Statistical Tools in Finance

Pavel Čížek, Wolfgang Härdle, Rafael Weron

August 6, 2003

Contents

Part I

Finance

1 Stable distributions in finance

Szymon Borak, Wolfgang Härdle, Rafał Weron

1.1 Introduction

Stable laws – also called α-stable or Levy-stable – are a rich family of probability distributions that allow skewness and heavy tails and have many interesting mathematical properties. They appear in the context of the Generalized Central Limit Theorem which states that the only possible non-trivial limit of normalized sums of independent identically distributed variables is α-stable. The Standard Central Limit Theorem states that the limit of normalized sums of independent identically distributed terms with finite variance is Gaussian (α-stable with $\alpha = 2$).

It is often argued that financial asset returns are the cumulative outcome of a vast number of pieces of information and individual decisions arriving almost continuously in time (McCulloch, 1996; Rachev and Mittnik, 2000). Hence, it is natural to consider stable distributions as approximations. The Gaussian law is by far the most well known and analytically tractable stable distribution and for these and practical reasons it has been routinely postulated to govern asset returns. However, financial asset returns are usually much more leptokurtic, i.e. have much heavier tails. This leads to considering the non-Gaussian ($\alpha < 2$) stable laws, as first postulated by Benoit Mandelbrot in the early 60s (Mandelbrot, 1997).

Apart from empirical findings, in some cases there are solid theoretical reasons for expecting a non-Gaussian α-stable model. For example, emission of particles from a point source of radiation yields a Cauchy distribution ($\alpha = 1$), hitting times for a Brownian motion yield a Levy distribution ($\alpha = 0.5, \beta = 1$), the gravitational field of stars yields the Holtsmark distribution ($\alpha = 1.5$), for a review see Janicki and Weron (1994) or Uchaikin and Zolotarev (1999).

1.2 α-stable distributions

Stable laws were introduced by Paul Levy during his investigations of the behavior of sums of independent random variables in (Levy, 1925). Note that a sum of two independent random variables having an α-stable distribution with index α is again α-stable with the same index α. However, this invariance property does not hold for different α's, i.e. a sum of two independent stable random variables with different α's is not α-stable. In fact, this property is fulfilled for a more general class – the class of infinitely divisible distributions, which are a limiting law for sums of independent (but not identically distributed) variables.

The α-stable distribution requires four parameters for complete description: an index of stability $\alpha \in (0, 2]$ also called the tail index, tail exponent or characteristic exponent, a skewness parameter $\beta \in [-1, 1]$, a scale parameter $\sigma > 0$ and a location parameter $\mu \in \mathbb{R}$. The tail exponent α determines the rate at which the tails of the distribution taper off, see Figure 1.1. When $\alpha = 2$, a Gaussian distribution results. When $\alpha < 2$, the variance is infinite. When $\alpha > 1$, the mean of the distribution exists and is equal to μ. In general, the pth moment of a stable random variable is finite if and only if $p < \alpha$. When the skewness parameter β is positive, the distribution is skewed to the right, i.e. the right tail is thicker, see Figure 1.2. When it is negative, it is skewed to the left. When $\beta = 0$, the distribution is symmetric about μ. As α approaches 2, β loses its effect and the distribution approaches the Gaussian distribution regardless of β. The last two parameters, σ and μ, are the usual scale and location parameters, i.e. σ determines the width and μ the shift of the mode (the peak) of the distribution.

1.2.1 Characteristic function representation

Due to the lack of closed form formulas for densities for all but three distributions (see Figure 1.3), the α-stable law can be most conveniently described by its characteristic function $\phi(t)$ – the inverse Fourier transform of the probability density function. However, there are multiple parameterizations for α-stable laws and much confusion has been caused by these different representations. The variety of formulas is caused by a combination of historical evolution and the numerous problems that have been analyzed using specialized forms of the stable distributions. The most popular parameterization of the characteristic function of $X \sim S_\alpha(\sigma, \beta, \mu)$, i.e. an α-stable random variable with parameters

