.52
maximum partial autocorrelation	0.51
maximum LP coefficient	0.22

and thus at a reasonable level, none of these tests would reject independence.

4. CONCLUDING REMARKS

There is very little evidence arguing against the hypothesis of independence and it seems McNeil's presumption that the data were independent was a safe assumption to make for this data set. Independence is not that common among teletraffic of finance data in my experience and thus should be treasured in the present insurance context. Fitting dependent data with a heavy tailed stationary time series model can be a frustrating business (see Resnick, 1996b; Feigin and Resnick, 1996) so when one concludes the data can be modelled as iid, a loud sigh of relief is heard.

The sensitivity of the estimation and fitting methods to the choice of threshold or the choice of the number of order statistics used in estimation is a persistent and troubling theme in McNeil's and my remarks. This seems inherent in the heavy tail and extreme value methods. It is not clear at this point how much the techniques can be improved to reduce sensitivity to choice of *k* or threshold. Smoothing techniques and alternate plotting help but are not a universal panacea.

It is encouraging to see the accumulating mass of theoretical and software tools which can be used to analyze such data sets.

REFERENCES

- BROCKWELL, P. and DAVIS, R., *Time Series: Theory and Methods*, 2nd edition, Springer-Verlag, New York, 1991.
- BROCKWELL, P. and DAVIS, R. *ITSM: An Interactive Time Series Modelling Package for the PC*, Springer-Verlag, New York, 1991.
- CASTILLO, ENRIQUE, *Extreme Value Theory in Engineering*, Academic Press, San Diego, California, 1988.

- DAVIS, R. and RESNICK, S., Limit theory for moving averages of random variables with regularly varying tail probabilities., *Ann. Probability* **13** (1985a), 179-195.
- DAVIS, R. and RESNICK, S., *More limit theory for the sample correlation function of moving averages*, *Stochastic Processes and their Applications*, **20** (1985b), 257-279.
- DAVIS, R. and RESNICK, S., Limit theory for the sample covariance and correlation functions of moving averages, *Ann. Statist.* **14** (1989), 533-558.
- DEKKERS, A., EINMAHL, J., and HAAN, L. DE, *A moment estimator for the index of an extreme value distribution*, *Ann. Statist.* **17** (1989), 1833-1855.
- DEKKERS, A. and HAAN, L. DE, On the estimation of the extreme value index and large quantile estimation, *Ann. Statist.* **17** (1989), 1795-1832.
- EMBRECHTS, P., KLUPPELBERG, C. and MIKOSCH, T., *Modelling Extremal Events for Insurance and Finance*, To appear, Springer-Verlag, Heidelberg, 1997.
- FEIGIN, P. and RESNICK, S., *Limit distributions for linear programming time series estimators*, *Stochastic Processes and their Applications* **51** (1994), 135-166.
- FEIGIN, P. and RESNICK, S., *Pitfalls of fitting autoregressive models for heavy-tailed time series*, Available at <http://www.orie.cornell.edu/trlist/trlist.html> as TR1163.ps.Z (1996).
- FEIGIN, P., RESNICK, S., and Stărică, C., *Testing for independence in heavy tailed and positive innovation time series*, *Stochastic Models* **11** (1995), 587-612.
- HAAN, L. DE, *On Regular Variation and its Application to the Weak Convergence of Sample Extremes*, *Mathematical Centre Tract 32*, Mathematical Centre, Amsterdam, Holland, 1970.
- HAAN, L. DE, *Extreme Value Statistics*, Lecture Notes, Econometric Institute, Erasmus University, Rotterdam (1991).
- HAAN, L. DE and RESNICK, S., *On asymptotic normality of the Hill estimator*, TR1155.ps.Z available at <http://www.orie.cornell.edu/trlist/trlist.html> (1996).
- HILL, B., A simple approach to inference about the tail of a distribution, *Ann. Statist.* **3** (1975), 1163-1174.
- HSING, T., Extreme value theory for suprema of random variables with regularly varying tail probabilities., *Stoch. Proc. and their Appl.* **22** (1986), 51-57.
- KRATZ, M. and RESNICK, S., *The qq-estimator and heavy tails*, *Stochastic Models* **12** (1996), 699-724.
- LEADBETTER, M., LINDGREN, G. and ROOTZEN, H., *Extremes and Related Properties of Random Sequences and Processes*, Springer Verlag, New York, 1983.
- MASON, D., Laws of large numbers for sums of extreme values, *Ann. Probability* **10** (1982), 754-764.
- MCNEIL, A., Estimating the tails of loss severity distributions using extreme value theory, Preprint: Dept. Mathematics, ETH Zentrum, CH-8092 Zürich (1966).
- RESNICK, S., *Extreme Values, Regular Variation, and Point Processes*, Springer-Verlag, New York, 1987.
- RESNICK, S., *Heavy tail modelling and teletraffic data*, Available as TR1134.ps.Z at <http://www.orie.cornell.edu/trlist/trlist.html>, *Ann. Statist.* (1995), (to appear).
- RESNICK, S., *Why non-linearities can ruin the heavy tailed modeler's day*, Available as TR1157.ps.Z at <http://www.orie.cornell.edu/trlist/trlist.html>, A PRACTICAL GUIDE TO HEAVY TAILS: Statistical Techniques for Analysing Heavy Tailed Distributions (Robert Adler, Raisa Feldman, Murad S. Taqqu, ed.), Birkhauser, Boston 1996, (to appear).
- RESNICK, S. and Stărică, C., *Consistency of Hill's estimator for dependent data*, *J. Applied Probability* **32** (1995), 139-167.
- RESNICK, S. and Stărică, C., *Smoothing the Hill estimator*, To appear: *J. Applied Probability* (1996a).
- RESNICK, S. and Stărică, C., *Smoothing the moment estimator of the extreme value parameter*, Available as TR1158.ps.Z at <http://www.orie.cornell.edu/trlist/trlist.html>, Preprint (1996b).
- ROOTZEN, H., LEADBETTER, M. and DE HAAN, L., *Tail and quantile estimation for strongly mixing stationary sequences*, Technical Report 292, Center for Stochastic Processes, Department of Statistics, University of North Carolina, Chapel Hill, NC 27599-3260 (1990).
- ROOTZEN, H., *The tail empirical process for stationary sequences*, Preprint 1995:9 ISSN 1100-2255, *Studies in Statistical Quality Control and Reliability*, Chalmers University of Technology (1995).

SIDNEY I. RESNICK
 Cornell University,
 School of Operations Research and Industrial Engineering
 Rhodes Hall 223
 Ithaca, NY 14853 USA
 E-mail: sid@orie.cornell.edu