Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Banking & Finance 29 (2005) 1017–1035
Journal of BANKING & FINANCE
www.elsevier.com/locate/econbase

The choice of the distribution of asset returns: How extreme value theory can help?

François Longin

Department of Finance, ESSEC Graduate Business School, Avenue Bernard Hirsch, B.P. 105, Cergy-Pontoise Cédex 95021, France

Available online 7 October 2004

Abstract

One of the issues of risk management is the choice of the distribution of asset returns. Academics and practitioners have assumed for a long time (for more than three decades) that the distribution of asset returns is a Gaussian distribution. Such an assumption has been used in many fields of finance: building optimal portfolio, pricing and hedging derivatives and managing risks. However, real financial data tend to exhibit extreme price changes such as stock market crashes that seem incompatible with the assumption of normality. This article shows how extreme value theory can be useful to know more precisely the characteristics of the distribution of asset returns and finally help to chose a better model by focusing on the tails of the distribution. An empirical analysis using equity data of the US market is provided to illustrate this point.

JEL classification: C51; G28

0. Introduction

The statistical distribution of asset returns plays a central role in financial modeling. Assumptions on the behavior of market prices are necessary to test asset pricing

E-mail address: longin@essec.fr

0378-4266/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.jbankfin.2004.08.011

theories, to build optimal portfolios by computing risk/return efficient frontiers, to value derivatives and define the hedging strategy over time, and to measure and manage financial risks. However, neither economic theory nor statistical theory exist to assess the exact distribution of returns. Distributions used in empirical and theoretical research are always the result of an assumption or estimation using data. The paradigm adopted in finance is the Gaussian distribution. In the 1950s and 1960s, Markowitz (1952) and Sharpe (1964) assume normality for asset returns when studying portfolio selection and deriving the capital asset pricing model. In the beginning of the 1970s, Black and Scholes (1973) and Merton (1973) derived the price and the hedging strategy of an option by assuming a Brownian motion for the price of the underlying asset implying a Gaussian distribution for returns. More recently, with the changes in the banking and financial regulation on risks and capital, value-at-risk models developed and implemented by financial institutions also rely intensively on the Gaussian distribution.

Although normality is the paradigm in financial modeling, several alternatives have been considered. The main reason for looking at other models is that there is a growing evidence that the Gaussian distribution tends to underestimate the weight of the extreme returns contained in the distribution tails. For example, the stock market crashes of 1929 and 1987, corresponding to daily market drops of more than 10% and 20% respectively, are very unlikely in a world governed by normality. Several other candidates have been proposed in the academic literature and used with more or less success by practitioners: a mixture of Gaussian distributions, stable Paretian distributions, Student-*t* distributions and the class of ARCH processes. One problem with these alternatives is that they are *not nested* and then not directly comparable (by carrying out a likelihood ratio test for example).

In this paper I propose a method which allows one to discriminate among these different models. I look at the two extreme parts of the distribution: the left tail and the right tail. The form of the tails is different for the models cited above as the weight of extremes varies. I use *extreme value theory*, which provides a measure of the importance of extremes in the distribution of returns. [1] This measure called the *tail index* is used to build a formal test to discriminate among the models commonly used.

The remainder of the paper is organized as follows: Section 1 presents extreme value theory while Section 2 gives the different methods of estimation of the statistical distribution of the extremes. Section 3 describes the application of extreme value theory for discriminating among distributions of returns. The empirical analysis is then presented in Section 4. The last section concludes.

[1] Earlier work applying extreme value theory in finance can be found in Jansen and De Vries (1991), Longin (1993) and Loretan and Phillips (1994). These studies focus on the distribution tails of the US stock market returns.

1. Extreme value theory

This section presents the main results for extreme value theory. Two approaches are usually considered to define extremes: the minimum/maximum approach and the negative/positive exceedance approach.

1.1. The distribution of returns

Starting with the notations, R will stand for the (logarithmic) return of the asset or portfolio computed over a given time-interval, f_R and F_R, respectively the density probability and cumulative distribution functions of the random variable R. The support of the density function is noted as $[l, u]$, the lower and upper bounds, l and u, being possibly equal to infinity (it is the case for the Gaussian distribution). Let $R_1, R_2, \ldots, R_n$ be n returns observed on n time-intervals of frequency f.

1.2. Extremes defined as minimal and maximal returns

Extremes can be defined as the minimum and the maximum of the n random variables $R_1, R_2, \ldots, R_n$. I note Y_n the highest return (the maximum) and Z_n the lowest return (the minimum) observed over n trading time-intervals. [2]

The extreme value theorem (EVT) is interested in the statistical behavior of the minimum and maximum of random variables. It is analogous to the central limit theorem (CLT), which is interested in the statistical behavior of the sum of random variables. Both theorems consider the asymptotic behavior of the variables in order to get results that are independent of the initial distribution. In the EVT framework extremes will have to be selected from very long time interval such as the sum is computed over very long time-interval in the CLT framework. In order to get non-degenerated limiting distributions, the variables of interest have first to be standardized. This is illustrated below in the EVT case.

If the variables $R_1, R_2, \ldots, R_n$ are statistically independent and drawn from the same distribution (hypothesis of the random walk for stock market prices), then the exact distribution of the maximum Y_n is simply given by

$$F_{Y_n}(r) = (F_R(r))^n. \tag{1}$$

The distribution of extremes depends mainly on the properties of F_R for large values of r. Indeed, for small values of r, the influence of $F_R(r)$ decreases rapidly with n. Hence, the most important information about the extremes is contained in the tails of the distribution of R. From formula (1), it can be concluded that the limiting distribution of Y_n is null for r less than the upper bound u and equal to one for r greater than u. It is a degenerate distribution.

[2] In the remainder of the paper, theoretical results are presented for the maximum only, since the results for the minimum can be directly deduced from those of the maximum of the opposite variable using the following relation: $Z_n(R) \equiv \mathrm{Min}(R_1, R_2, \ldots, R_n) = -\mathrm{Max}(-R_1, -R_2, \ldots, -R_n) \equiv -Y_n(R)$.