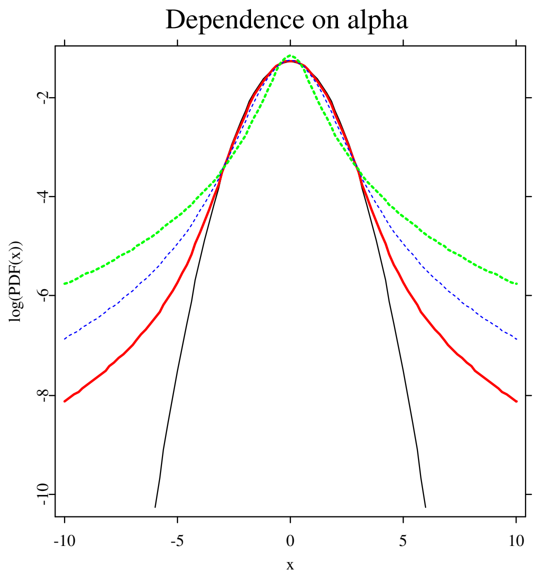

Figure 1.1: A semilog plot of symmetric ($\beta = \mu = 0$) α-stable probability density functions for $\alpha = 2$ (thin black), 1.8 (red), 1.5 (thin, dashed blue) and 1 (dashed green). Observe that the Gaussian ($\alpha = 2$) density forms a parabola and is the only α-stable density with exponential tails.

XCSstab01.xpl

α, σ, β and μ, is given by (Samorodnitsky and Taqqu, 1994; Weron, 1996):

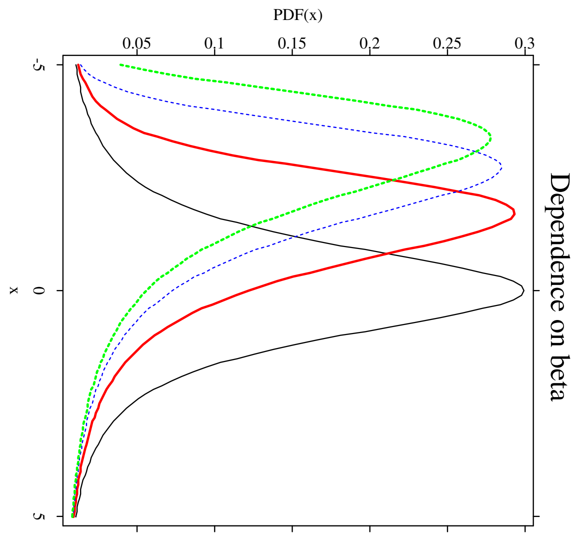

Figure 1.2: α-stable probability density functions for $\alpha = 1.2$ and $\beta = 0$ (thin black), 0.5 (red), 0.8 (thin, dashed blue) and 1 (dashed green).

XCSstab02.xpl

$$\log \phi(t) = \begin{cases} -\sigma^{\alpha}|t|^{\alpha}\{1 - \Im\beta \mathtt{sign}(t) \tan\frac{\pi\alpha}{2}\} + \Im\mu t, & \alpha \neq 1, \\ \\ -\sigma|t|\{1 + \Im\beta \mathtt{sign}(t) \frac{2}{\pi} \log|t|\} + \Im\mu t, & \alpha = 1. \end{cases} \tag{1.1}$$

For numerical purposes, it is often useful (Fofack and Nolan, 1999) to use a different parameterization:

$$\log \phi_0(t) = \begin{cases} -\sigma^{\alpha}|t|^{\alpha}\{1 + \Im\beta \mathtt{sign}(t) \tan\frac{\pi\alpha}{2}[(\sigma|t|)^{1-\alpha} - 1]\} + \Im\mu_0 t, & \alpha \neq 1, \\ \\ -\sigma|t|\{1 + \Im\beta \mathtt{sign}(t) \frac{2}{\pi} \log(\sigma|t|)\} + \Im\mu_0 t, & \alpha = 1. \end{cases} \tag{1.2}$$

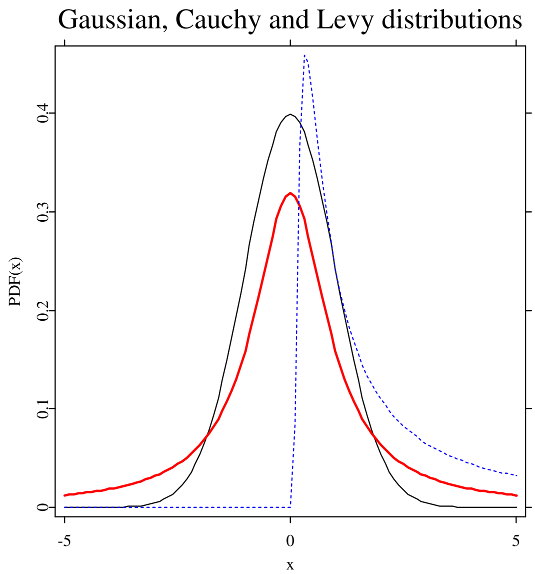

Figure 1.3: Closed form formulas for densities are known only for three distributions: Gaussian ($\alpha = 2$; thin black), Cauchy ($\alpha = 1$; red) and Levy ($\alpha = 0.5, \beta = 1$; thin, dashed blue). Note that the latter is a totally skewed distribution, i.e. its support is the positive real line only. In general, for $\alpha < 1$ and $\beta = 1$ (-1) the distribution is totally skewed to the right (left).

