Exact inference and optimal invariant estimation for the stability parameter of symmetric α-stable distributions*

JEAN-MARIE DUFOUR [†] JEONG-RYEOL KURZ-KIM[‡]

July 12, 2007

Abstract

Hill estimation (Hill, 1975), the most widespread method for estimating tail thickness of heavy-tailed financial data suffers from two drawbacks: one of them is that the *unknown* number of observations in the tail is involved in the estimation, which diminishes the empirical relevance of the Hill estimation. The other is that the hypothesis test for finite samples whether the underlying data lie in the domain of attraction of an α-stable law ($\alpha < 2$) or of a normal law ($\alpha \geq 2$) is performed on the basis of the asymptotic distribution, which can be different from those for finite samples. In this paper, using the Monte Carlo technique, we propose an exact test method for the stability parameter of α-stable distributions which is able to provide exact confidence intervals for finite samples. Our exact test method includes automatically an estimation procedure which does not need the assumption of a known number of observations on the distributional tail. Empirical applications demonstrate the advantages of the MC method in comparison with the Hill estimation.

JEL Classification: C1; C15; C22; G0; E0

*The views expressed in this paper are those of the authors and not necessarily those of the Deutsche Bundesbank. Research support from the Alexander-von-Humboldt Foundation is gratefully acknowledged by the two authors.

[†]Corresponding author, C.R.D.E. and Département de Sciences Economique, Université de Montréal, Montréal, Québec H3C 3J7, Canada, E-mail: jean.marie.dufour@umontreal.ca

[‡]Deutsche Bundesbank, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main, Germany, E-mail: jeong-ryeol.kurz-kim@bundesbank.de

Keywords: α-stable distribution; Hill estimator; Monte Carlo test; Hodges-Lehmann estimator; exact confidence interval; finite sample.

1 Introduction

Since the influential work by Mandelbrot (1963), α-stable distributions have often been considered a more realistic distribution for high-frequency variables, such as financial data, than the normal distribution, because asset returns, for example, are typically heavy-tailed and excessively peaked around zero—phenomena that can be captured by α-stable distributions with $\alpha < 2$.

Statistical inferences for estimations and hypothesis tests under the α-stable distributional assumption depend crucially on α (DuMouchel, 1971).[1] Therefore, one of the most important tasks in using the α-stable distribution is to precisely estimate α and to find exact confidence intervals for finite samples for estimated α. To discriminate $\alpha < 2$ (domain of attraction of an α-stable law, Paretian case) from $\alpha \geq 2$ (domain of attraction of normal law), or rather $\alpha = 2$ (Gaussian case), would be one example of how important it is to know what the exact confidence interval for the estimated α is.

Because of its well-developed asymptotic properties, the Hill estimator (Hill, 1975) is the most popular method for estimating tail thickness of empirical data, the stability parameter in the context of α-stable distributions. It is a simple nonparametric estimator based on order statistics. A severe drawback of the Hill estimator, however, is that the number of the observations on the distributional tails must be known. In practice, the number of observations on the distributional tail is generally unknown and depends on an unknown α. One more drawback is that its confidence interval for finite samples can be given only based on the asymptotic distribution, which generally differ from the finite sample distributions. The Pickands (Pickands, 1975) and Dekkers, Einmahl and de Haan estimator (Dekkers et al., 1989) are variations on the Hill estimator. For a rough check, the quantile estimation of McCulloch (1986) may be also used. Some modifications are also considered by some authors: Huisman et al. (2001), for example, propose a weighted Hill estimator that takes into account the trade-off between bias and variance of the Hill estimator. However, for all the modified estimators of the Hill estimator, say Hill-type estimators, the assumption of a known number of observations on the distributional tail is necessary, and confidence intervals for finite samples can be given only based on the asymptotic

[1] Kurz-Kim and Loretan (2007), for example, revisit the CRSP data used in Fama and French (1992) and show that the empirical conclusion about the Capital Asset Pricing Model driven by Fama and French (1992) is not robust depending on the distributional assumption for the underlying data.

distribution.

In this paper, using the Monte Carlo (MC) technique, we propose an exact test method which automatically includes an estimation procedure[2] (henceforth referred to as the MC test or the MC estimation) for the stability parameter of α-stable distributions. This is because an exact confidence interval for finite samples can be constructed in the estimation procedures, or rather an estimate in the test procedure. Our MC method therefore improves on the Hill estimation in two ways: first, the number of observations on the distributional tail does not need to be assumed to be known for our estimator. Second, our estimator provides exact confidence intervals for finite samples.

The rest of the paper is structured as follows. Section 2 gives a brief summary of α-stable distributions and the Hill estimator. In Section 3, the MC estimation and test procedure are explained. In Section 4, we perform simulations to study the size of the usual asymptotic test and power of the MC test for finite samples. An empirical application is given in Section 5. Section 6 summarizes the paper.

2 Framework

2.1 A brief summary of α-stable distributions

A random variable ($r.v.$) X is said to be stable if, for any positive numbers A and B, there is a positive number C and a real number D such that $AX_1 + BX_2 \stackrel{d}{=} CX + D$, where X_1 and X_2 are independent $r.v.s$ with $X_i \stackrel{d}{=} X, i = 1, 2$; and " $\stackrel{d}{=}$ " denotes equality in distribution. Moreover, $C = (A^\alpha + B^\alpha)^{1/\alpha}$ for some $\alpha \in (0, 2]$, where the exponent α is called a stability parameter. A stable $r.v.$, X, with a stability parameter α is called α-stable. The α-stable distributions are described by four parameters denoted by $S(\alpha, \beta, \mu, \sigma)$. Although the α-stable laws are absolutely continuous, their densities can be expressed only by a complicated special function except in three special cases.[3] Therefore, the logarithm of the characteristic function of the α-stable distribution is the best way of characterizing all members of this

[2] An estimation procedure which is based on the MC method is termed a Hodges-Lehmann estimation in the literature. See Hodges and Lehmann (1963) for the basic idea.

[3] The three special cases, in which the densities are expressible via elementary functions, are (i) the Gaussian distribution $S(2, 0, \mu, \sigma) \equiv N(\mu, 2\sigma^2)$, (ii) the symmetric Cauchy distribution $S(1, 0, \mu, \sigma)$, and (iii) the Lévy distribution $S(0.5, \pm 1, \mu, \sigma)$; see Zolotarev (1986).

family and is given as

$$\ln \int_{-\infty}^{\infty} e^{ist} d\mathrm{P}(S < s) = \begin{cases} -\sigma^{\alpha}|t|^{\alpha}[1 - i\beta \ \mathrm{sign}(t) \tan\frac{\pi\alpha}{2}] + i\mu t, & \text{for } \alpha \neq 1, \\ -\sigma|t|[1 + i\beta\frac{\pi}{2} \ \mathrm{sign}(t) \ln|t|] + i\mu t, & \text{for } \alpha = 1. \end{cases}$$

The shape of the α-stable distribution is determined by the stability parameter α. For $\alpha = 2$ the α-stable distribution reduces to the normal distribution, which is the only member of the α-stable family with finite variance. If $\alpha < 2$, moments of order α or higher do not exist and the tails of the distribution become thicker, i.e. the magnitude and frequency of outliers (from the viewpoint of the Gaussian) increase as α decreases. Skewness is governed by $\beta \in [-1, 1]$. If $\beta = 0$, the distribution is symmetric. The location and scale of the α-stable distributions are denoted by μ and σ. The standardized version of the α-stable distribution is given by $S((x - \mu)/\sigma; \alpha, \beta, 0, 1)$.

A strong argument in favor of the α-stable distribution as a distributional assumption for heavy-tailed empirical data is that only the α-stable distribution can serve as the limiting distribution of sums of independent identically distributed (i.i.d.) *r.v.s*, which is proved by Zolotarev (1986). For more details on the α-stable distributions, see Zolotarev (1986) and Samorodnitsky and Taqqu (1994); and for discussions of the role of the α-stable distribution in financial markets and macroeconomic modelling, see McCulloch (1996), Kim et al. (1997) and Rachev et al. (1999).

2.2 Hill estimator and choice of the number of observations on the distributional tail

2.2.1 Hill estimator

The most popular estimation for α is the Hill estimator (Hill, 1975), which is a simple nonparametric estimator based on order statistics. Because of its simplicity and popularity, we use the Hill estimator for constructing our test statistic.[4] Given a sample of n observations, $X_1, X_2, \ldots, X_n$, the Hill estimator is given as

$$\hat{\alpha}_H = \left[k^{-1} \sum_{j=1}^{k} (\ln X_{n+1-j:n} - \ln X_{n-k:n}) \right]^{-1}, \tag{1}$$

[4]Basically, any consistent estimator can be used to formulate our test statistic.

with standard error

$$SD(\hat{\alpha}_H) = \frac{k\hat{\alpha}_H}{(k-1)\sqrt{(k-2)}}, \tag{2}$$

where k is the number of observations which lie on the tails of distributions of interest and $X_{j:n}$ denotes the j-order statistic of the sample size n. As pointed above, the Hill estimation contains an unknown value of k. The k is to be chosen depending on the sample size, n, and the stability parameter, α, as $k = k(n, \alpha)$. In empirical works, however, the α is again unknown. The asymptotic properties of the Hill estimator have been studied by many authors and are now well developed: Mason (1982) and Hsing (1992) consider weak consistency of the Hill estimator for independent and dependent cases, respectively. The strong consistency is proved by Deheuvels et al. (1988). Goldie and Smith (1987) prove asymptotic normality of the Hill estimator, i.e.

$$\sqrt{k}(\hat{\alpha}_H^{-1} - \alpha^{-1}) \sim N(0, \hat{\alpha}^{-2}). \tag{3}$$

The confidence interval of an estimated stability parameter for a finite sample is based on the asymptotic distribution in (3) as is given in (2).