As explained in Longin (1996), the exact formula of the extremes and the limiting distribution are not, however, especially interesting. In practice, the distribution of the parent variable is not precisely known and, therefore, if this distribution is not known, neither is the exact distribution of the extremes. For this reason, the asymptotic behavior of the maximum Y_n is studied. Tiago de Oliveira (1973) argues, "As, in general, we deal with sufficiently large samples, it is natural and in general sufficient for practical uses to find limiting distributions for the maximum or the minimum conveniently reduced and use them". To find a limiting distribution of interest, the random variable Y_n is transformed such that the limiting distribution of the new variable is a non-degenerate one. The simplest transformation is the standardization operation. The variate Y_n is adjusted with a scaling parameter α_n (assumed to be positive) and a location one β_n. In the remainder of the paper, the existence of a sequence of such coefficients ($\alpha_n > 0$, β_n) is assumed. Extreme value theory specifies the possible non-degenerate limit distributions of extreme returns as the variable n tends to infinity. [3] In statistical terms, a limit cumulative distribution function denoted by G_{Y_n} satisfies the following condition:

$$\lim_{n\to+\infty} \sup_{l<y<u} |F_{Y_n}(y) - G_{Y_n}(y)| = 0.$$

Gnedenko (1943) shows that the extreme value distribution (EVD) is the only non-degenerate distribution which approximates the distribution of extreme returns F_{Y_n}. The limit distribution function G_{Y_n} is given by

$$G_{Y_n}(y) = \exp\left(-\left(1 + \tau \cdot \left(\frac{y - \beta_n}{\alpha_n}\right)\right)\right)^{-(1/\tau)}. \tag{2}$$

The parameter τ, called the tail index, gives a precise characterization of the tail of the distribution of returns. Distributions with a power-declining tail (fat-tailed distributions) correspond to the case $\tau > 0$, distributions with an exponentially-declining tail (thin-tailed distributions) to the case $\tau = 0$, and distributions with no tail (finite distributions) to the case $\tau < 0$. [4] The extreme value distribution is called a Fréchet distribution, a Gumbel distribution and a Weibul distribution. The Gumbel distribution can be regarded as a transitional limiting form between the Fréchet and the Weibull distributions.

The extreme value theorem gives an interesting result: whatever the distribution of the parent variable R, the limiting distribution of the extremes always has the same form. The distribution of the extremes for two different parent processes is differentiated by the values of the standardizing coefficients α_n and β_n and the tail index τ.

More interestingly, the same limiting distribution is obtained if the i.i.d. hypothesis is relaxed. Berman (1963) shows the same result stands if the variables are correlated and if the series of the squared correlation coefficients is finite. A common

[3] Proofs of the extreme values theorem and other claims can be found Gnedenko (1943) and in Gumbel (1958, Chapters 5 and 7) and Galambos's (1978) text books. See also Embrechts et al. (1997) and Reiss and Thomas (1997).

[4] Note that a different convention is sometimes used for the sign of the tail index.

model is a discrete mixture of Gaussian distributions. In this particular case, the Gumbel distribution is still the limiting distribution of the extremes (see Leadbetter et al., 1983). De Haan et al. (1989) show that if the returns followed an ARCH(1) process, the variable Y_n would have a limiting Fréchet distribution. Following their research I detail below the relationship between the parameters of the ARCH process and those of the distribution of the extremes. Recall that an ARCH(1) process is given by two equations:

$$R_t = E_{t-1}(R_t) + \varepsilon_t, \tag{3a}$$

$$h_t = E_{t-1}(R_t - E_{t-1}(R_t))^2 = a_0 + a_1 \cdot \varepsilon_{t-1}^2. \tag{3b}$$

The realized return R_t observed at time t is decomposed into an expected part noted as $E_{t-1}(R_t)$ computed one period before at time $t-1$ and an unexpected part noted as ε_t known at time t only. The expected variance h_t varies over time and is conditioned upon the past value of the innovation ε_{t-1}. ARCH models reflect quite well the time-varying behavior of volatility and especially the clustering of extremes. After a big shock (i.e. a large value for ε_{t-1}) one expects a high level of variance and then more big shocks in the future. The coefficient a_1 reflects the persistence of volatility (or the correlation of absolute returns). A high value of a_1 implies a high level of persistence, a lot of clusters of extremes and finally a fat-tailed unconditional distribution of returns. The tail index τ is related to the degree of persistence a_1 by the following formula:

$$E((a_1 \cdot \varepsilon^2)^{1/\tau}) = 1. \tag{4}$$

Assuming a conditional Gaussian distribution for the innovation ε, Eq. (4) becomes

$$\Gamma\left(\frac{1}{\tau} + 0.5\right) = \sqrt{\pi} \cdot (2a_1)^{-\frac{1}{\tau}}, \tag{5}$$

where Γ is the gamma function and π the constant Pi. For a given value of the parameter a_1, a unique value of τ is obtained by solving Eq. (5). For example, for a_1 equal to 0.5, the tail index τ is equal to 0.42.

These results show that the assumption of independence is less important for extreme values than it would seem at first sight. Let us note that the extremes are (asymptotically) drawn from an unconditional distribution, even if the parent variable is drawn from a conditional distribution.

1.3. Extremes defined as return exceedances

Extremes can also be defined in terms of exceedances with reference to a threshold denoted by θ. For example, positive θ-exceedances correspond to all observations of R greater than the threshold θ. As results for negative exceedances can be deduced from those for positive exceedances by consideration of symmetry, I focus on the case ($R > \theta$) which defines the right tail of the distribution of returns. The probability that a return R is higher than θ, denoted by probability p_θ, is linked to the threshold θ and the distribution of returns F_R by the relation $p_\theta = 1 - F_R(\theta)$.

