Using finite moment log-stable distributions for modelling financial risk

G.K. Robinson

Version 5 — September 2012

Abstract

This paper concentrates on the stable distributions which have maximum skewness. The exponentials of such stable distributions with maximum skewness to the left are called finite moment log-stable distributions. They have the property that all moments are finite, and are of interest in financial options pricing as an alternative to log-normal distributions.

One difficulty which has hampered the practical use of stable distributions has been that their density and distribution functions are difficult to compute. The main novelty of this paper is to show how this difficulty can be overcome by using interpolation formulae in two variables to compute the density and distribution functions rapidly and to high precision. Another difficulty is that it is easy to make mistakes when manipulating the parameters of stable distributions. This difficulty can be reduced by using objects which contain information about the distributions and performing computational procedures on those objects. Both of these developments have been implemented in an `R` package called `FMStable`.

Possible use of finite moment log-stable distributions for options pricing is discussed. The most important qualitative difference from the Black-Scholes log-normal model for options pricing is that dynamic hedging appears to reduce portfolio risk by a much smaller amount. This suggests that finite moment log-stable distributions could be used to provide more conservative assessments of portfolio risk.

1 Introduction

Stable distributions were originally defined as the solution to a theoretical problem. They have the property that the sum of several independent and identically distributed random variables has the same distribution as the individual random variables, except that the location and scale parameters may be different. According to Feller (1966, page 169), Lèvy (1924) first found the Fourier transforms of the family of stable distributions.

There are four parameters to the family of stable distributions: a location parameter, a scale parameter, a skewness parameter which is usually denoted by β and a parameter denoted by α which is called the characteristic exponent. If two independent random variables X and Y have identical stable probability distributions then the sum $X+Y$ has a stable probability distribution with a scale which is $2^{1/\alpha}$ times the scale of X.

For financial applications, I consider it useful to restrict attention to the maximally skew distributions (i.e. those for which β is either $+1$ or -1). Zolotarev (1986) used the adjective "extremal" to distinguish them. If X has a maximally skew stable distribution which is skewed to the right then $\exp(-X)$ can be said to have a log maximally skew stable distribution. This terminology is based on the fact that if X has a normal distribution then $\exp(X)$ is said to have a log-normal distribution. The log maximally skew stable distributions have also been called "finite moment log-stable distributions" by Carr and Wu (2003), highlighting their important property that all moments are finite.

Sections 2 and 3 of this paper give technical details of how to achieve good precision when translating between parametrizations, when computing values of the distribution function or probability density by numerical integration, and when when computing values of the distribution function or probability density by interpolation. These details need not be understood by a user concerned with financial applications, but further work on these aspects of stable distributions is required before computer software dealing with stable distributions can be relied upon. Section 4 discusses the use of these basic methodologies for calculating option prices. Section 5 discusses consequences.

Propagation of rounding errors

Some familiarity with numerical analysis is needed to understand parts of this paper. In particular, it must be appreciated that mathematically equivalent formulae sometimes provide quite different computational accuracy.

Computer systems generally provide facilities for computing common mathematical functions such as log, exp, sin, arcsin, cos, arccos, sinh and asinh to approximately the relative precision with which floating-point numbers are stored: about 1 part in 2^{52} or 2.2×10^{-16}. The functions log1p and expm1 are also available in most computer languages which are used for mathematical computing. In some circumstances where the argument of the function might sensibly be specified as a deviation from a standard value or the quantity required is the function plus or minus a constant, we need to be very conscious of the precision that might be lost when numbers are subtracted. Some examples of computational alternatives are listed below. In all cases, δ denotes a small positive number such as $\delta = 1 \times 10^{-13}$, which is known to good relative precision.