2.2.2 Choice of optimal k

Before we estimate the stability parameter using the Hill estimator, one practical problem needs to be solved: how to choose optimally the number of observations on the distributional tail, k, which are used for the Hill estimation. Note that the choice of k involves a trade-off, because it must be small enough for the observation, $X_{n-k:n}$ to be the first (smallest) observation in the tail of the distribution. If it is too small, however, the estimator will lack precision. The theoretical relationship between k and α can be derived from the second-order property as (see de Haan and Peng, 1998)

$$k = n\, \frac{\alpha - 1}{\Gamma(2-\alpha)\sin(\pi * \alpha/2)}, \quad \alpha \in (1,\ 2). \tag{4}$$

One problem with using this relationship for empirical work is that k is again a function of the unknown α. In practice, the true value of α is generally unknown, which means that the relationship between k and α is only of theoretical interest. In practice, the k is determined more or less by intuition or rather arbitrarily in the Hill estimation. This is a severe drawback for the Hill estimator in terms of empirical relevance and, as far as we know there is no statistical consensus to determine k. DuMouchel (1983) proposes a 10% fraction of samples as k independent of *alpha* and n, which, as will be shown, can be only optimal when α is very small and n is very large. Therefore, from the practical point of view, our MC estimation for which k does not need to be known is of empirical relevance.

In order to demonstrate the problem of the dependence of k on the unknown α, we show the so-called 'Hill horror plot'. Figure 1 shows Hill estimates for which the data come from an α-stable distribution with index $\alpha = 1.75$ and a sample size of 1,000.[5] The rule of choosing the k for the Hill estimation is to take the ten largest observations in the first estimation and to add the ten (1% of the 1,000 observations) next largest observations recursively until 990 (99% of the 1,000 observations) are used as the tail area. For every 99 cases, 100,000 replications are made.

Figure 1 somewhere here.

The solid line shows mean of the 100,000 estimates for each k and the dashed lines over and upper of the decreasing solid line shows 95% Monte Carlo confidence intervals. The straight solid line means the true value of α is 1.75.

Figure 1 shows clearly that the Hill estimate is extremely sensitive to the choice of k. On the other hand, Figure 1 shows also that estimates of tail index (stability parameter) greater than 2 in empirical works are, therefore, not evidence against infinite-variance stable distributions ($\alpha < 2$), as is pointed out in McCulloch (1997). A bad choice of k is often misleading about the true tail-thickness.

3 MC estimation and test procedure

In this section, we introduce the MC estimation and test procedure. The technique of the MC method was originally proposed by Dwass (1957) for implementing permutation tests and was later extended by Barnard (1963), and has recently been revisited by Dufour (2005). It provides an attractive method of building exact tests from statistics whose finite sample distribution is intractable but can be simulated. The most promising advantage of the MC method – unlike bootstrap techniques and other conventional test methods, which have only asymptotic justification – is that an exact finite-sample inference can be obtained. Consequently, the validity

[5]The results for other α values and sample sizes are the same as that for 1,000 with respect to the main conclusion.

of this MC-method-based exact randomized test does not depend on the number of replications made. For more details on the MC method, see Birnbaum (1974), and Dufour (2005).

We now test our random sample, $\{X_1, X_2, \ldots, X_n\}$, from a symmetric α-stable ($S_\alpha S$) distribution[6] for

$$H_0(\alpha_0) : \alpha = \alpha_0. \tag{5}$$

To perform this test, we need a test statistic which is free of nuisance parameters under the null hypothesis. A possible statistic can be given as

$$ST = \hat{\alpha} - \alpha_0, \tag{6}$$

where $\hat{\alpha}$ may be any consistent estimator for α. The fact that $\hat{\alpha}$ may be any consistent estimator for α means that our MC method can provide any consistent estimator with an exact confidence interval for finite samples.

Specifically, we apply the Hill estimator to our test statistic as

$$ST_H = \hat{\alpha}_H - \alpha_0. \tag{7}$$

The MC estimate can be given as

$$\hat{\alpha}_H(\alpha_0) = \left[k^{-1} \sum_{j=1}^{k(\alpha_0)} \left(\ln |\tilde{X}_{n+1-j:n}| - \ln |\tilde{X}_{n-k:n}| \right) \right]^{-1}, \tag{8}$$

with $\tilde{X}_i := X_i - X^{med}$, where $\tilde{X}_{j:n}$ stands for the j-th order statistic of the sample size n. Note that the distributional tail of the Hill estimator in (8) no longer depends on an unknown α as in (1), but on α_0 under the null hypothesis. This enables us to use the theoretical relationship between k and n as in (4) and/or the optimal k/n ratio tabulated in Rachev and Mittnik (2000, p. 114). In this sense, the MC estimation is optimal.

Two practical points in the test statistic above should be mentioned: the use of absolute values and the median centering. The use of absolute values after a median centering enables us to deal with asymmetric cases. This is because, as we will see in the following proposition, it avoids the dependence of the stability parameter α on the skewness parameter β. The other practical problem is how to choose a centering

[6] In case of asymmetric data, our procedure can be also applied in the same way, as in the Hill estimation, for only one tail. Based on the Kolmogorov-Smirnov statistic, Dufour et al. (2007) propose an estimation and test for the asymmetric parameter of α-stable distributions.

parameter in order to relocate the empirical data for the Hill estimator. Note that the Hill estimator is scale-invariant, but not location-invariant, which means X has to be centered properly at the beginning of the estimation. Despite the existence of the first moment for $1 < \alpha < 2$, the mean often cannot serve optimally as a centering parameter because of its fluctuation; especially α is small. Therefore, the median is an alternative choice as a centering parameter.

Regarding centering, we perform a simulation study. The simulation shows the efficiency of the Hill estimator among three centerings; true mean, sample mean and sample median. The case for true mean is not of empirical relevance, but it serves as a benchmark for the other two sample statistics. The simulation is designed as $\alpha = 1.0, 1.25, 1.5, 1.75, 1.95, 2$, $\beta = 0$, $\mu = 0$ and $\sigma = 1$ with sample size of $n = 100, 250, 500, 1,000, 5,000$. For each combination, 10,000 replications were made. For estimation we use the usual Hill estimator, where the sample is relocated by true mean, by estimated sample mean and by the estimated sample median. The pseudo-α-stable *r.v.s* were generated with the improved algorithm of Chambers et al. (1976) by Weron (1996).[7] The results of the simulations are summarized in Table A in the Appendix. Table 1 shows that using the median as a centering parameter is almost as efficient as using the true mean for all α and n adopted in the simulation. Furthermore, it is clearly shown that using the median as a centering parameter is more efficient than using the mean in the sense of mean square error for all α and n adopted in the simulation. The difference of the two root mean squares for the case median and mean is larger as α becomes smaller, which is expected because of the large fluctuation of sample means for small values of α, and the difference remains even for a large sample size (n = 5,000). For this reason, in the literature a trimmed mean as a centering is also recommended. But, the median centering seems to be mostly appropriate for our purpose.

To estimate the stability parameter using our MC method, the test statistic in (7) should be nuisance-free. Because the estimator in (8) is location and scale-invariant, the test statistic in (7) is pivotal as proved in the following lemma.

Proposition 1 *[Invariance] Let $X_1, X_2, \ldots, X_n$ be i.i.d. random variables which follow a $S(\alpha, \beta, \mu, \sigma)$ distribution, and let*

$$\hat{\alpha} = a(X_1, X_2, \ldots, X_n)$$

[7] The same random generator will be used for all the following simulations.

be an estimator of α. *If the estimator* $\hat{\alpha}$ *is scale-invariant, i.e.*

$$\hat{\alpha} = a(cX_1, \ldots, cX_n) = a(X_1, X_2, \ldots, X_n)\ , \ \textit{for all } c > 0\,, \tag{9}$$

then the estimator $\hat{\alpha}$ *has a distribution which depends only on* α, β *and* μ/σ. *If, furthermore,* $\hat{\alpha}$ *is location-scale-invariant, i.e.*

$$\hat{\alpha} = a(cX_1{+}d, \ldots, cX_n{+}d) = a(X_1, X_2, \ldots, X_n)\ , \ \textit{for all } c > 0 \textit{ and } d \in \mathbf{R}, \tag{10}$$

then the estimator $\hat{\alpha}$ *has a distribution which depends only on* β.