As explained in Longin and Solnik (2001), the cumulative distribution of θ-exceedances, denoted by F_R^θ and equal to $(F_R(x) - F_R(\theta))/(1 - F_R(\theta))$ for $x > \theta$, is exactly known if the distribution of returns F_R is known. However, in most financial applications, the distribution of returns is not precisely known and, therefore, neither is the exact distribution of return exceedances. For empirical purposes, the *asymptotic* behavior of return exceedances needs to be studied. Extreme value theory addresses this issue by determining the possible non-degenerate limit distributions of exceedances as the threshold θ tends to the upper point u of the distribution. In statistical terms, a limit cumulative distribution function denoted by G_R^θ satisfies the following condition:

$$\lim_{\theta\to u} \sup_{\theta<x<u} |F_R^\theta(x) - G_R^\theta(x)| = 0.$$

Balkema and De Haan (1974) and Pickands (1975) show that the generalized Pareto distribution (GPD) is the only non-degenerate distribution which approximates the distribution of return exceedances F_R^θ. The limit distribution function G_R^θ is given for $x > \theta$ by

$$G_R^\theta(x) = 1 - \left(1 + \tau \cdot \left(\frac{x-\theta}{\sigma}\right)\right)^{1/\tau}, \tag{6}$$

where σ, the dispersion parameter, depends on the threshold θ and the distribution of returns F_R, and τ, the tail index, is intrinsic to the distribution of returns F_R.

2. Estimation of the tail index

This section deals with the statistical estimation of the tail index. Two approaches are considered. First, in the so-called parametric approach, the parametric form of the asymptotic distribution of extremes is assumed to hold even though the database contains a finite number of observations. The parameter of the distribution of extremes, including the tail index, are directly estimated by classical methods such as the maximum likelihood method. Second, in the so-called non-parametric approach, no parametric distribution is assumed for the extremes.

2.1. The parametric approach

The parametric approach assumes that minimal returns and maximal returns selected over a given period are exactly drawn from the extreme value distribution given by Formula (2) or alternatively that negative and positive return exceedances under or above a given threshold are exactly drawn from the distribution given by formula (6). With either definition of extremes, the asymptotic distribution contains three parameters: τ, α_n and β_n. for extremes defined as minimal or maximal returns selected from a period containing n returns, or alternatively, τ, σ_θ and p_θ for extremes defined as negative or positive return exceedances under or above a given threshold θ. Under the assumption that the limit distribution holds, the maximum likelihood method gives unbiased and asymptotically normal estimators (see Tiago de Oliveira (1973) for the system of equations). The system of non-linear equations can be solved numerically using the Newton–Raphson iterative method. Note that the regression method (see Gumbel, 1958) gives biased estimates of the parameters but may be used to get initial values for the maximum likelihood algorithm.

In practice, the extreme value distributions can be estimated with different values of the number of returns contained in the selection period n (for minimal and maximal returns) and alternatively, with different values of the threshold θ (for negative and positive return exceedances). A goodness-of-fit test such as a Sherman test can then be carried out in order to choose the most relevant values from a statistical point of view.

2.2. The non-parametric approach

The previous methods assume that the extremes are drawn exactly from the extreme value distribution. Estimators for the tail index τ, which do not assume that the observations of extremes follow exactly the extreme value distribution, have been developed by Pickands (1975) and Hill (1975). These estimators are based on order statistics of the parent variable R.

Pickands's estimator for the right tail is given by

$$\tau_{\text{Pickands}} = -\frac{1}{\ln 2} \cdot \ln \frac{R'_{N-q+1} - R'_{N-2q+1}}{R'_{N-2q+1} - R'_{N-4q+1}} \tag{7}$$

where $(R'_t)_{t=1,N}$ is the series of returns ranked in an increasing order and q is an integer depending on the total number of returns contained in the database N. Pickands's estimator is consistent if q increases at a suitably rapid pace with N (see Dekkers and De Haan, 1989). Pickands's statistic is asymptotically normally distributed with mean τ and variance $\tau^2 \cdot\ (2^{-2\tau+1} + 1)/[2(2^{-\tau} - 1) \cdot \text{Log}\, 2]^2$. Pickands's estimator is the most general estimator since it can be used for all types of distributions.

Hill's estimator for the right tail is given by

$$\tau_{\text{Hill}} = \frac{1}{q-1} \sum_{i=1}^{q-1} \ln R'_{N-i} - \ln R'_{N-q}. \tag{8}$$

Hill's estimator can be used in the case of the Fréchet distribution only ($\tau > 0$). In this situation, Hill's estimator is consistent and the most efficient estimator. Consistency is still obtained under weak dependence in the parent variable R. Hill's statistic is asymptotically normally distributed with mean τ and variance τ^2.

In practice, as the database contains a finite number of return observations, the number of extreme returns, q, used for the estimation of the model is finite. As largely discussed in the extreme value theory literature, the choice of its value is a critical issue (see Danielsson et al. (2001) and Huisman et al. (2001) for a discussion). On the one hand, choosing a high value for q leads to few observations of extreme returns and implies inefficient parameter estimates with large standard errors. On the other hand, choosing a low value for q leads to many observations of extreme returns but induces biased parameter estimates as observations not belonging to the tails are included in the estimation process. To optimize this trade-off between bias and inefficiency, I use a Monte Carlo simulation method inspired by Jansen and De Vries (1991). Return time-series are simulated from a known distribution for which the tail index can be computed. For each time-series, the tail index value is estimated with a different number of extreme returns. The choice of the optimal value is based on the mean square error (MSE) criterion which allows one to take into account the trade-off between bias and inefficiency. The procedure is detailed in Appendix A.

3. Application of extreme value theory to discriminate among the distributions of returns

This section reviews the different models for the distribution of returns and shows how extreme value theory can be used to discriminate among these models by focusing on the distribution tails.