XCSstab03.xpl

The $S^0_\alpha(\sigma, \beta, \mu_0)$ parameterization is a variant of Zolotariev's (M)-parameterization (Zolotarev, 1986), with the characteristic function and hence the density and the distribution function jointly continuous in all four parameters, see Figure 1.4. In particular, percentiles and convergence to the power-law tail vary in a continuous way as α and β vary. The location parameters of the two representations are related by $\mu = \mu_0 - \beta\sigma \tan \frac{\pi\alpha}{2}$ for $\alpha \neq 1$ and $\mu = \mu_0 - \beta\sigma \frac{2}{\pi} \log \sigma$ for $\alpha = 1$.

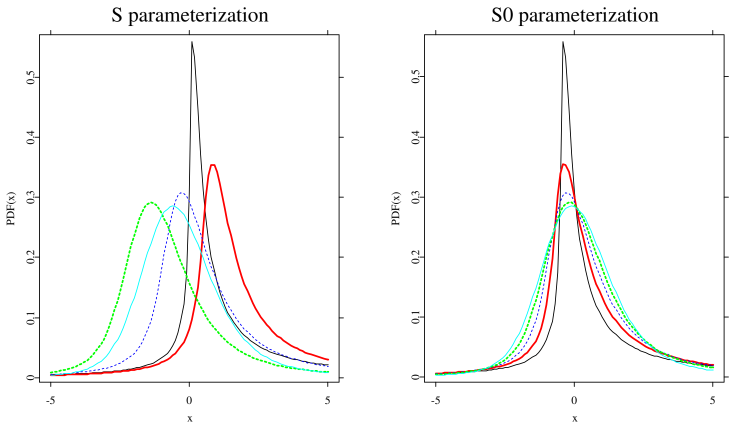

Figure 1.4: Comparison of S and S^0 parameterizations: α-stable probability density functions for $\beta = 0.5$ and $\alpha = 0.5$ (thin black), 0.75 (red), 1 (thin, dashed blue), 1.25 (dashed green) and 1.5 (thin cyan).

XCSstab04.xpl

1.2.2 Computing α-stable distributions

```
y = pdfstab(x,alpha,sigma,beta{,mu,{,n}})
```
computes the probability density function (pdf) of a generic stable distribution

```
y = pdfcauch(x,sigma,mu)
```
computes the pdf of the Cauchy distribution

```
y = pdflevy(x,sigma,mu)
```
computes the pdf of the Levy distribution

```
y = cdfstab(x,alpha,sigma,beta{,mu{,n}})
```
computes the cumulative distribution function (cdf) of a generic stable distribution

```
y = cdfcauch(x,sigma,mu)
```
computes the cdf of the Cauchy distribution

```
y = cdflevy(x,sigma,mu)
```
computes the cdf of the Levy distribution

The quantlet `pdfstab` computes the probability density function of values from `x` where `alpha`, `sigma`, `beta`, and `mu` (default `mu`=0) are parameters of the stable distribution and `n` (default `n`=2000) is the number of subintervals in internal integral calculations. Quantlets `pdfcauch` and `pdflevy` calculate values of the Cauchy and Levy distributions, respectively. `x` is the input array; `sigma` and `mu` are the scale and location parameters of these distributions.

The quantlet `cdfstab` computes the cumulative distribution function of values from `x` where `alpha`, `sigma`, `beta`, and `mu` (default `mu`=0) are parameters of the stable distribution and `n` (default `n`=2000) is the number of subintervals in internal integral calculations. Quantlets `cdfcauch` and `cdflevy` calculate values of the Cauchy and Levy distributions, respectively. `x` is the input array; `sigma` and `mu` are the scale and location parameters of these distributions.

Quantlets `pdfstab` and `cdfstab` utilize Nolan's (1997) integral formulas for the density and the cumulative distribution function. The larger the value of `n` the more accurate and time consuming (!) the numerical integration.

1.2.3 Simulation of α-stable variables

The complexity of the problem of simulating sequences of α-stable random variables results from the fact that there are no analytic expressions for the inverse F^{-1} of the cumulative distribution function. The first breakthrough was made by Kanter (1975), who gave a direct method for simulating $S_\alpha(1,1,0)$ random variables, for $\alpha < 1$. It turned out that this method could be easily adapted to the general case. Chambers, Mallows and Stuck (1976) were the first to give the formulas.