Proof 1 *To obtain the first result, we observe that*

$$X_i/\sigma \sim S(\alpha, \beta, \mu/\sigma, 1)\,,\ i = 1, \ldots, n\,.$$

Then, using the scale-invariance property (9) with $c = 1/\sigma$, *we can write*

$$\hat{\alpha} = a(X_1/\sigma, \ldots, X_n/\sigma)\ ,$$

from which we see that the distribution of $\hat{\alpha}$ *depends only on* α, β, *and* μ/σ. *Similarly, under the location-scale invariance condition (10), we observe the following:*

$$(X_i - \mu)/\sigma \sim S(\alpha, \beta, 0, 1)\,,\ i = 1, \ldots, n.$$

Hence, taking $c = 1/\sigma$ *and* $d = -\mu/\sigma$,

$$\hat{\alpha} = a(X_1^*, \ldots, X_n^*)\ ,$$

where $X_i^* = (X_i - \mu)/\sigma$, $i = 1, \ldots, n$.

The MC estimation and test procedure can be summarized in six steps as follows, given a random sample $\{X_1, X_2, \ldots, X_n\}$ from a $S_\alpha S$ distribution.

1 Determine the set of possible α under the null hypothesis. From the viewpoint of empirical relevance it stands to reason that $\alpha_0 \in [1\ 2]$.

2 Calculate test statistics $(\hat{\alpha}_H - \alpha_0)$ for every α_0, where the step length of two neighborhoods of α_0 may be 0.01, for example.

3 Generate typically 99 or 999 samples for every element of the set by a stable random variable generator and calculate the test statistics.

4 Compute p-values under all possible null hypotheses.

5 Take the $\hat{\alpha}_H(\alpha_0)$ as the estimate of α at which the $(1-p)$ value has its minimum (usually zero). This is a Hodges-Lehmann estimate.

6 Take the $\hat{\alpha}_l$ and $\hat{\alpha}_r$ as the left and right limit of the $\eta\%$ confidence interval at which the p-value is $(1-\eta)/100$, where $\hat{\alpha}_l < \hat{\alpha}_r$. This is now the exact confidence interval for the Hodges-Lehmann estimate in step 5.

4 A simulation study: exact and asymptotic confidence interval

4.1 Size distortion

A strong advantage of the MC method is that it provides exact confidence intervals for finite samples. In practice, the asymptotic normality as given in (3) is usually used for finite samples. In order to see size distortion of the asymptotic test, we perform a simulation study. The simulation is designed as $\alpha = 1.0, 1.25, 1.5, 1.75, 2$, $\beta = 0$, $\mu = 0$ and $\sigma = 1$ with a sample size of $n = 100, 250, 500, 1,000, 5,000$. For each combination, 10,000 replications were made. We use the usual Hill estimator relocating the sample by true mean, by estimated sample mean and by the estimated sample median. The result of the simulation is summarized in Table 1, where the numbers in the table are percentage points of rejection. Although we only report for the 95% significance level, the other usual confidence levels show very similar results.

Table 1 somewhere here.

The result of the asymptotic test is somehow surprising: a Hill estimation with a centering of true mean or median shows a (relatively) good size. However, a Hill estimation with a centering of sample mean shows size distortion if α is small. Note that the confidence intervals from the MC method are, by construction, exact.

4.2 Power function of the Monte Carlo method based test

The theoretical size and power of the MC test is considered in Dufour (2005). Although the discrepancy of the correct size and the superior power of the MC test over the conventional test goes to zero as the sample size approaches infinity, the behavior of the power function for the finite sample is usually of interest.

To check the power of our MC test, we perform a simulation study by drawing from symmetric α-stable pseudo-$r.v.s$ relocated by the median. As pseudo-empirical data we use the same α-stable random sample generated earlier, and test $H_0 : \alpha = \alpha_0$, where α_0 is assumed to take on values from 1.0 to 2.0 in steps of 0.1. Sample sizes of n = 100, 250, 500 1,000, 2,000, 5,000 and 10,000 are selected, and the number of replications is 10,000. To calculate the test statistic in (6) containing the Hill estimator, we use the ratios tabulated in Rachev and Mittnik (2000, p. 114) as the optimal k/n. To demonstrate the power function, we select a usual significance level of 95%. Figure 1 shows the power functions for the selected α, n and percentage points as described above.

Figure 2 somewhere here.

As expected, the power converges to the corresponding ideal value for each given significance level as the sample size grows. A sample size of 2,000 gives a rather satisfactory power. A large loss of power can be observed for extremely small sample sizes.

5 Empirical applications

To illustrate the use of the Monte Carlo method in practice, we employ the German stock index from its beginning (1 October 1959) to 30 September 2006, (47 years). For a deeper look, we consider them in three different frequencies, namely daily (11,796 observations), weekly (2,453 observations) and monthly (564 observations), where the observations for the weekly and monthly data are those of the end of the period, i.e., the Friday values for the weekly data and the value at end of the each month for the monthly data. Figure 3 shows the empirical data.

Figure 3 somewhere here.

The volatility cluster looks –to a large extent– similar in three different frequencies. But, a careful look reveals that many of single outliers in the daily returns are no longer seen in the weekly returns and *vice versa*. The same is also valid between the weekly returns and the monthly ones. This is because the weekly and/or monthly data do not come from a moving average of the daily data. (Even if the low frequency data come from a moving average of a higher frequency data, the two dynamics are not necessarily the same or very similar.)

Figure 4 shows the empirical densities of the three time series (solid line) compared with the normal density (dotted line).

Figure 4 somewhere here.

Each of the empirical densities appears excessively peaked around the mean and, at the same time, the tails are thicker than those of the normal density, which are the typical features of α-stable densities. This phenomenon is the most striking in the daily data, namely high-frequency data, as usually observed and reported in the literature. The *pseudo*-kurtosis[8] for the three data are 10.67 for the daily returns, 5.36 for the weekly returns and 5.60 for the monthly returns.

Next, we estimate the confidence interval and the stability parameter of the three times series by means of our MC estimation and test procedure. Figure 5 illustrates the estimates and the corresponding confidence intervals, where the solid line gives $1 - p$ values at given $H_0 : \alpha = \alpha_0$ of the empirical data and the three dashed lines (from bottom to top) give simulated quantiles of 90%, 95% and 99% for the estimate $\hat{\alpha} = \alpha_0$.

Figure 5 somewhere here.

The results of the estimates are numerically summarized in Table 2 again.

Table 2 somewhere here.

[8]Under the assumption of α-stable distributions with $\alpha < 2$, there exists no fourth moment.

Some comments on the empirical results are in order. First, the changes in the probability at given $H_0 : \alpha = \alpha_0$ shown in Figure 5 are the smoother as T increases. For a very small sample size, one can even observe a non-monotone probability curve. This is because (even if the same unit and exponential r.v. are used in the transformation into α-stable r.v.) for different α values the k changes for each α. Second, the increasing and decreasing of the probability around $H_0 : \alpha = \alpha_0$ are not necessarily symmetric. This means that, as the results show, the exact confidence intervals from the MC test for finite samples can be asymmetric, which is another advantage of our exact confidence interval for finite samples. Note that the asymptotic distribution of the Hill estimate, namely the normal distribution, is symmetric for all sample sizes and all quantiles. And there is no statistical background that the distribution of estimates for finite samples must be symmetric. Third, the estimated stability parameter for the daily data is 1.69, but it reduces to 1.84 when the observation frequency decreases to a weekly interval, where both of the two estimates are via our exact confidence intervals highly significant for $\alpha < 2$. Theoretically, however, the stability parameter should not change depending on sample size and/or observation frequency. This is only the case when the data are i.i.d. However, most of empirical financial data do not conform to this condition and are highly correlated in the second moment, as is well known. The changes of frequency from weekly to monthly, however, are not the case. The estimate for the monthly data, via the exact confidence interval is not (strongly) significant for the hypothesis $\alpha < 2$, and, probably the increasing of estimation inefficiency is more dominant than the effect of dependency in the data. Fourth, the confidence intervals becomes smaller and smaller as sample size increases, as it should be. The comparison of the two confidence intervals from our method and the normal distribution shows that the densities of the MC estimates and, hence, the confidence intervals for small sample sizes seem to be (right-skewed) asymmetric. This can be seen for the sample of the monthly (563 observations) and still somehow the weekly data (2,452 observations) because the right limits of the exact confidence intervals (90%, 95% and 99% or rather equally speaking 95%, 97.5% and 99.5% quantiles) are larger than those of the normal distribution, although the left limits of the two confidence intervals are approximately the same. This asymmetry vanishes when the sample size is very large. This (right-skewed) asymmetry can often be observed in the density of the Hill-type estimations, such as Hill-, Pickands- and Dekkers, Einmahl and de Haan estimate, as documented in Rachev and Mittnik (2000, Ch. 3).

The main result from the empirical application, which should be emphasized with respect to our exact confidence intervals, is that the hypothesis of $\alpha < 2$ for the monthly data cannot be accepted by the exact confidence intervals, but, by the asymptotic confidence interval, the hypothesis is still valid at a significance level of 97.5%. This means that results of the hypothesis for or against $\alpha < 2$ for finite samples, especially small sample sizes, can be misleading if they are concluded by the asymptotic normal distributions. As already discussed, for large samples (here, daily data with a size of 11,795) the usual confidence intervals from the exact test method and the normal distribution are almost the same, whereas for middle-large size (here, weekly data with size of 2,452) the two confidence intervals are slightly different, especially in the right limits.