3.1. Distributions of returns

Several distributions for stock returns have been proposed in the financial literature. Most of the empirical works in finance assume that continuously compounded rates of return on common stock or on a portfolio are normally distributed with a constant variance. The Gaussian distribution is consistent with the log-normal diffusion model made popular by the Black–Scholes–Merton option pricing formula. Moreover, most of the statistical tests lie on the hypothesis of normality. Unfortunately, there is now strong evidence that the distribution of the stock returns departs from normality. High kurtosis usually found in the data implies that the distribution is leptokurtic. The empirical distribution is fat-tailed; there are more extreme observations than predicted by the normal model. This is of great importance because the tails of the density function partly determine the level of the volatility. And volatility is certainly a most important variable in finance.

I review below the alternative models to the Gaussian distribution, and show how these models can be discriminated using the tail index.

Mandelbrot (1963) first suggests that the variance of certain speculative price returns could not exist. Studying cotton prices, he concludes that the stable Paretian distributions fitted the data better than the Gaussian distribution did. Fama (1965) extend this approach to stock market prices.

If stock returns usually present fat tails, this does not imply that the variance is infinite. The mixture of Gaussian distributions and the unconditional Student-*t* distributions present an excess of kurtosis but still possess finite variance. Such models are proposed for stock prices by Praetz (1972) and Press (1967). A mixed distribution models the heterogeneity of the random phenomenon. The returns are drawn from different Gaussian distributions. Such a model has been used to take into account extreme price movements such as stock market crashes, which do not fit in model with a single distribution. Such events are assumed to be drawn from a distribution with a negative mean and high variance. Anomalies in the stock market like the "day effect" can also motivate this model.

The volatility varies in fact much more over time. Mandelbrot (1963) first finds a "clustering effect" in volatility, and he points out that large changes in prices tend to be followed by large changes of either sign, and similarly that small changes tend to be followed by small changes of either sign. The ARCH process proposed by Engle (1982) models this feature and tends to fit quite well the behavior of volatility.

3.2. Test based on extreme value theory

An extreme value investigation allows one to discriminate among these non-nested models. Although all processes of returns lead to the same form of distribution of extreme returns, the values of the parameters of the distribution of extremes are in general different for two different processes. Especially, the value of the tail index τ allows to discriminate these processes. A tail index value equal to 0 implies a Gumbel distribution obtained for thin-tailed distributions of returns. A negative value for the tail index implies a Weibul distribution obtained for distributions of returns with finite tails. A positive value for the tail index implies a Fréchet distribution obtained for fat-tailed distributions of returns. More precisely, a value of τ greater than 0.5 is consistent with the stable Paretian distribution. The Cauchy distribution corresponds to the special case $\tau = 1$. A value of τ less than 0.5 is consistent with the ARCH process or Student's distribution. An interesting feature of the tail index is that it is related to the highest existing moment of the distribution. The tail index τ and the highest existing moment denoted by k are simply related by: $k = 1/\tau$ (for τ positive). When τ is equal to 0, then all moments are defined ($k = +\infty$). This is the case of the Gaussian distribution and of the mixture of Gaussian distributions. For the stable Paretian distribution, k is lower than 2 (the variance is not defined) and equal to the characteristic exponent. For the Student-t distributions, k is greater than 2 and equal to the number of degree of freedom. Table 1 summarizes these results.

The tail index provides us with a straightforward test. Two particular unconditional distributions are considered below: the thin-tailed Gaussian distribution and the fat-tailed stable Paretian distribution.

Table 1
Tail index and highest existing moment for different models for returns

Models of returns	Type	Tail index τ	Highest existing moment k
Gaussian distribution	Gumbel	$\tau = 0$	$k = +\infty$
Mixture of Gaussian distributions	Gumbel	$\tau = 0$	$k = +\infty$
Stable Paretian distributions	Fréchet	$\tau > 0.5$	$k < 2$
Student-t distributions	Fréchet	$0 < \tau < 0.5$	$k \geqslant 2$
ARCH processes	Fréchet	$0 < \tau < 0.5$	$k \geqslant 2$

Note: The type of extreme value distribution, the tail index value and the highest existing moment for different models of returns commonly used in financial modeling are given. The tail index τ and the highest existing moment k are related by $k = 1/\tau$. The last two columns indicate the constraints on the coefficients τ and k imposed by each model.

3.2.1. The Gaussian distribution

As the Gaussian distribution for returns implies a Gumbel distribution for extreme returns, the tail index can be used for testing normality. The null hypothesis is stated as

$$H_0 : \tau = 0.$$

If the tail index τ is significantly different from 0, then the asymptotic distribution of extreme returns is not a Gumbel distribution. As a consequence, the Gaussian distribution for returns can be rejected. Alternatively, if the tail index is not different from 0, then the asymptotic distribution is the Gumbel distribution. Such a result is not inconsistent with the normal model.

3.2.2. The stable Paretian distribution

As the Paretian distribution for returns implies a Fréchet distribution for extreme returns (with a constraint on the tail index value greater than 0.5), the tail index can also be used for testing the Paretian model. The null hypothesis is stated as

$$H_0 : \tau > 0.5.$$

If the tail index τ is significantly lower than 0.5, then the asymptotic distribution of extreme returns is not a Fréchet one with high tail index value. As a consequence, the stable Paretian distribution for returns can be rejected. Alternatively, if the tail index τ is not significantly lower than 0.5, then the asymptotic distribution is the Fréchet distribution with high tail index value. Such a result is not inconsistent with the stable Paretian model.

4. Empirical results

4.1. Data

I use logarithmic daily percentage returns of the S&P500 index based on closing prices. Data are obtained from www.economagic.com. The database covers the period January 1954–December 2003 and contains 12.587 observations of daily returns.

The daily returns have a slightly positive mean (0.030%) and a high standard deviation (0.853). The values of the skewness (−1.34) and the excess kurtosis (35.20) suggest departure from the Gaussian distribution. The first order auto-correlation (generally attributed to a non-trading effect) is small (0.078) but significantly positive. Little serial correlation is found at higher lags. For the second moment, I find a strong positive serial correlation: 0.123 at lag 1. The correlation decreases slowly and remains significant even with a lag of 20 days (0.053), which suggests a strong persistence in volatility.