- $\sin(\delta)$ is better than $\sin(\pi - \delta)$.
- $2\arcsin(\sqrt{0.5\delta})$ is better than $\arccos(1 - \delta)$.
- $\operatorname{asinh}(\sqrt{\delta(2 - \delta)})$ is better than $\operatorname{acosh}(1 + \delta)$.
- $\operatorname{log1p}(\delta)$ is better than $\log(1 + \delta)$.
- $\operatorname{expm1}(\delta)$ is better than $\exp(\delta) - 1$.
- $2\sin^2(0.5\delta)$ is better than $1 - \cos(\delta)$.
- $2\sinh^2(0.5\delta)$ is better than $\cosh(\delta) - 1$.

One example of sensitivity to deviations from a standard value involving stable distributions is that the right-hand-tail probability for a maximally-skew stable distribution is asymptotically proportional to $2c_\alpha x^{-\alpha}$ where $c_\alpha = \Gamma(\alpha)\sin(\frac{\pi}{2}\alpha)/\pi$. Mathematically, $\sin(\frac{\pi}{2}\alpha)$ is the same as $\sin(\frac{\pi}{2}\varepsilon)$ where $\varepsilon = 2 - \alpha$. However from the point of view of computational precision, if α is near to 2 and $\varepsilon = 2 - \alpha$ is known to good relative precision, then it is better to use $c_\alpha = \Gamma(\alpha)\sin(\frac{\pi}{2}\varepsilon)/\pi$.

Another example where deviations from a standard value matter is in numerical quadrature. Computing W using formula (1) on page 6 is much more accurate in circumstances where $\theta - \phi_0$ and $1 - \alpha$ are small in absolute value if they can somehow be obtained to good relative precision rather than being computed by subtraction from θ and α.

Translating between parametrizations

One annoying feature of stable distributions is that there are many different parametrizations. See, for instance, Zolotarev (1986) or Samorodnitsky and Taqqu (1984). Different parametrizations are often convenient for different purposes. Hall (1980) pointed out that mistakes concerning the direction of skewness of different parametrizations had been made by many people, including himself. Here, we are concerned with translating numerical values of parameters between different parametrizations.

For computations involving translating between different parametrizations, it is often more accurate to work with deviations from standard values rather than the parameter values themselves. For instance, let β_A denote the value of β for the A (or S1) parametrization and β_C denote the value of β for the C parametrization. The mathematical relationship between them is usually written $\beta_C = \arctan(\beta_A \tan(\frac{\pi}{2}\alpha))/(\frac{\pi}{2}\alpha)$. However, greater computational accuracy can be achieved by alternative, mathematically equivalent formulae in appropriate circumstances.

- If α is near unity and its difference from unity is known more accurately than α itself, compute $\tan(\frac{\pi}{2}\alpha)$ as $1/\tan(\frac{\pi}{2}(1-\alpha))$.
- If α is near two and its difference from two is known more accurately than α itself, compute $\tan(\frac{\pi}{2}\alpha)$ as $-\tan(\frac{\pi}{2}(2-\alpha))$.
- If β_A is near unity and $1 - \beta_A$ is known more accurately than β_A then rather than computing $\beta_C = \arctan(\beta_A \tan(\frac{\pi}{2}\alpha))/(\frac{\pi}{2}\alpha)$ in the obvious way, first compute

$$\tan(\tfrac{\pi}{2}\alpha(1-\beta_C)) = \frac{\tan(\frac{\pi}{2}\alpha) - \tan(\frac{\pi}{2}\alpha\beta_C)}{1 + \tan(\frac{\pi}{2}\alpha)\tan(\frac{\pi}{2}\alpha\beta_C)} = \frac{(1-\beta_A)\tan(\frac{\pi}{2}\alpha)}{1 + \beta_A \tan^2(\frac{\pi}{2}\alpha)}$$

  and then compute $1 - \beta_C$ by taking the arctangent of this quantity. Similar computation of $1 + \beta_C$ from $1 + \beta_A$ can be done when β_A is near -1.

When working with finite moment log-stable distributions I often use the mean and standard deviation, denoted by μ and σ respectively, as parameters rather than the location and scale, often denoted by δ and γ, of the corresponding stable distribution.