Now, we turn to the problem of choice of k in empirical works. In order to demonstrate the behavior of the Hill estimation especially depending on the choice of the optimal k, we estimate the three empirical data sets using the Hill estimation with $k := \tau n, \tau \in (0, 1)$ and n being sample size. Specifically, we choose $\tau = 0.1 : 0.0001 : 0.5$.[9] Figure 6 shows the 4001 Hill estimates (solid line) and the MC estimate (dashed line) for each set of the empirical data, where the y-axis shows the estimated α and the x-axis shows the chosen k for the Hill estimation.

Figure 6 somewhere here.

Figure 6 demonstrates again that the hill estimation could be empirically almost useless, unless the k is known, because it is so sensitive to the choice of the k. The bias seems to have a linear relationship to k for large samples. In the case of the daily (weekly, monthly) data, we have a positive [negative] bias if we have chosen τ smaller [larger] than 0.4243 (0.4345, 0.4336).

In the second part of our empirical applications, we apply the test for constancy of tail thickness suggested by Quintos et al. (2001) to the same weekly return data as before. The aim of this part is again to demonstrate that empirical results can crucially depend on the choice of unknown k. In other words, because our MC procedure automatically uses an optimal k it is able to provide correct results regarding underlying tail thickness parameter.

[9] As the Hill horror plot (Figure 1) shows it seems to make no sense to consider $\tau < 0.1$ and $\tau > 0.5$ when the α is relatively high, say $\alpha > 1.5$.

Quintos et al. (2001) propose three tests for the constancy of tail thickness, namely recursive, rolling and sequential tests. We only employ the rolling test[10], by which we can expect the largest fluctuation in the estimates for the tail thickness. This is because the Hill estimator is conditional on the largest k observations, so that, for the recursive and the sequential test, the outlier behavior that appears in the initial sample remains in the selection of the k largest observations in the latter part of the sample, whereas, for the rolling test, the k largest observations can more probably change from one subsample to the next subsample. Under the assumption of an optimal k they formulate a statistic for the rolling test as

$$V_T(i) = \frac{n_i k(n_i)}{T} \left(\frac{\hat{\alpha}_i}{\hat{\alpha}_T} - 1 \right), \tag{11}$$

where $n_i, i = 1, 2, ..., N$ is size of N subsamples; $k(n_i)$ is the number of the observations on the distributional tail in a subsample; T is size of the whole sample; $\hat{\alpha}_i$ is the estimate of α for the i-th sample; and $\hat{\alpha}_T$ is the estimate of α for the whole sample. The test statistic for the rolling test focuses on the maximum of the statistic (11) for $k(n_i)/T \in \mathrm{IR}_\pi := [\pi;\ 1 - \pi]$, where π represents some small value (usually 0.1) commonly used in the construction of structure constancy tests, see Andrews (1993), and is given as

$$Q^* = sup_{r \in \mathrm{IR}_\pi} V_T([Tr]) \xrightarrow{d} sup_{r \in \mathrm{IR}_\pi} \bar{W}(r, \gamma_0), \tag{12}$$

where $\bar{W}(r, \gamma_0) = W(r, \gamma_0) - (r - s)W(1, 1)$ with $W(r)$ denoting a standard Wiener process. The statistic in (12) converges to a non-standard distribution and the critical values from the non-standard distribution are tabulated in Quintos et al. (2001).

Specifically, we choose two rolling subsamples: 10-years and 20-years subsample, i.e., our first subsample for the 10-years (20-years) subsample contains the weekly return from October 1959 to September 1969 (September 1979), and our last subsample contains the weekly returns from October 1996 (October 1986) to September 2006. Therefore, the number of subsamples is 37 and 27 for the 10-years and 20-years subsample, respectively. The size of subsamples is either 523 (27 times) or 522 (10 times) for the 10-years subsample and (1044) (14 times) or 1043 (13 times) for the 20-years subsample. For calulation of the test statistic, therefore, we take 523 and 1044 as size of a rolling sample for the 10-years subsample and the 20-years

[10] The recursive and sequential test show mainly the same result.

subsample. Figure 7 and 8 show empirical test statistics based on the MC method and the Hill estimation, where Figure 7 (8) corresponds to the 20-years (10-years) subsample.

Figure 7 and 8 somewhere here.

For both graphs, the mountain-looking solid curve shows the empirical test statistic in (12) based on the 4001 Hill estimates for $\tau = [= .1 : 0.001 : 0.5]$ and the dashed line based on the MC estimates. The three straight solid lines show 99%, 95% and 90% critical values of 3.075, 2.25 and 1.865 for the 20-years subsample in Figure 7 and 2.37, 1.81 and 1.54 for the 10-years subsample in Figure 8[11].

Figures 7 and 8 show that, according our MC estimate, the stability of tail thickness of the DAX weekly returns seems to be constant during the whole period for both subsample lengths. However, the results of Hill estimate for both cases depend crucially on which $k(= 2452\tau)$, or rather τ is chosen for the estimation. If τ is chosen smaller 0.3740 and 0.3862 (except some very small τ values for both cases) for the 20-years and 10-years subsamples, respectively, we would not accept the null hypothesis of constancy of the tail thickness at a significance level of 95%.

When comparing the two cases, the results for the 10-years subsample, i.e., a sample size of 523 (Figure 8) is more erratic than those for the 20-years subsample, i.e., a sample size of 1044. This can be seen in the range of even relatively higher k/n, say $k/n > 0.4$.

6 Summary

In this paper we have considered an exact test and estimation method for the stability parameter of α-stable distributions using the Monte Carlo technique. Specifically, we have employed the Hill estimator for constructing the MC test statistic. Using our MC test and estimation, we improved the Hill estimation in two ways: our MC estimation does not need to assume that the number of observations on the distributional tail is known. Moreover, our MC test provides exact confidence intervals for finite samples.

[11] The critical value is calculated by a linear interpolation from the table in Quintos et al. (2001) p. 662.

The empirical applications demonstrate that empirical conclusions about the hypothesis test for $\alpha < 2$ for finite samples and/or for constancy of tail thickness depend crucially on the choice of the number of observations on the distributional tail and the exact confidence intervals.

References

Andrews, D.W.K. (1993): Tests for parameter stability and structural change with unknown change point, *Econometrica* 59, 817–858.

Barnard, G.A. (1963): Comment on "the spectral analysis of point processes" by M.S. Bartlett, *Journal of the Royal Statistical Soceity* Series B 25, 294.

Birnbaum, Z.W. (1974): Computers and unconventional test-statistics, in *Reliability and biometry*, ed. F. Proschan and R.J. Serfling, 441–458.

Candelson, B. and S. Straetmans (2006): testing for multiple breaks in the tail behavior of emerging currency returns, *Journal of International Money and Finance* 25, 1187–1205.

de Haan, L. and L. Peng (1998): Comparison of tail index estimators, *Statistics Neerlandica* 52, 60–70.

Dekkers, A.L.M., J.H.J. Einmahl and L. de Haan (1989): A moment estimator for the index of an extreme-value distribution, *Annals of Statistics* 17, 1833–1855.

Dufour, J.-M. (2005): Monte Carlo tests with nuisance parameters: a general approach to finite-sample inference and standard asymptotics, forthcoming in *Journal of Econometrics*

Dufour, J.-M., L. Khalaf and M.-C. Beaulieu (2007): Exact confidence set estimation and goodness-of-fit test methods for asymmetric heavy tailed stable distributions, Université Laval, Manuscript.

DuMouchel, W.H. (1971): Stable distributions in statistical inference, Ph.D. Thesis, Yale University.

DuMouchel, W.H. (1983): Estimating the stable index α in order to measure tail thickness: A critique, *Annals of Statistics* 11, 1019–1031.

Dwass, M. (1957): Modified randomization tests for nonparametric hypotheses, *Annals of Mathematical Statistics* 28, 181-187.

Deheuvels, P., E. Häusler and D.M. Mason (1988): Almost sure convergence of the Hill estimator, *Math. proc. Cambridge Philos. Soc.* 104, 372–381.

Fama, E.F. and K.R. French (1992): The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.

Goldie, C.M. and R.L. Smith (1987): Slow variation with remainder: a survey of the theory and its applications, *Quarterly Journal of Mathematics Oxford* 38, 45–71.

Hill, B.M. (1975): A simple general approach to inference about the tail of a distribution, *Annals of Statistics* 3, 1163–1174.

Hodges, J.L.Jr. and E.L. Lehmann (1963): Estimates of location based on rank tests, *The Annals of Mathematical Statistics* 34, 598–611.

Hsing, T. (1991): On tail index estimation using dependent data, *Annals of Statistics* 19, 1547–1569.

Huisman, R., K.G. Koedijk, C.J.M. Kool and F. Palm (2001): Tail-index estimates in small samples, *Journal of Business & Economic Statistics* 19, 208–216.

Kim, J.-R., S. Mittnik and S.T. Rachev (1997): Detecting asymmetries in observed time series and unobserved disturbances, *Studies in Nonlinear Dynamics and Econometrics* 1, 131–143.

Kurz-Kim, J.-R. and M. Loretan (2007): A note on the coefficient of determination in models with infinite-variance variable, Discussion paper 10/2007, Deutsche Bundesbank.

Mandelbrot, B. (1963), The variation of certain speculative prices, *Journal of Business* 36, 394–419.

Mason, D.M. (1982): Laws of large numbers for sums of extreme values, *Annals of Probability* 10, 756–764.