I now give some statistics about the extremes. Let us first consider the definition of extremes as minimal and maximal returns selected over a given time-period. Considering yearly extremes, I get 50 observations for each type of extreme over the period January 1954–December 2003. The top 10 yearly largest daily market falls and

Table 2
Top 10 yearly minimal and maximal daily returns in the S&P500 index

Yearly largest daily market falls		Yearly largest daily market rises	
October 19 1987	−22.90	October 21 1987	8.71
October 27 1997	−7.11	July 24 2002	5.57
August 31 1998	−7.04	October 28 1997	4.99
January 8 1988	−7.01	September 8 1998	4.96
May 28 1962	−6.91	May 27 1970	4.90
September 26 1955	−6.85	January 3 2001	4.89
October 13 1989	−6.31	March 16 2000	4.65
April 14 2000	−6.00	August 17 1982	4.65
September 17 2001	−5.05	May 29 1962	4.54
September 11 1986	−4.93	October 9 1974	4.49

Note: The 10 lowest yearly minimal daily returns and the 10 highest yearly maximal daily returns in the S&P500 index over the period January 1954–December 2003 are given. Yearly extreme returns are selected over non-overlapping years (containing 260 trading days on average).

market rises are reported in Table 2. Both types of extreme are widely spread. For the largest declines, the minimal value (−22.90%) is obtained in October 1987 and the second and third minimal values (−7.11% and −7.04%) during the Asian crisis in 1997 and the Russian crisis in 1998. The lowest yearly minimal daily returns (−1.33%) is observed in 1972. For the largest rises, the maximal value (+8.71%) is observed in October 1987 a few days after the market crash. Let us then consider the definition of extremes as negative and positive return exceedances under or above a given threshold. The top 10 largest daily market falls and market rises are reported in Table 3. As expected, the two definitions of extremes lead to similar sets of extreme observations. However, due to some clustering effect, extreme returns tend to appear around the same time. [5] This effect seems limited to the stock market crash of October 1987. Among the top 10 largest daily market falls, this event appears three times: October 16 (−5.30%), October 19 (−22.90%) and October 26 (−8.64%). The same remark applies to top 10 largest daily market rises. The period of extreme volatility following the crash on October 19 contains three top positive return exceedances: October 21 (+8.71%), October 20 (+5.20%) and October 29 (+4.81%).

4.2. Tail index estimates

The approaches described in Section 2 are now used to estimate the tail index. The empirical results are reported in Table 4 for parametric estimates using minimal and maximal returns, Table 5 for parametric estimates using negative and positive return exceedances and Table 6 for non-parametric estimates.

Let us begin to analyze the results for each estimation method as the tail index value tend to vary according to the method used and also to the parameter used

[5] Note that this effect can be estimated by incorporating another parameter in the extreme value distribution called the extremal index (see Longin, 2000).

Table 3
Top 10 negative and positive return exceedances in the S&P500 index

Largest daily market falls		Largest daily market rises	
October 19 1987	−22.90	October 21 1987	8.71
October 26 1987	−8.64	July 24 2002	5.57
October 27 1997	−7.11	July 29 2002	5.27
August 31 1998	−7.04	October 20 1987	5.20
January 8 1988	−7.01	October 28 1997	4.99
May 28 1962	−6.91	September 8 1998	4.96
September 26 1955	−6.85	May 27 1970	4.90
October 13 1989	−6.31	January 3 2001	4.89
April 14 2000	−6.00	October 29 1987	4.81
October 16 1987	−5.30	August 17 1982	4.65

Note: The 10 lowest negative daily return exceedances and the 10 highest positive daily return exceedances observed over the period January 1954–December 2003 are given. The threshold level is low enough to obtain at least 10 return exceedances (say ±4%).

Table 4
Parametric estimates of the tail index using minimal and maximal returns

Length of the selection period	Tail index estimate	
	Left tail	Right tail
One month	0.209	0.205
(600)	(0.031)	(0.037)
	[26.776] {1.000}	[26.675] {1.000}
One quarter	0.244	0.145
(200)	(0.048)	(0.061)
	[0.782] {0.783}	[1.201] {0.885}
One semester	0.452	0.076
(100)	(0.100)	(0.078)
	[1.221] {0.889}	[1.576] {0.942}
One year	0.511	0.164
(50)	(0.153)	(0.140)
	[0.660] {0.745}	[1.321] {0.907}

Note: The tail index estimates using minimal and maximal returns observed over a given time-period are given. Minimal and maximal returns are selected over period of different length: from one month to one year. The number of minimal or maximal returns used in the estimation process is given below in parentheses. The parameters of the distributions of minimal and maximal returns are estimated by the maximum likelihood method (only the tail index estimates are reported). Asymptotic standard errors are given below in parentheses. The result of Sherman's goodness-of-fit test is given in brackets with the *p*-value (probability of exceeding the test-value) given next in curly brackets. The 5% confidence level at which the null hypothesis of adequacy (of the estimated asymptotic distribution of extreme returns to the empirical distribution of observed extreme returns) can be rejected, is equal to 1.645.

to implement a particular method (i.e. the length of the selection period, the threshold value and the number of tail observations). Let us consider the left tail for example. For the left tail, the tail index estimate varies between 0.209 and 0.511 for the parametric method using minimal returns observed a given period, from 0.137 to 0.742 for the parametric method using negative return exceedances under a given

Table 5
Parametric estimates of the tail index using negative and positive return exceedances

Threshold used to select exceedances	Tail index estimate	
	Left tail	Right tail
±1%	0.137	0.083
(1.266) (1.178)	(0.032)	(0.026)
	[68.329] {1.000}	[63.900] {1.000}
± 2%	0.331	0.121
(238) (239)	(0.085)	(0.065)
	[11.671] {1.000}	[10.369] {1.000}
± 3%	0.742	−0.007
(55) (64)	(0.235)	(0.124)
	[1.075] {0.859}	[1.776] {0.962}
± 4%	0.367	0.051
(19) (19)	(0.314)	(0.241)
	[0.021] {0.508}	[0.174] {0.569}

Note: The tail index estimates using negative and positive return exceedances under or above a given threshold is given. Return exceedances are selected with different threshold values: from ±1% to ±4%. The number of return exceedances used in the estimation process is given below in parentheses for both negative and positive return exceedances. The parameters of the distributions of negative and positive return exceedances are estimated by the maximum likelihood method (only the tail index estimates are reported). Asymptotic standard errors are given below in parentheses. The result of Sherman's goodness-of-fit test is given in brackets with the *p*-value (probability of exceeding the test-value) given next in curly brackets. The 5% confidence level at which the null hypothesis of adequacy (of the estimated asymptotic distribution of extreme returns to the empirical distribution of observed extreme returns) can be rejected, is equal to 1.645.