The most obvious standard values for μ and σ are both unity. Figure 1 shows these standard stable distributions for a range of values of the characteristic exponent, α. The figure illustrates the following features.

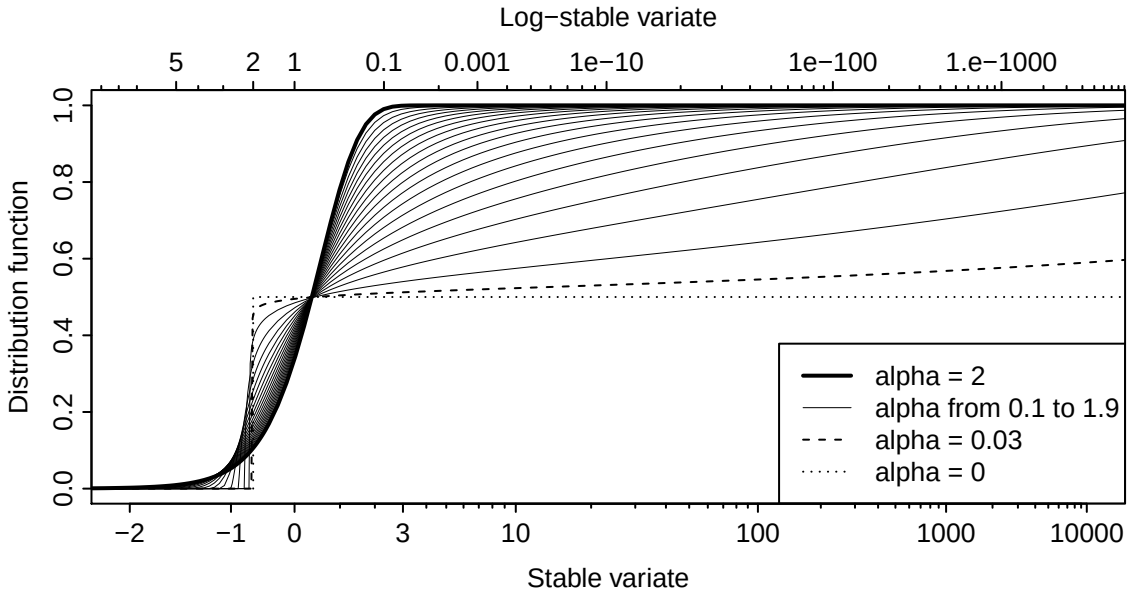

Figure 1: **Distribution functions for various α for finite moment log-stable distributions which have mean 1 and standard deviation 1.**

- For $\alpha = 2$ the stable variate has a normal distribution and the log-stable variate has a log-normal distribution.
- The left-hand tails of the stable distributions are less heavy than the left-hand tail of the normal distribution which corresponds to $\alpha = 2$. The right-hand tails of the stable distributions are heavier than the right-hand tail of the normal distribution which corresponds to $\alpha = 2$.
- The limit as $\alpha \to 0$ for the log-stable variate has probability $\frac{1}{2}$ at $X = 0$ and probability $\frac{1}{2}$ at $X = 2$. The stable variate has probability $\frac{1}{2}$ at $X = -\log(2)$ and probability $\frac{1}{2}$ at $X = \infty$.
- Dealing with these distributions can often require use of numbers outside the ranges used for storing real numbers on computers (e.g. 4.9×10^{-324} to 1.8×10^{308}). For several of the log-stable distributions shown there is a substantial probability of values smaller than 10^{-1000}.