McCulloch, J.H. (1986): Simple consistent estimators of stable distribution parameters, *Communications in Statistics, Computation and Simulation* 15, 1109–1136.

McCulloch, J.H. (1996): Financial applications of stable distributions, *Handbook of Statistics - Statistical Methods in Finance* 14, 393–425, (ed. Maddala, G.S. and C.R. Rao), Elsevier Science B.V., Amsterdam.

McCulloch, J.H. (1997): Measuring tail thickness to estimate the stable index α: A critique, *Journal of Business and Economic Statistics* 15, 74–81.

Pickands, J. (1975): Statistical inference using extreme order statistics, *Annals of Staistics* 3, 119-131.

Quintos, C., Z. Fan and P.C.B. Phillips (2001): Structural change tests in tail behaviour and the Asian crisis, *Review of Economic Studies* 68, 633-663.

Rachev, S.T., J.-R. Kim and S. Mittnik (1999): Stable Paretian econometrics part I and II, *The Mathematical Scientists* 24, 24–55 and 113–127.

Rachev, S.T. and S. Mittnik (2000): Stable Paretian Models in Finance, John Wiley & Sons: Chichester.

Samorodnitsky, G. and M.S. Taqqu (1994): *Stable Non-Gaussian Random Processes*, Chapman & Hall: New York.

Zolotarev, V.M. (1986), *One-Dimensional Stable Distributions*, Translations of Mathematical Monographs, American Mathematical Society, Vol. 65, Providence.

Appendix

Table A. Root mean square error of Hill estimator with different centering

α	Centering / Sample size	True	Mean	Median
1.0	100	0.1660	0.3338	0.1654
	250	0.1235	0.2890	0.1234
	500	0.0953	0.2697	0.0952
	1000	0.0727	0.2711	0.0727
	5000	0.0442	0.1763	0.0442
1.25	100	0.1616	0.2488	0.1625
	250	0.1063	0.2090	0.1063
	500	0.0762	0.1845	0.0764
	1000	0.0546	0.1771	0.0545
	5000	0.0253	0.1263	0.0253
1.5	100	0.1808	0.1958	0.1813
	250	0.1126	0.1306	0.1135
	500	0.0806	0.0949	0.0809
	1000	0.0579	0.0761	0.0578
	5000	0.0263	0.0368	0.0262
1.75	100	0.2048	0.2053	0.2068
	250	0.1291	0.1306	0.1299
	500	0.0909	0.0912	0.0903
	1000	0.0640	0.0655	0.0641
	5000	0.0290	0.0294	0.0290
1.95	100	0.2258	0.2261	0.2259
	250	0.1421	0.1414	0.1419
	500	0.0985	0.0987	0.0992
	1000	0.0696	0.0698	0.0699
	5000	0.0318	0.0317	0.0317
2.0	100	0.2303	0.2292	0.2313
	250	0.1463	0.1452	0.1455
	500	0.1004	0.1010	0.1016
	1000	0.0716	0.0714	0.0711
	5000	0.0330	0.0328	0.0327

Table 1. Asymptotic test at 95% significance level

		α				
n	Centering	1	1.25	1.5	1.75	2
100	mean	8.55	6.43	3.53	4.13	3.47
	median	4.08	4.48	4.31	4.97	4.26
200	mean	12.47	10.54	4.68	3.60	3.84
	median	4.26	4.62	4.26	3.85	4.08
500	mean	15.34	13.25	5.12	3.80	3.77
	median	4.49	4.76	3.96	3.82	3.70
1000	mean	19.39	16.23	5.44	4.08	3.56
	median	5.24	4.43	4.11	4.05	3.40
5000	mean	16.53	20.04	5.73	4.27	4.02
	median	4.95	4.71	4.39	3.94	4.02

Table 2. Estimated stability parameters and exact confidence intervals[a]

Quantile / Data	0.5%	2.5%	5%	$\hat{\alpha}_H(\alpha_0)$	95%	97.5%	99.5%
Daily	1.66	1.67	1.67	**1.69**	1.72	1.72	1.74
	(1.65)	(1.66)	(1.66)		(1.72)	(1.72)	(1.73)
Weekly	1.74	1.76	1.78	**1.84**	1.91	1.93	1.98
	(1.74)	(1.77)	(1.78)		(1.90)	(1.91)	(1.94)
Monthly	1.61	1.66	1.69	**1.82**	2	2	2
	(1.62)	(1.67)	(1.69)		(1.95)	(1.97)	(2)

[a]The values in parentheses are corresponding confidence intervals based on the normal distribution given in equation (2).