Table 6
Non-parametric estimates of the tail index

Estimator	Tail index estimate	
	Left tail	Right tail
Pickands	0.178	(0.127)
(203) (203)	(0.120)	0.000
Hill	0.294	(0.030)
(75) (75)	(0.034)	0.263

Note: The tail index estimates based on non-parametric methods developed by Pickands (1975) and Hill (1975) are given. For each method the optimal number of tail observations is computed by simulation (see Appendix A). It is given in parentheses for both the left and right tails below the method name in the first column. Asymptotic standard errors of the tail index estimates are given below in parentheses.

threshold, and from 0.178 to 0.294 for the non-parametric methods. As the parametric approach assumes that the asymptotic distribution holds for finite samples, it is important to check the goodness-of-fit of the distribution to empirical data. For minimal returns, the Sherman test (reported in Table 4) shows that it seems cautious to select the extremes over a period longer than a semester. Similarly, for negative return exceedances, the Sherman test (reported in Table 5) shows that it seems cautious to select extremes under a threshold value lower than −3%. Looking at the non-paramertic approach, Pickands's estimate is positive suggesting that Hill's estimator can be used as it is restrained to the case of a positive tail index. Under this assumption, Hill's estimator is more precise than Pickands's estimator: the standard error of Hill's estimate is almost four times lower than the one of Pickands's estimate (see Table 6). For the remaining of the study, I will retain the values, which are the most relevant from a statistical point of view. For the left tail: 0.511 (parametric method based on minimal returns), 0.367 (parametric method based on negative return exceedances) and 0.294 (non-parametric Hill method). For the right tail: 0.164 (parametric method based on maximal returns), 0.051 (parametric method based on positive return exceedances) and 0.263 (non-parametric Hill method).

The first result is about the sign of the tail index, which determines the type of extreme value distribution. All tail index estimates (except one) are positive implying that the distribution of extreme returns is a Fréchet distribution consistent with fat-tailed distribution of returns.

The second result is about the relative asymmetry between the left tail and the right tail. The tail index estimates for the left tail are systematically higher than the one for the right tail. This statement can be formalized by testing the null hypothesis H_0: $\tau^{max} = \tau^{min}$. For usual confidence level (say 5%) this hypothesis is sometimes rejected by the data indicating that the left tail is heavier than the right tail.

4.3. *Choice of a distribution of stock market returns*

Two particular unconditional distributions are considered: the Gaussian distribution and the stable Paretian distribution by testing respectively the null hypotheses H_0: $\tau = 0$ and H_0: $\tau > 0.5$. Empirical results are reported in Table 7. Three confidence level are considered: 1%, 5% and 10%. The lowest the confidence level, the hardest to reject the null hypothesis.

4.3.1. *The Gaussian distribution*

Although the tail index estimates are always different from zero, they may not be significantly different from zero. Results reported in Table 7 show that the null hypothesis is often rejected even at conservative confidence levels such as 1%. The Gumbel distribution for extreme returns consistent with thin-tailed distributions for returns is then rejected. As the Gaussian distribution for returns implies a Gumbel distribution for extreme returns, this leads to the rejection of the Gaussian distribution.

4.3.2. *The stable Paretian distribution*

Although the tail index estimates are always lower than 0.5 (though positive), the null hypothesis H_0: $\tau > 0.5$ may not be significantly rejected. Results reported in Table 7 show that the null hypothesis is often rejected even at conservative confidence level such as 1%. The Fréchet distribution for extreme returns with a tail index value higher than 0.5 consistent with heavy-tailed distributions for returns is then rejected. As the stable Paretian distribution for returns implies a Fréchet distribution

Table 7
Choice of the distribution of returns based on the tail index

Estimator	Test of the null hypothesis	
	Left tail	Right tail
Panel A: The Gaussian distribution		
Parametric (ML) (minimal and maximal returns)	1%: rejected	1%: not rejected
	5%: rejected	5%: not rejected
	10%: rejected	10%: not rejected
Parametric (ML) (return exceedances)	1%: not rejected	1%: not rejected
	5%: not rejected	5%: not rejected
	10%: not rejected	10%: not rejected
Non-parametric Hill	1%: rejected	1%: rejected
	5%: rejected	5%: rejected
	10%: rejected	10%: rejected
Panel B: The stable Paretian distribution		
Parametric (ML) (minimal and maximal returns)	1%: not rejected	1%: not rejected
	5%: not rejected	5%: rejected
	10%: not rejected	10%: rejected
Parametric (ML) (return exceedances)	1%: not rejected	1%: rejected
	5%: not rejected	5%: rejected
	10%: not rejected	10%: rejected
Non-parametric Hill	1%: rejected	1%: rejected
	5%: rejected	5%: rejected
	10%: rejected	10%: rejected

Note: The result of the choice of a particular distribution for returns based on the tail index are given. Two particular distributions are considered: the Gaussian distribution characterized with a tail index value equal to 0 (Panel A) and the stable Paretian distribution characterized with a tail index value higher than 0.5 (Panel B). For the Gaussian distribution the null hypothesis is H_0: $\tau = 0$. For the stable Paretian distribution the null hypothesis is H_0: $\tau > 0.5$. Three estimators are used: the parametric maximum likelihood (ML) estimator based on minimal and maximal returns and return exceedances and the Hill non-parametric estimator. Three confidence levels are considered: 1%, 5% and 10%.

for extreme returns with a tail index value higher than 0.5, this leads to the rejection of the stable Paretian distribution.