The moment generating function for log maximally skew stable distributions is given by Samorodnitsky and Taqqu (1994, page 15, proposition 1.2.12). It can be found by taking the analytic continuation of the characteristic function to values of t on the imaginary axis. For the A or S1 parametrization , it is

$$M_A(t) = \mathrm{E}[\exp(-t(\delta + \gamma X_A))] = \exp\left(-\gamma^\alpha t^\alpha / \cos(\tfrac{\pi}{2}\alpha) - \delta t\right).$$

For the M or S0 parametrization, it is

$$H_M(t) = \exp\left(-\gamma^\alpha t^\alpha / \cos(\tfrac{\pi}{2}\alpha) + t\tan(\tfrac{\pi}{2}\alpha) - \delta t\right),$$

as can be seen by replacing δ by $\delta - \tan(\frac{\pi}{2}\alpha)$. When $\alpha = 1$, the moment generating function can be found as the limit of this expression as $\varepsilon \to 0$, namely $\exp(\frac{2}{\pi} t \log t) = (\gamma t)^{2t/\pi - \delta t}$.

The C parametrization has a scale which is different by a factor of $\cos^{1/\alpha}(\frac{\pi}{2}\alpha)$ from the A (or S1) parametrization, so its moment generating function is

$$M_C(t) = \exp\left(-(\gamma t)^\alpha - \delta t\right).$$

The relationship between the location and scale parameters, δ and γ, and the mean and standard deviation of the finite moment log-stable distribution, μ and σ, is readily computed in either direction. For the M parametrization, the first moment is

$$\exp(-\delta)\exp\left(\frac{\gamma}{\sin(\frac{\pi}{2}\varepsilon)}\left[\gamma^{\alpha-1} - \sin(\tfrac{\pi}{2}\alpha)\right]\right)$$

and the second moment is

$$\exp(2-\delta)\exp\left(\frac{2\gamma}{\sin(\frac{\pi}{2}\varepsilon)}\left[(2\gamma)^{\alpha-1} - \sin(\tfrac{\pi}{2}\alpha)\right]\right).$$

For the C parametrization, the first moment is

$$\exp(-\delta)\exp\left(-\gamma^\alpha\right)$$

and the second moment is

$$\exp(2-\delta)\exp\left(-(2\gamma)^\alpha\right).$$

In both cases, the ratio of the second moment to the square of the first moment (which ratio we will denote by r) is a function of γ which does not involve δ. This can be used to find the value for the parameter γ for specified first moment, μ, and second moment, $\mu^2+\sigma^2$. The ratio must have the value $r = 1+(\sigma/\mu)^2$. For the M (or S0) parametrization, provided that $\alpha \neq 1$, the equation for γ is of the form

$$\exp\left(\frac{2\gamma}{\sin(\frac{\pi}{2}\varepsilon)}\left[(2\gamma)^{\alpha-1} + \sin(\tfrac{\pi}{2}\alpha)\right]\right) \div \exp\left(\frac{2\gamma}{\sin(\frac{\pi}{2}\varepsilon)}\left[\gamma^{\alpha-1} + \sin(\tfrac{\pi}{2}\alpha)\right]\right) = r.$$

Taking logarithms of both sides and simplifying:

$$\frac{2\gamma}{\sin(\frac{\pi}{2}\varepsilon)}\left[(2\gamma)^\varepsilon - \gamma^\varepsilon\right] = \log(r).$$

This is satisfied when

$$\gamma^\alpha = \frac{\log(r)\sin(\frac{\pi}{2}\varepsilon)}{2(2^\varepsilon - 1)} = \frac{\log(r)\sin(\frac{\pi}{2}\varepsilon)}{2\,\mathrm{expm1}(\varepsilon\log 2))}.$$

For the case $\alpha = 1$, the corresponding equation is $\gamma = \pi\log(r)/(4\log 2)$. For the C parametrization, the equation for γ is of the form

$$\exp\left(-(2\gamma)^\alpha\right) \div \exp\left(-2(\gamma)^\alpha\right) = r.$$

This also has an explicit solution:

$$\gamma = \left[\log(r)/(2-2^\alpha)\right]^{1/\alpha}.$$

Code in R

I have written a package called `FMStable` for R which implements many of the ideas discussed in this paper. One important aspect of this implementation is that it uses the C parametrization when $\alpha < \frac{1}{2}$ and to use the M or S0 parametrization when $\alpha \geq \frac{1}{2}$, but hides this from the user.