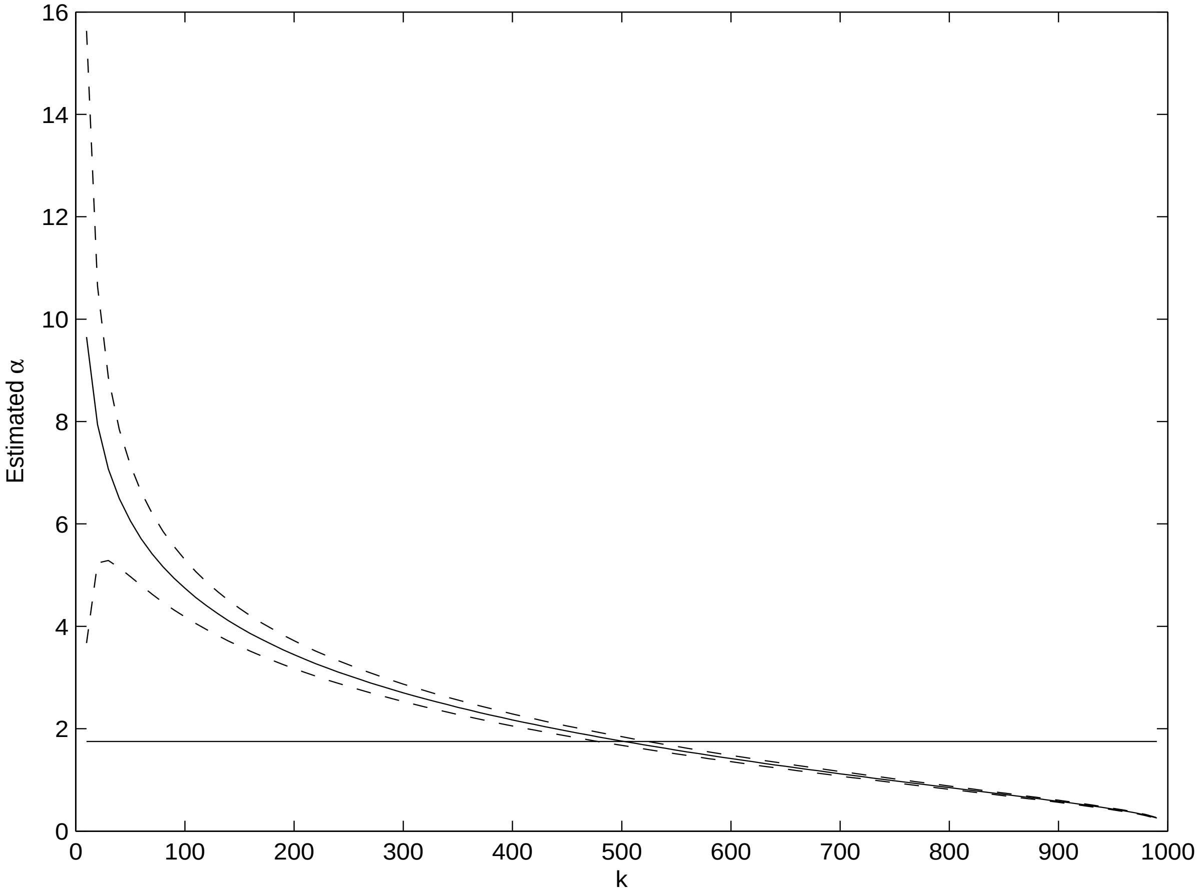

Figure 1: Hill horror plot

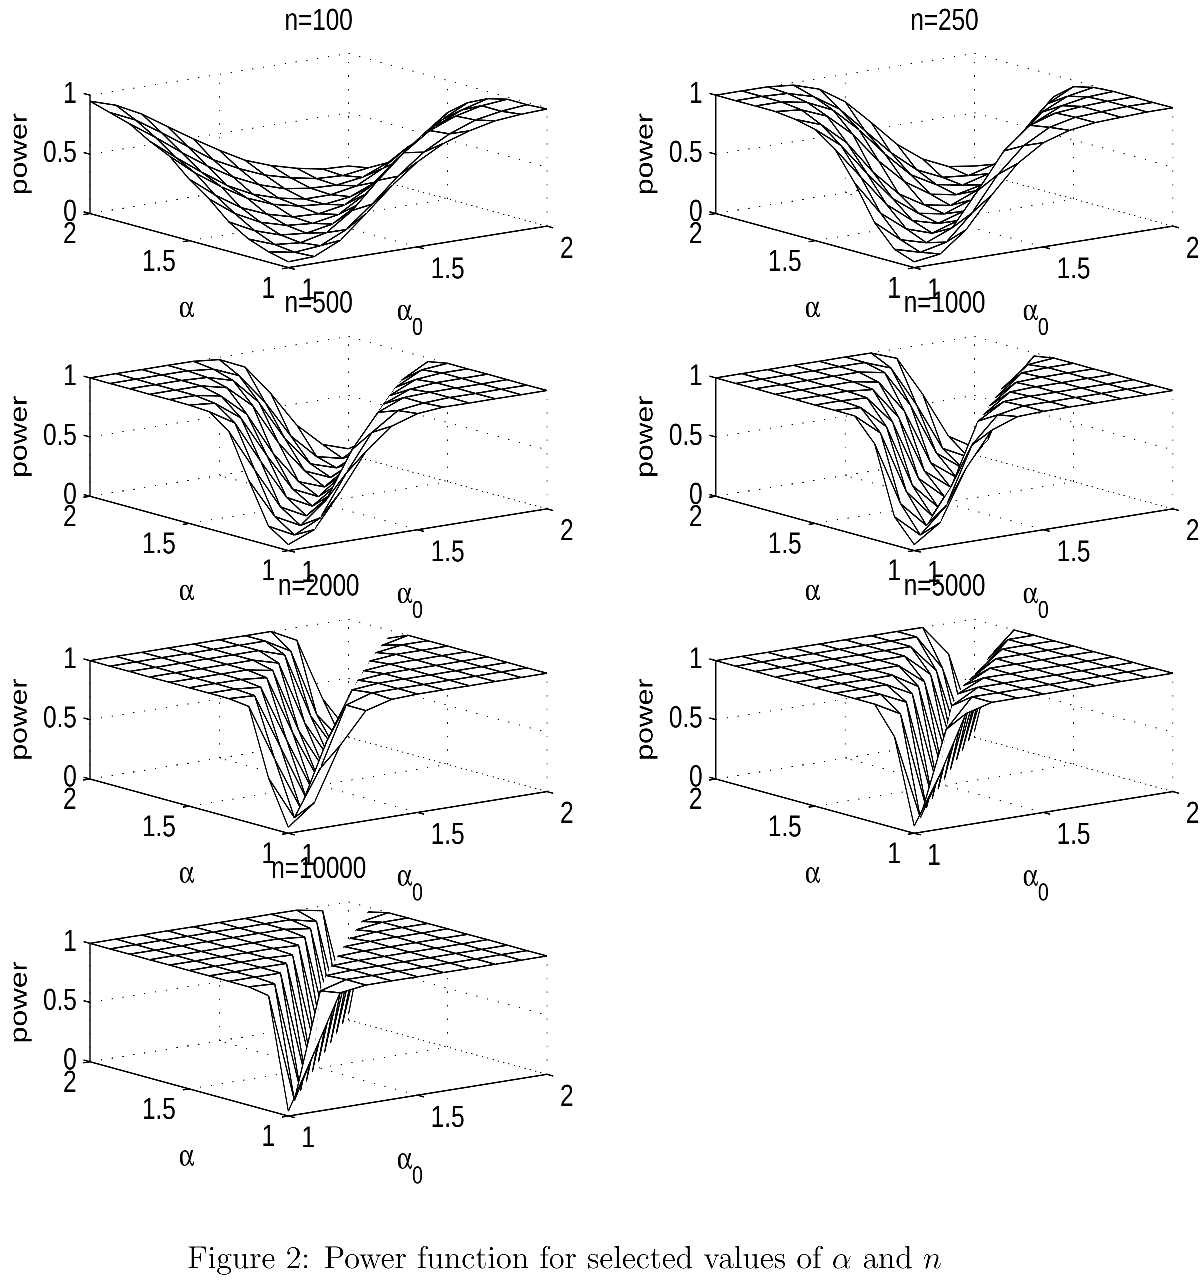

Figure 2: Power function for selected values of α and n

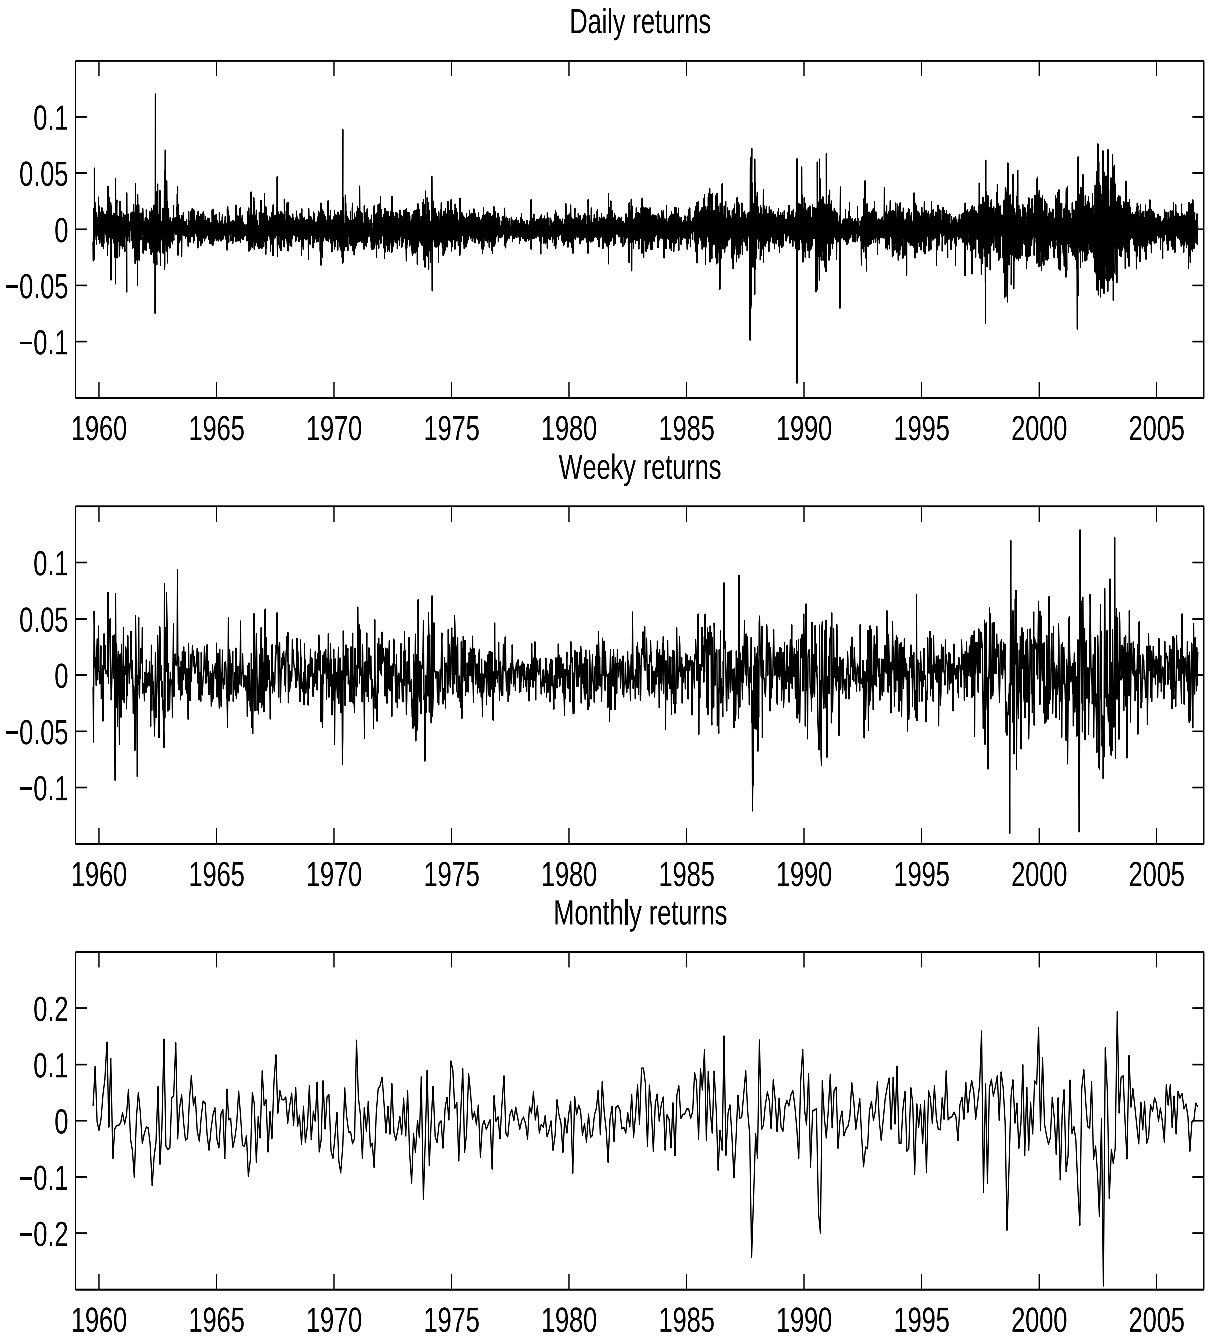