Both the Gaussian distribution and the stable Paretian distribution seem rejected by the data. In terms of moments, the variance appears to be defined although not all moment are defined. The highest existing moment is determined next.

4.4. Highest existing moment

The tail index can be used to compute the highest defined moment of the distribution of returns. Technically, it corresponds to the highest integer k such that $E(R^k)$ is finite. I proceed as follows: I consider a set of null hypotheses $H_0(k)$ defined by: $\tau < 1/k$. If the null hypothesis $H_0(k)$ is rejected at a given confidence level, then the moment of order k is not defined at this level. The null hypothesis $H_0(+\infty)$ defined by $\tau \leqslant 0$ serves as a limiting case. If the null hypothesis $H_0(+\infty)$ is rejected, then not all moments are defined.

Table 8
Maximal existing moment of the distribution of the S&P500 index returns

Estimator	Maximal existing moment	
	Left tail	Right tail
Parametric (ML) (minimal and maximal returns)	1%: sixth	1%: all
	5%: third	5%: all
	10%: third	10%: all
Parametric (ML) (return exceedances)	1%: 10th	1%: all
	5%: seventh	5%: all
	10%: sixth	10%: all
Non-parametric Hill	1%: fourth	1%: fifth
	5%: fourth	5%: fourth
	10%: third	10%: fourth

Note: The highest existing moment of the distribution of stock market returns by investigating the weight of extreme price movements. is given. For a given level of confidence, equal to 1%, 5% and 10%, the null hypotheses $H_0(k)$ defined by: $\tau > 1/k$ where k is equal to $1, 2, 3, \ldots$ is studied. A t-test and its associated p-value are computed. The null hypothesis $H_0(+\infty)$: $\tau < 0$ serves as the limiting case. The highest integer k for which $H_0(k)$ is not rejected at the given level, is reported. If the null hypothesis $H_0(+\infty)$ is not rejected, then all the moments are defined. Three estimators are used: the parametric maximum likelihood (ML) estimator based on minimal and maximal returns and return exceedances and the Hill non-parametric estimator. Three confidence levels are considered: 1%, 5% and 10%.

Table 8 gives the empirical results concerning the highest existing moment by looking at each tail independently. Three confidence levels are considered: 1%, 5% and 10%. The lowest the confidence level, the easiest to accept the existence of lower moments. As expected, the conclusion of the test depends on the method used for estimating the tail index. However, general results emerged. The first result is that the second moment (the variance) seems to be always defined as the null hypothesis $H_0(2)$ is never rejected. The fourth moment seems however not always defined. The second result is the relative asymmetry between the left tail and the right tail. The highest existing moment by considering the left tail is always lower than the highest existing moment by considering the right tail suggesting that the left tail is heavier than the right tail. Moreover, by looking at the right tail, all moments seem defined in most of tests.

5. Conclusion

Extreme value theory gives a simple way to discriminate among the distributions of returns. The distributions commonly proposed in the literature can be differentiated by the tails or in other words by the frequency of extreme price movements.

Empirical results for the US stock market lead to the rejection of the Gaussian distribution and the stable Paretian distributions as well. The former contains too few extremes while the later too many. Although the distribution of stock market returns is fat-tailed, the variance appears to be well-defined. Only the Student-t distribution and the class of ARCH processes are not rejected by the data. This suggests that for the US stock market, a Student-t distribution could be used in a unconditional modeling of returns and that an ARCH process could be used in a conditional modeling of returns.

Appendix A. Computation of the optimal value q for non-parametric estimators

I compute the optimal value of q by carrying out a Monte Carlo study as done by Jansen and De Vries (1991). I proceed as follows: I simulate 12.587 return observations (the total number of daily returns in the database) drawn from different return distributions: a Cauchy distribution and Student-t distributions with degrees of freedom equal to 2, 3 and 4. The fatness of these four distributions is different and corresponds to tail indices τ equal to 1, 0.5, 0.33 and 0.25. The Cauchy distribution gives a lot of extreme values while a Student-t distribution with four degrees of freedom very few. Then I estimate the tail index using Pickands's or Hill's formula with different values of q ranging from 1 to 2.500 (about 20% of the observations). I repeat this simulation 10.000 times. For each distribution i (chacharterized by a tail index value τ_i) and each value of q, I get a series of 10.000 observations of the tail index estimates. Then for each distribution i, I compute the mean square error (MSE) of this series and I choose the value of q, written q_i^{opt}, which minimizes the MSE. As explained by Theil (1971, pp. 26–32) the MSE criterion allows one to take explicitly into account the two effects of bias and inefficiency. The mean square error of S simulated observations $\widetilde{X}_s$ of the estimator of a parameter X can be decomposed as follows:

$$\mathrm{MSE}((\widetilde{X}_s)_{s=1,S'}X) = (\overline{X} - X)^2 + \frac{1}{S}\sum_{s=1}^{S}(\widetilde{X}_s - X)^2,$$

where $\overline{X}$ represents the mean of S simulated observations. The first part of the decomposition measures the bias and the second part the inefficiency.

Table 9 reports the minimizing q-levels and associated MSEs using Hill's estimator. Along the diagonal are the minimal MSEs; the theoretical MSE value equal to τ^2/q is also reported. As noted by Jansen and De Vries (1991), there is a U-shaped relationship between MSE and q. It reflects the trade-off between inefficiency and bias: when few observations are used (q low), the bias in the estimation of τ is negligible as most of the observations are extreme but the variance of the estimator is high; when a lot of observations are used (q high), a bias is introduced in the estimation of τ because of the inclusion of more central values but the variance of the estimator is low.