2 Precise computation of density and distribution functions

I have written code in Fortran90 for computing the density and distribution function of stable distributions. This code can be run in using either 64-bit precision or 128-bit precision for the floating point numbers.

The mathematical formulae behind this computation can be understood most easily relative to the simulation method of Chambers, Mallows and Stuck (1976) which uses the C parametrization. For $\alpha \neq 1$, the distribution function is

$$F_\alpha(x) = \int \exp(-W(\theta, x))\, d\theta$$

where

$$W = \cos(\theta - \alpha(\theta - \phi_0)) \sin(\alpha(\theta - \phi_0))^{\frac{\alpha}{1-\alpha}} \cos(\theta)^{-\frac{1}{1-\alpha}} x^{-\frac{\alpha}{1-\alpha}}. \tag{1}$$

The integration limits are some combination of $-\frac{\pi}{2}$, $\frac{\pi}{2}$ and $\phi_0 = -\frac{\pi}{2}\beta k(\alpha)/\alpha$ where $k(\alpha) = 1 - |1 - \alpha|$, depending on the values of α, β and x. For $\alpha = 1$, the distribution function is of the form

$$F_1(x) = \int_{-\frac{\pi}{2}}^{\frac{\pi}{2}} \exp(-W(\theta, x))\, d\theta$$

where

$$W = \tfrac{2}{\pi} \exp\left[\{(\tfrac{\pi}{2} + \beta\theta) \tan\theta - \tfrac{2}{\pi}x\}/\beta\right] (\tfrac{\pi}{2} + \beta\theta) \sec\theta.$$

Formally differentiating these expressions with respect to x gives the density as a one-dimensional definite integral.

The accuracy of this Fortran90 code was checked during its development by comparing results obtained using 64-bit and 128-bit precision. It was often found useful to store critical quantities relative to more than one origin so that the most appropriate could readily be used. For instance, as well as storing α, the values of $\alpha - 1$ and $2 - \alpha$ were also stored. Similarly, values of $\beta + 1$ and $1 - \beta$ were stored as well as β. For numerical integration with respect to θ, the quantities $\theta \pm \frac{\pi}{2}$, $\alpha(\theta - \phi_0) \pm \pi$ and $\theta - \alpha(\theta - \phi_0) \pm \frac{\pi}{2}$ were also stored as well as θ.

Later, I wrote some code in R which used a strategy for reducing rounding errors which has the same effect but is compatible with existing software for automatic numerical integration. The variable of integration is taken to range from -1 to $+1$, so that even when the range of integration is halved up to 52 times there will be no rounding error in the computer representation of the endpoints of subintervals. Then the quantities θ, $\theta \pm \frac{\pi}{2}$, $\alpha(\theta - \phi_0) \pm \pi$ and $\theta - \alpha(\theta - \phi_0) \pm \frac{\pi}{2}$ for the intermediate points. This code is included in an `R` package called `FMStable`.

3 Interpolation approximation to distribution function and probability density function in various regions of the parameter space

The main novel contribution of this paper is to discuss how the probability distribution function and probability density can be computed satisfactorily by interpolation. The value judgement that only the maximally skew stable distributions are of primary interest is important. This reduction in the scope of the problem means that only two-dimensional interpolation is required.

In order to provide good approximations to the density, the distribution function and the right tail probability for all α and all x, it was found necessary to use different mathematical forms in each of several different regions. The regions where the various approximations have been used are indicated on Figure 2. In some cases, there is overlap between adjacent regions of validity of interpolation formulae, but this is not important.