The algorithm for constructing a random variable $X \sim S_\alpha(1,\beta,0)$, in representation (1.1), is the following (Weron, 1996):

- generate a random variable V uniformly distributed on $(-\frac{\pi}{2},\frac{\pi}{2})$ and an independent exponential random variable W with mean 1;
- for $\alpha \neq 1$ compute:

$$X = S_{\alpha,\beta} \times \frac{\sin\alpha(V+B_{\alpha,\beta})}{\{\cos(V)\}^{1/\alpha}} \times \left[\frac{\cos V - \alpha(V+B_{\alpha,\beta})}{W}\right]^{(1-\alpha)/\alpha}, \quad (1.3)$$

  where

$$\begin{aligned} B_{\alpha,\beta} &= \frac{\arctan(\beta\tan\frac{\pi\alpha}{2})}{\alpha}, \\ S_{\alpha,\beta} &= \left\{1+\beta^2\tan^2\left(\frac{\pi\alpha}{2}\right)\right\}^{1/(2\alpha)}; \end{aligned}$$

- for $\alpha = 1$ compute:

$$X = \frac{2}{\pi}\left\{\left(\frac{\pi}{2}+\beta V\right)\tan V - \beta\log\left(\frac{\frac{\pi}{2}W\cos V}{\frac{\pi}{2}+\beta V}\right)\right\}. \quad (1.4)$$

Given the formulas for simulation of a standard α-stable random variable, we can easily simulate a stable random variable for all admissible values of the parameters α, σ, β and μ using the following property: if $X \sim S_\alpha(1,\beta,0)$ then

$$Y = \begin{cases} \sigma X + \mu, & \alpha \neq 1, \\ \\ \sigma X + \frac{2}{\pi}\beta\sigma \log \sigma + \mu, & \alpha = 1, \end{cases} \tag{1.5}$$

is $S_\alpha(\sigma, \beta, \mu)$. Although many other approaches have been presented in the literature, this method is regarded as the fastest and the most accurate.

`z = rndstab(alpha,sigma,beta,mu,d1{,...,dn})`
generates arrays up to eight dimensions of pseudo random variables with a stable distribution

`z = rndsstab(alpha,sigma,d1{,...,dn})`
generates arrays up to eight dimensions of pseudo random variables with a symmetric stable distribution

Quantlets `rndstab` and `rndsstab` use formulas (1.3)-(1.5) and provide pseudo random variables of stable and symmetric stable distributions, respectively. Parameters `alpha` and `sigma` in both quantlets and `beta` and `mu` in the first one determine the parameters of the stable distribution. `d1,...,dn` are responsible for the output dimensions.

1.2.4 Tail behavior

Levy (1925) has shown that when $\alpha < 2$ the tails of α-stable distributions are asymptotically equivalent to a Pareto law. Namely, if $X \sim S_{\alpha<2}(1, \beta, 0)$ then as $x \to \infty$:

$$\begin{aligned} P(X > x) = 1 - F(x) &\quad\to\quad C_\alpha(1+\beta)x^{-\alpha}, \\ P(X < -x) = F(-x) &\quad\to\quad C_\alpha(1-\beta)x^{-\alpha}, \end{aligned} \tag{1.6}$$

where

$$C_\alpha = \left(2\int_0^\infty x^{-\alpha} \sin x dx\right)^{-1} = \frac{1}{\pi}\Gamma(\alpha)\sin\frac{\pi\alpha}{2}.$$

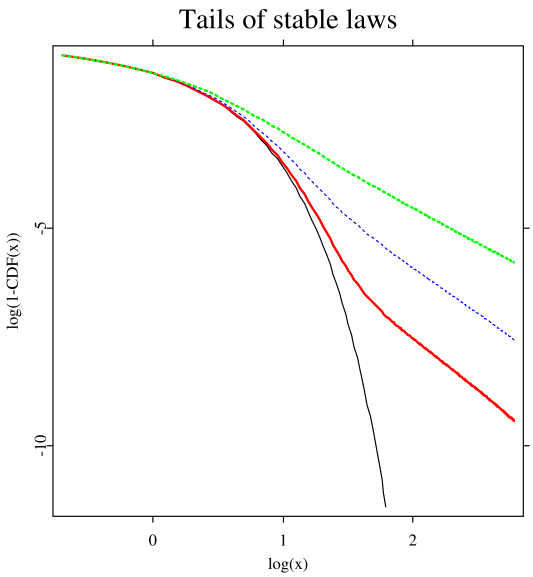

Figure 1.5: Right tails of symmetric α-stable distribution functions for $\alpha = 2$ (thin black), 1.95 (red), 1.8 (thin, dashed blue) and 1.5 (dashed green) on a double logarithmic paper. For $\alpha < 2$ the tails form straight lines with slope $-\alpha$.