Figure 3: DAX returns in difference frequencies

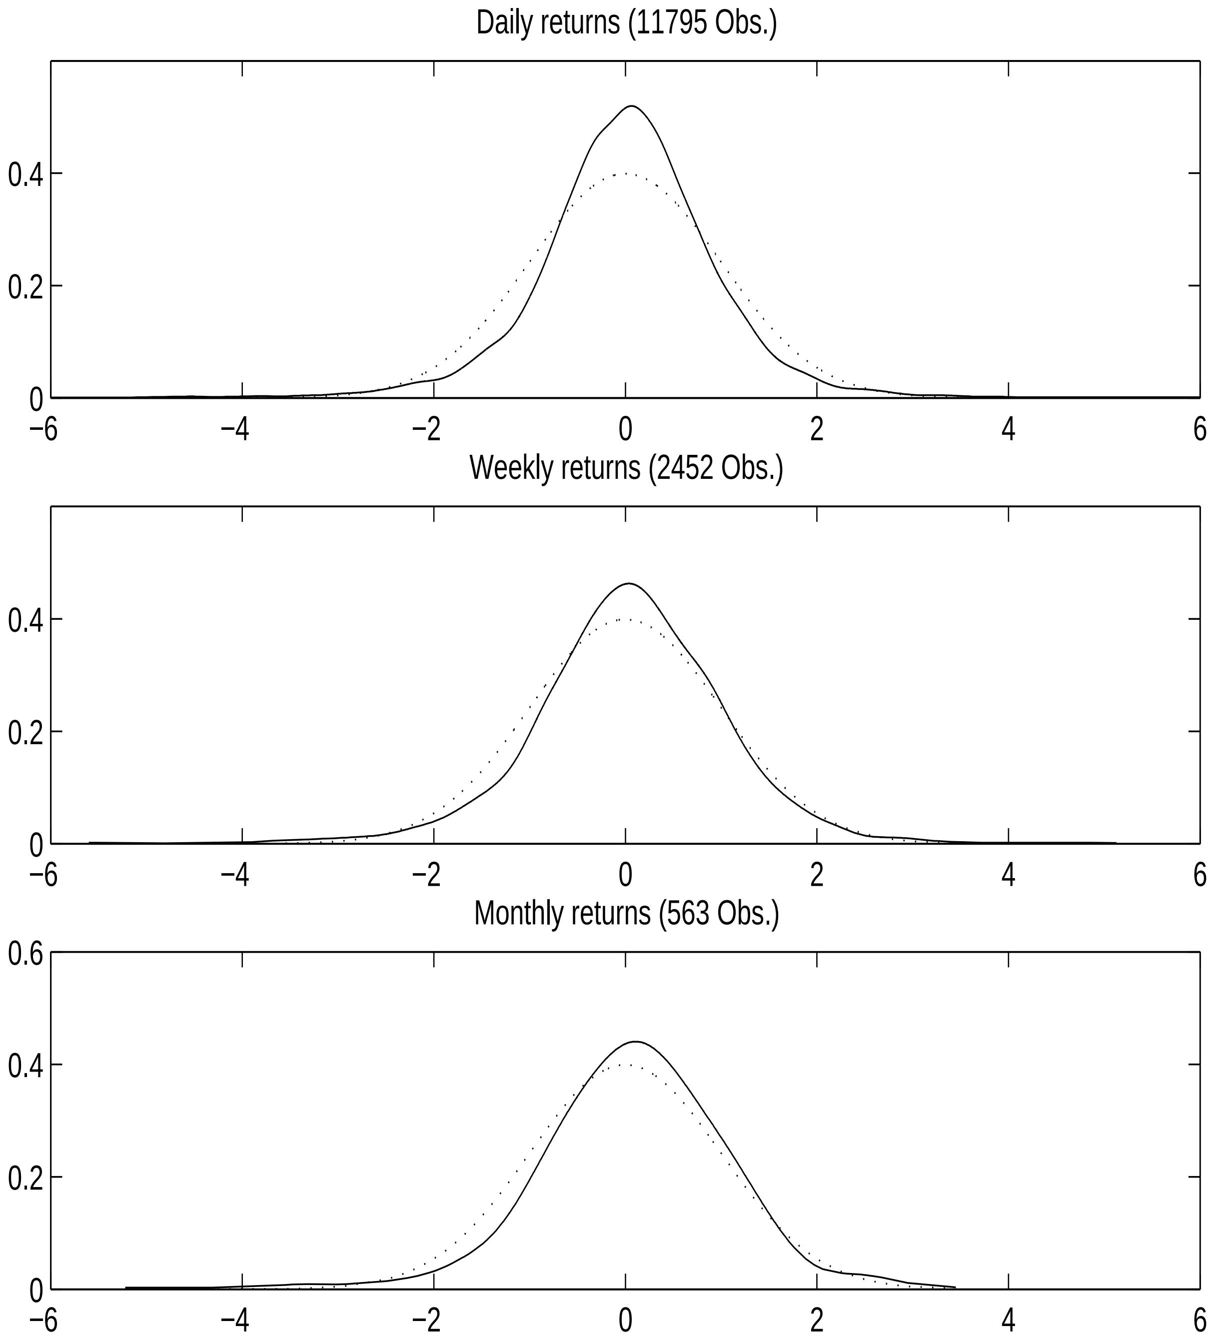

Figure 4: Empirical densities

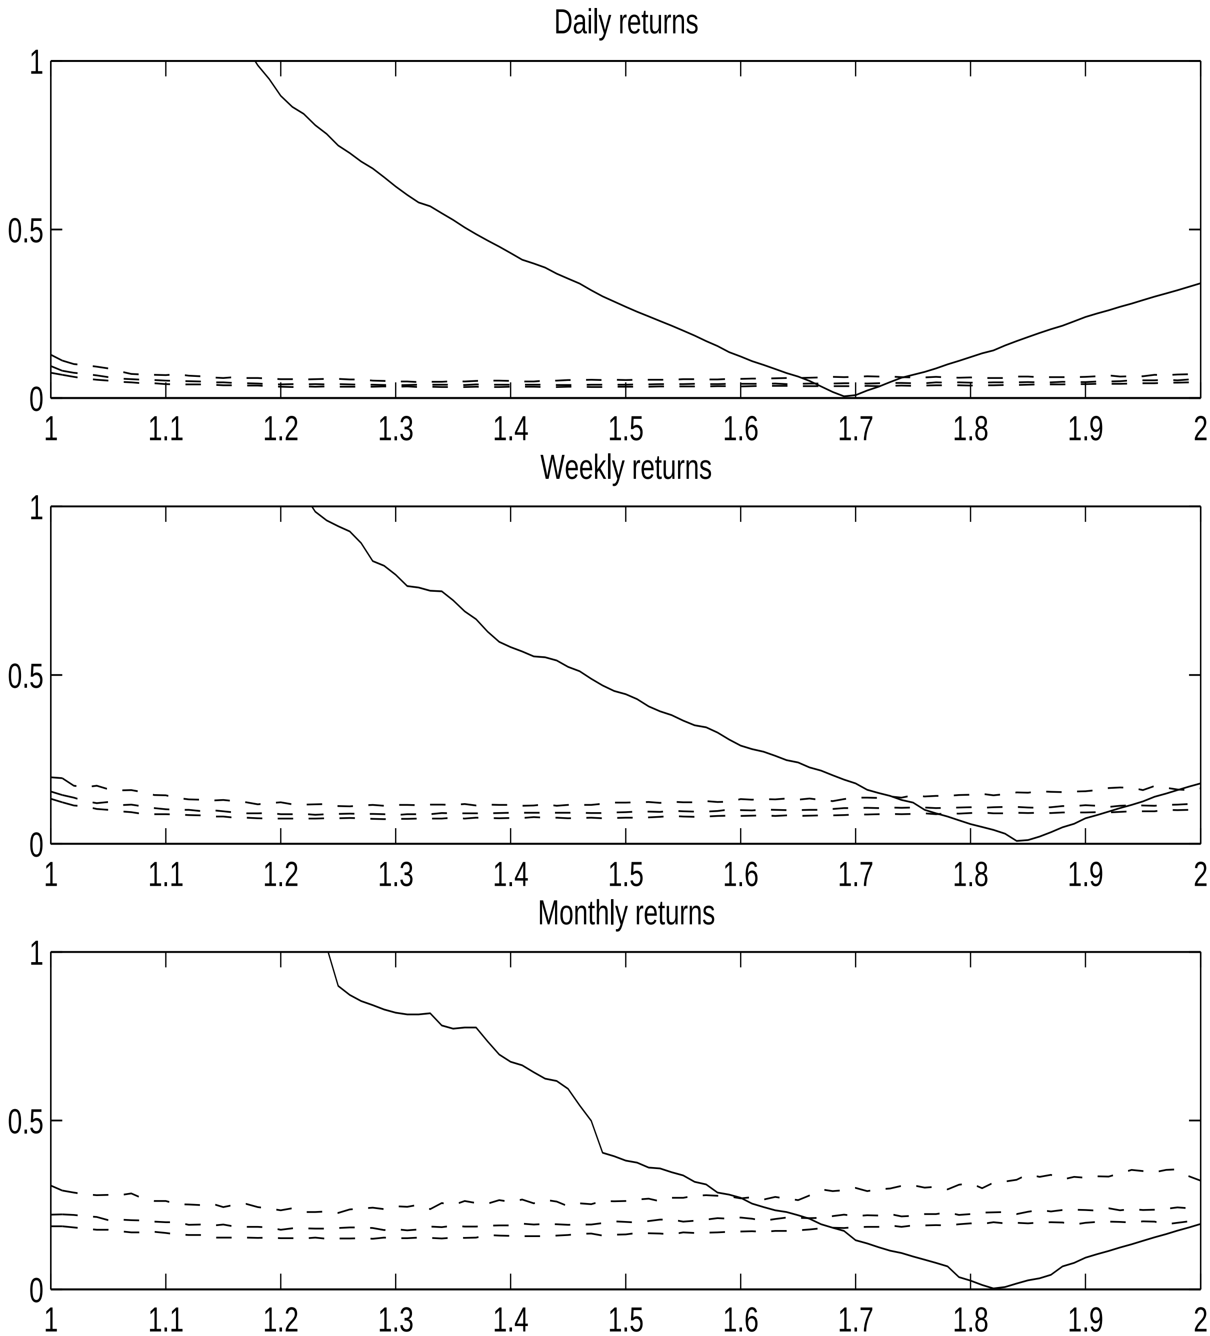

Figure 5: Exact confidence intervals and estimates

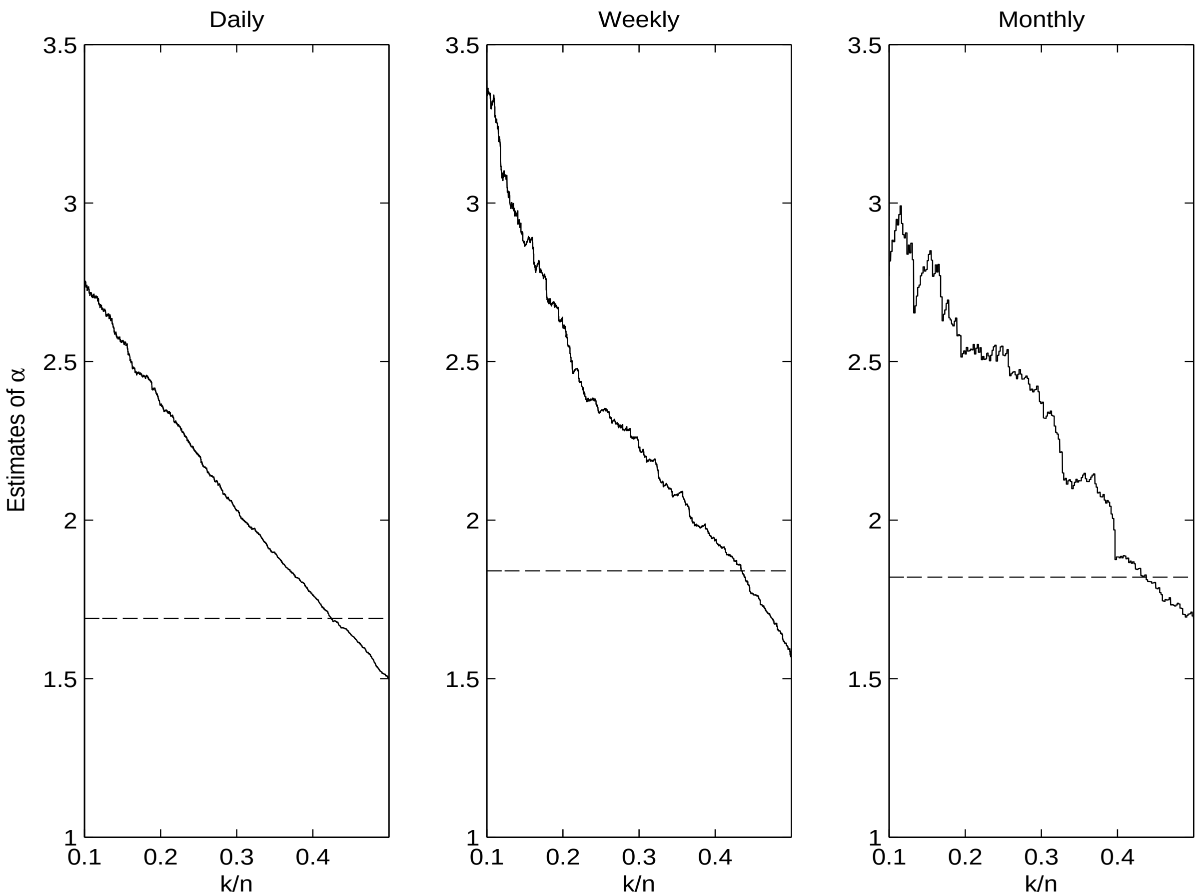