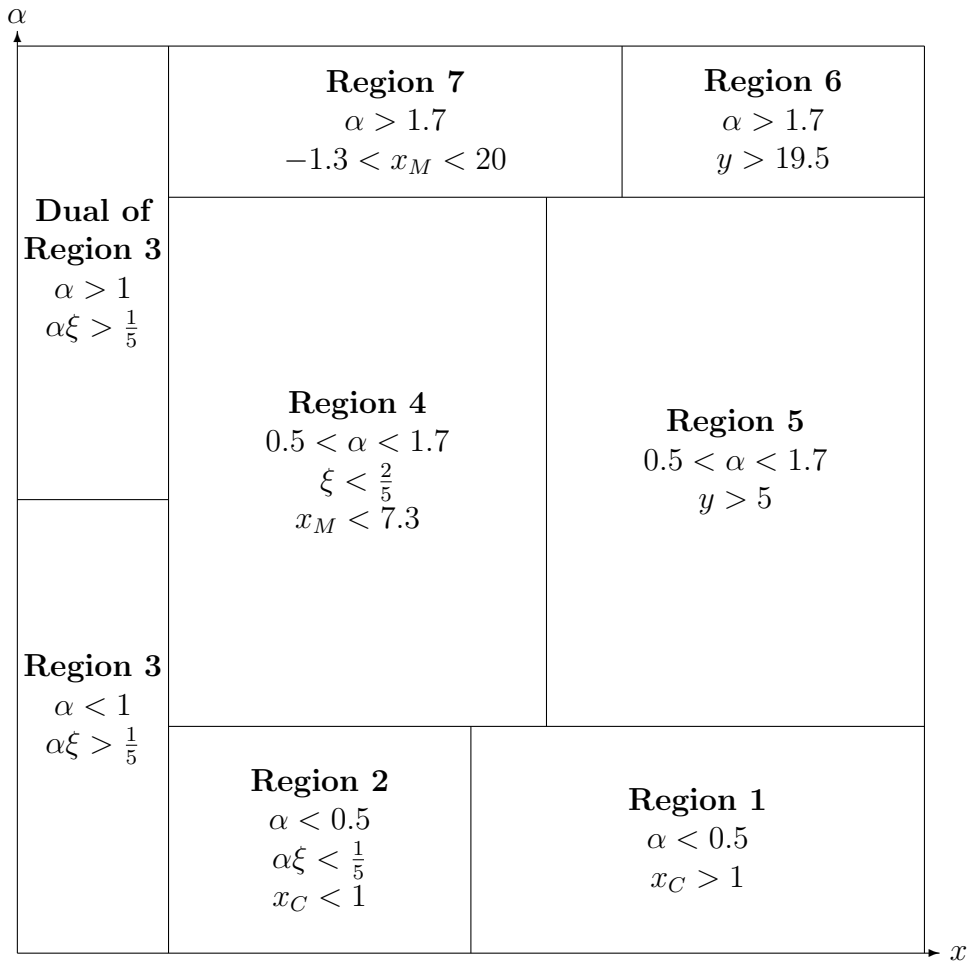

Figure 2: *Regions where different approximations were used for the density and distribution function for log maximally-skewed stable distributions*

Region 1

From Holt and Crow (1973) section 2.21 or Zolotarev (1986) equations 2.4.3 and 2.4.8, the probability density at x for the C parametrization is given by the convergent series

$$f_\alpha(x) = \frac{1}{\pi x}\sum_{k=1}^{\infty}\frac{(-1)^{k-1}}{k!}\Gamma(\alpha k+1)x^{-\alpha k}\sin(\tfrac{\pi}{2}k\alpha).$$

The probability in the right tail is

$$1-F_\alpha(x) = \frac{1}{\pi\alpha}\sum_{k=1}^{\infty}\frac{(-1)^{k-1}}{k\times k!}\Gamma(\alpha k+1)x^{-\alpha k}\sin(\tfrac{\pi}{2}k\alpha). \tag{2}$$

These series suggest that $xf_\alpha(x)$ and $1-F_\alpha(x)$ can be interpolated as functions of $x^{-\alpha}$ and α. Such interpolation was found to be reasonably accurate (i.e. relative errors apparently less than 10^{-14}) over the range $\alpha < 0.5$ and $x > 1$ in the C parametrization with 20 Chebyshev-spaced nodes in each of the variables. This is Region 1 on Figure 2.