XCSstab05.xpl

The convergence to a power-law tail varies for different α's (Mandelbrot, 1997, Chapter 14) and, as can be seen in Figure 1.5, is slower for larger values of the tail index. Moreover, the tails of α-stable distribution functions exhibit a crossover from an approximate power decay with exponent $\alpha > 2$ to the true tail with exponent α. This phenomenon is more visible for large α's (Weron, 2001).

1.3 Estimation of parameters

The estimation of stable law parameters is in general severely hampered by the lack of known closed–form density functions for all but a few members of the stable family. Most of the conventional methods in mathematical statistics, including the maximum likelihood estimation method, cannot be used directly in this case, since these methods depend on an explicit form for the density. However, there are numerical methods that have been found useful in practice and are discussed in this section.

All presented methods work quite well assuming that the sample under consideration is indeed α-stable. However, if the data comes from a different distribution, these procedures may mislead more than the Hill and direct tail estimation methods. Since there are no formal tests for assessing the α-stability of a data set we suggest to first apply the "visual inspection"or non-parametric tests to see whether the empirical densities resemble those of α-stable laws.

Given a sample $x_1, ..., x_n$ from $S_\alpha(\sigma, \beta, \mu)$, in what follows, we will provide estimates $\hat{\alpha}$, $\hat{\sigma}$, $\hat{\beta}$ and $\hat{\mu}$ of α, σ, β and μ, respectively.

1.3.1 Tail exponent estimation

The simplest and most straightforward method of estimating the tail index is to plot the right tail of the (empirical) cumulative distribution function (i.e. $1 - F(x)$) on a double logarithmic paper, as in Figure 1.6. The slope of the linear regression for large values of x yields the estimate of the tail index α, through the relation $\alpha = -\text{slope}$.

This method is very sensitive to the sample size and the choice of the number of observations used in the regression. Moreover, the slope around -3.7 may indicate a non-α-stable power-law decay in the tails or the contrary – an α-stable distribution with $\alpha \approx 1.9$. To illustrate this run the quantlet `XCSstab07`. First simulate (using equation (1.3) and quantlet `rndsstab`) samples of size $N = 10^4$ and 10^6 of standard symmetric ($\beta = \mu = 0$, $\sigma = 1$) α-stable distributed variables with $\alpha = 1.9$. Next, plot the right tails of the empirical distribution functions on a double logarithmic paper. The results are displayed in Table 1.1.

The true tail behavior (1.6) is observed only for very large (also for very small, i.e. the negative tail) observations, after a crossover from a temporary power-

Table 1.1: Tail index estimation

alpha = -slope
10^6 samples
-3.7881
-1.9309
10^4 samples
-3.7320

XCSstab07.xpl

like decay. Moreover, the obtained estimates still have a slight positive bias, which suggests that perhaps even larger samples than 10^6 observations should be used. In Figure 1.6 we used only the upper 0.15% of the records to estimate the true tail exponent. In general, the choice of observations used in the regression is subjective and can yield large estimation errors.

A well known method for estimating the tail index that does not assume a parametric form for the entire distribution function, but focuses only on the tail behavior was proposed by Hill (1975). The Hill estimator is used to estimate the tail index α, when the upper (or lower) tail of the distribution is of the form: $1 - F(x) = Cx^{-\alpha}$. Like the log-log regression method, the Hill estimator tends to overestimate the tail exponent of the stable distribution if α is close to two and the sample size is not very large, see Figure 1.7. For a review of the extreme value theory and the Hill estimator see Chapter 13 in Härdle, Klinke, and Müller (2000) or Embrechts, Klüppelberg and Mikosch (1997).

These examples clearly illustrate that the true tail behavior of α-stable laws is visible only for extremely large data sets. In practice, this means that in order to estimate α we must use high-frequency asset returns and restrict ourselves to the most "outlying" observations. Otherwise, inference of the tail index may be strongly misleading and rejection of the α-stable regime unfounded.

1.3.2 Sample Quantiles Methods

Let x_f be the f–th population quantile, so that $S_\alpha(\sigma, \beta, \mu)(x_f) = f$. Let $\hat{x}_f$ be the corresponding sample quantile, i.e. $\hat{x}_f$ satisfies $F_n(\hat{x}_f) = f$.