With real data I proceed as follows: I compute tail index estimates with the four optimal values previously obtained given as $(q_i^{\text{opt}})_{i=1,4}$ These values correspond to the four chosen values of the tail index given as $(\tau_i)_i)_{i=1,4}$. I retain the estimate that is the closest to the chosen value τ_i. To do this I compute the statistics $(\tau^{\text{Hill}}(q_i^{\text{opt}}) - \tau_i)/\sigma_i$ where $\tau^{\text{Hill}}(q_i^{\text{opt}})$ is the Hill's estimate computed with q_i^{opt} extremes and σ_i is the standard error of this estimate, and the associated p-value noted p_i. I finally retain the

Table 9
Optimal value for Hill's estimator of the tail index

	$k = 1$ ($\tau = 1.00$)	$k = 2$ ($\tau = 0.50$)	$k = 3$ ($\tau = 0.33$)	$k = 4$ ($\tau = 0.25$)
$q = 1.084$ $k = 1$ ($\tau = 1.00$)	0.0010 [0.0009]	0.0054	0.0142	0.0214
$q = 329$ $k = 2$ ($\tau = 0.50$)	0.0031	0.0009 [0.0008]	0.0019	0.0039
$q = 153$ $k = 3$ ($\tau = 0.33$)	0.0071	0.0016	0.0010 [0.0007]	0.0015
$q = 77$ $k = 4$ ($\tau = 0.25$)	0.0146	0.0034	0.0014	0.0011 [0.0007]

Note: The mean squared error (MSE) and the theoretical MSE in brackets obtained from simulations for different values of q used to compute Hill's estimate and for different values of the degrees of freedom α (or the tail index τ) are indicated. The case $k = 1$ corresponds to a Cauchy distribution and $k = 2$, 3 and 4 to Student-t distributions. This corresponds to tail index values respectively equal to 1.00, 0.50, 0.33 and 0.25. The whole period is assumed to contain 12.587 observations. Values of q minimizing the MSE are 1.084 for $\tau = 1.00$, 329 for $\tau = 0.50$, 153 for $\tau = 0.33$ and 77 for $\tau = 0.25$. Minimizing MSEs can be found in diagonal.

estimate for which the lowest value of p_i is obtained. In my study I keep 75 extreme returns to compute Hill's estimator (the same number for the left and right tails). Optimal value for Pickands's estimator is 203.

References

Balkema, A.A., De Haan, L., 1974. Residual life time at great age. Annals of Probability 2, 792–804.

Berman, S.M., 1963. Limiting theorems for the maximum term in stationary sequences. Annals of Mathematical Statistics 35, 502–516.

Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. Journal of Political Economy 81, 637–659.

Danielsson, J., De Haan, L., Peng, L., De Vries, C.G., 2001. Using a bootstrap method to choose the sample fraction in tail index estimation. Journal of Multivariate Analysis 76, 226–248.

Dekkers, A.L.M., De Haan, L., 1989. On the estimation of the extreme value index and large quantile estimation. The Annals of Statistics 17, 1795–1832.

De Haan, L., Resnick, I.S., Rootzèn, H., De Vries, C.G., 1989. Extremal behavior of solutions to a stochastic difference equation with applications to ARCH process. Stochastic Processes and Their Applications 32, 213–224.

Embrechts, P., Klüppelberg, C., Mikosch, T., 1997. Modelling Extremal Events for Insurance and Finance. Springer-Verlag, Berlin.

Engle, R.F., 1982. Auto-regressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation. Econometrica 50, 987–1007.

Fama, E.F., 1965. The behavior of stock market prices. The Journal of Business 38, 34–105.

Galambos, J., 1978. The Asymptotic Theory of Extreme Order Statistics. John Wiley and Sons.

Gnedenko, B.V., 1943. Sur la distribution limite du terme maximum d'une série aléatoire. Annals of Mathematics 44, 423–453.

Gumbel, E.J., 1958. Statistics of Extremes. Columbia University Press, New York.

Hill, B.M., 1975. A simple general approach to inference about the tail of a distribution. Annals of Statistics 46, 1163–1173.

Huisman, R., Koedijk, K., Kool, C.J.M., Palm, F., 2001. Tail index estimates in small samples. Journal of Business and Economic Statistics 19, 208–215.
Jansen, D.W., De Vries, C.G., 1991. On the frequency of large stock returns: putting booms and busts into perspectives. The Review of Economic and Statistics 73, 18–24.
Leadbetter, M.R., Lindgren, G., Rootzèn, H., 1983. Extremes and Related Properties of Random Sequences and Processes. Springer-Verlag, New York.
Longin, F., Volatility and extreme movements in fianncial markets. Ph.D. Thesis, HEC, 1993.
Longin, F., 1996. The asymptotic distribution of extreme stock market returns. Journal of Business 63, 383–408.
Longin, F., 2000. From VaR to stress testing: The extreme value approach. Journal of Banking and Finance 24, 1097–1130.
Longin, F., Solnik, B., 2001. Extreme correlation of international equity markets. Journal of Finance 56, 651–678.
Loretan, M., Phillips, P.C.B., 1994. Testing for covariance stationarity of heavy-tailed time series: An overview of the theory with applications to several financial datasets. Journal of Empirical Finance 2, 211–248.
Mandelbrot, B., 1963. The variation of certain speculative prices. Journal of Business 36, 394–419.
Markowitz, H., 1952. Portfolio selection. Journal of Finance 7, 77–91.
Merton, R., 1973. Theory of rational option pricing. Bell Journal of Economics Management Science 4, 141–183.
Pickands, J., 1975. Statistical inference using extreme order statistics. Annals of Statistics 3, 119–131.
Praetz, P.D., 1972. The distribution of share price changes. Journal of Business 45, 49–55.
Press, S.J., 1967. A compound events model for security prices. Journal of Business 40, 317–335.
Reiss, R.-D., Thomas, M., 1997. Statistical Analysis of Extreme Values. Birkhäuser Verlag, Basle.
Sharpe, W., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. Journal of Finance 19, 425–442.
Theil, H., 1971. Applied Economic Forecasting. North-Holland, Amsterdam.