The one-dimensional interpolation method used is always based on 16 nodes. If there are 8 nodes available on each side of a point for which an interpolated function value is required then the nearest 8 nodes on each side of that point are used. Otherwise, the nearest 16 nodes are used. This form of interpolation is moderately efficient and was kept constant while other aspects of the interpolation procedure were varied, in order to reduce the complexity of the search for good methods of interpolation.

Using the first term of series 2 and taking $c_\alpha = \Gamma(\alpha)\sin(\frac{\pi}{2}\alpha)/\pi$ gives the approximation

$$1-F_\alpha(x) \approx 2c_\alpha x^{-\alpha}.$$

This is useful as a first approximation when finding quantiles.

Region 2

From equation (2), as $\alpha \to 0$ the probability in the right tail tends to

$$\frac{1}{\pi\alpha}\sum_{k=1}^{\infty}\frac{(-1)^{k-1}}{k\times k!}x^{-\alpha k}\left[\tfrac{\pi}{2}k\alpha\right] = \frac{1}{2}\sum_{k=1}^{\infty}\frac{(-1)^{k-1}}{k!}x^{-\alpha k} \approx \tfrac{1}{2}\left(1-\exp(-x^{-\alpha})\right).$$

In Region 2, interpolation was done using the variable $1-\exp(-x^{-\alpha})$, rather than $\exp(-x^{-\alpha})$ as in Region 1. Again, 20 Chebyshev-spaced nodes were used in each of the variables.

The approximation

$$F_\alpha(x) \approx 1-2c_\alpha x^{-\alpha}.$$

was used as a first approximation when finding quantiles.

Region 3 and its dual

Zolotarev (1986) equation (2.5.17) tells us that the density of a maximally-skewed stable variable when $\alpha < 1$ for values of x in the C parametrization near zero is approximately

$$\frac{\nu}{\sqrt{2\pi\alpha}}\xi^{\frac{2-\alpha}{2\alpha}}\exp(-\xi)\left\{1+\sum_{n=1}^{\infty}Q_n(\alpha)(\alpha\xi)^{-n}\right\}$$

where $\xi = (1-\alpha)(x/\alpha)^{-\alpha/(1-\alpha)}$ and $\nu = (1-\alpha)^{-1/\alpha}$. The terms $Q_n(\alpha)$ are polynomials of degree $2n$.

Similarly, Zolotarev (1986) equation (2.5.20) tells us that the distribution function of a maximally-skewed stable variable when $\alpha < 1$ for values of x in the C parametrization near zero is approximately

$$\frac{1}{\sqrt{2\pi\alpha\xi}}\exp(-\xi)\left\{1+\sum_{n=1}^{\infty}\tilde{Q}_n(\alpha)(\alpha\xi)^{-n}\right\}$$

where the polynomials $\tilde{Q}_n(\alpha)$ are not the same as $Q_n(\alpha)$.

We do not need to evaluate the polynomials $Q_n(\alpha)$ or $\tilde{Q}_n(\alpha)$. These formulae suggest that interpolation as functions of α and $(\alpha\xi)^{-1}$ can be used to approximate the expressions in large brackets, at least when $\alpha\xi$ is large. Calculations suggest that 20 Chebyshev-spaced nodes in α and 70 Chebyshev-spaced nodes in $(\alpha\xi)^{-1}$ was adequate to achieve good accuracy provided that $\alpha\xi < \frac{1}{5}$.