Figure 6: Hill estimates depending unknown k

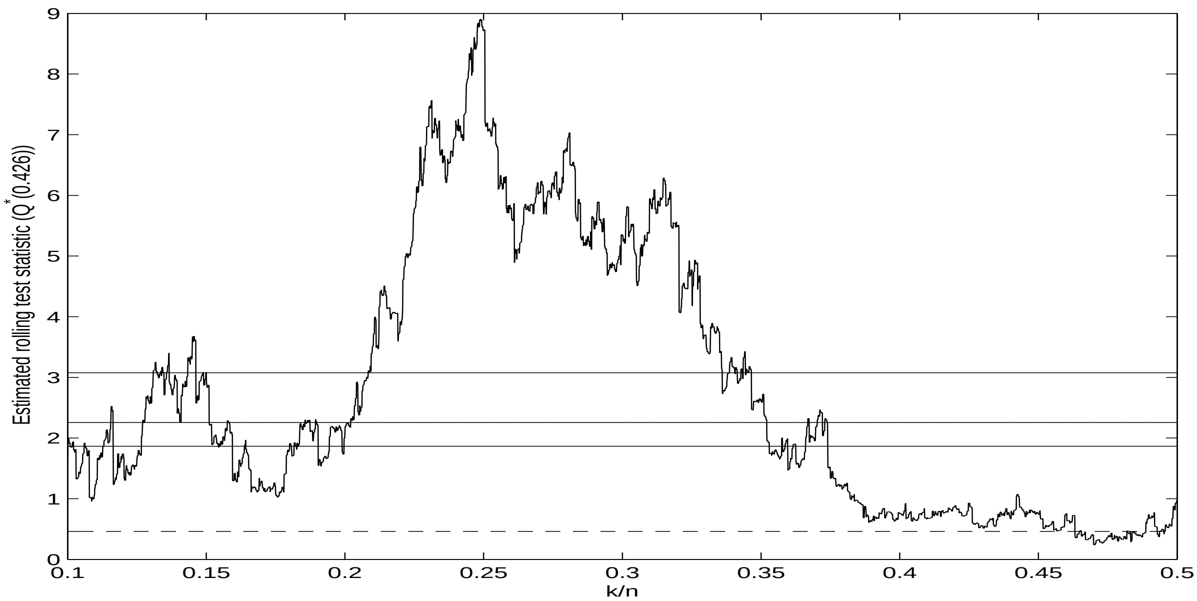

Figure 7: Rolling test for 20-years subsamples

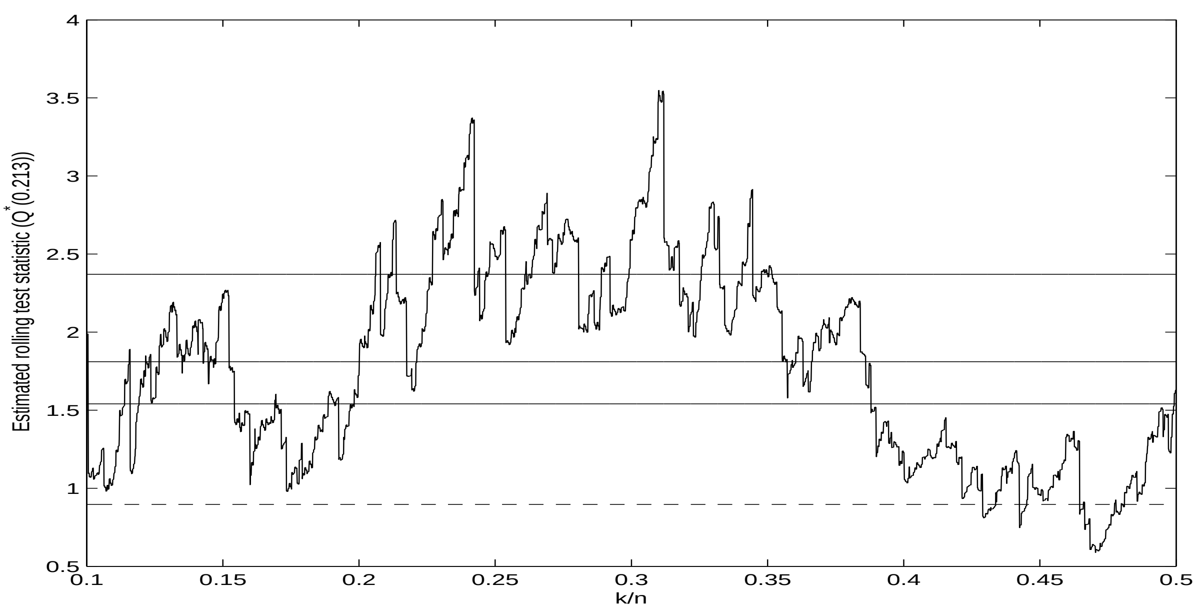

Figure 8: Rolling test for 10-years subsamples