Zolotarev (1986) says that this formula also applies when $\alpha > 1$ and $x \to \infty$ provided that α is replaced by $1/\alpha$ in the summation. This could also be shown by the principle of duality which is most simply stated in the C parametrization. See section 2.3 of Zolotarev (1986). The portions of maximally skew stable distributions for $\alpha > 1$ for positive x are related to portions of the maximally skew stable distributions for $\alpha < 1$. Denoting the distribution function by $F_\alpha(x)$ and the density function by $f_\alpha(x)$; if $\alpha > 1$ then

$$\alpha\,(1 - F_\alpha(x)) = F_{1/\alpha}(x^{-\alpha})$$

and

$$f_\alpha(x) = x^{-1-\alpha} f_{1/\alpha}(x^{-\alpha})$$

Note also that the value of ξ is the same for the points related by duality. For $\alpha > 1$, it should be noted that ξ as a function of x in the complex domain has an essential singularity at $x = 0$ except for the case when $\alpha = 2$. Hence this approximation cannot be expected to be useful for negative x or for x near to zero.

The formulae above are for the C parametrization. For the M (or S0) parametrization, x needs to be replaced by xs where $s = (1 + \tan^2(\frac{\pi}{2}k))^{-1/(2\alpha)}$ in the formula for the distribution function. For the density, there needs to be a factor of s as well as this replacement. When $\alpha < 1$,

$$\xi = (1-\alpha)\left(\frac{x + \tan(\frac{\pi}{2}\alpha)}{\alpha\cos(\frac{\pi}{2}\alpha)}\right)^{-\alpha/(1-\alpha)}$$

The inverse relationship for x in terms of ξ is

$$x = \alpha\left(\frac{1-\alpha}{\xi}\right)^{(1-\alpha)/\alpha}\left(\cos(\tfrac{\pi}{2}\alpha)\right)^{-1/\alpha} - \tan(\tfrac{\pi}{2}\alpha).$$

This relationship is not computationally practical for α near to 1. It can be rewritten as

$$x = \frac{\alpha}{\cos(\frac{\pi}{2}\alpha)}\text{expm1}\left(\frac{1-\alpha}{\alpha}\log\frac{1-\alpha}{\xi\cos(\frac{\pi}{2}\alpha)}\right) + \frac{\alpha}{\cos(\frac{\pi}{2}\alpha)} - \tan(\tfrac{\pi}{2}\alpha).$$

Computationally, this formula is handled by first calculating four quantities which are dependent only on α or on $\varepsilon = 1-\alpha$. It turns out that these formulae work for $\alpha > 1$ also, even though the earlier relationships would need to be modified by addition of some modulus signs and multiplication by $\text{sign}(1-\alpha)$.

$$C_1 = \frac{\alpha}{\cos(\frac{\pi}{2}\alpha)} = \frac{\alpha}{\sin(\frac{\pi}{2}\varepsilon)}$$

$$C_2 = \frac{1-\alpha}{\alpha} = \frac{\varepsilon}{1-\varepsilon}$$

$$C_3 = \frac{1-\alpha}{\cos(\frac{\pi}{2}\alpha)} = \frac{\varepsilon}{\sin(\frac{\pi}{2}\varepsilon)}$$

$$C_4 = \frac{\alpha}{\cos(\frac{\pi}{2}\alpha)} - \tan(\tfrac{\pi}{2}\alpha) = \frac{1-\varepsilon-\sin(\frac{\pi}{2}\alpha)}{\cos(\frac{\pi}{2}\alpha)} = \frac{2\sin^2\frac{\pi}{4}\varepsilon - \varepsilon}{\sin\frac{\pi}{2}\varepsilon}$$

Then translation between x and ξ for the M (or S0) parametrization can be handled using the equations

$$x = C_1\text{expm1}\left(C_2\log(C_3/\xi)\right) + C_4$$

$$\xi = C_3/\exp\left(\text{log1p}(\frac{x-C_4}{C_1})/C_2\right)$$

In these regions, an approximation to ξ for given F is found by approximately solving the equation

$$F = \frac{1}{\sqrt{2\pi\alpha\xi}}\exp(-\xi).$$

A first approximation is $\xi = -\log(F)$. This is refined using a single Newton-Raphson iteration.

Region 4

In this region it is not necessary to match the method of interpolation with any asymptotic behaviour. The logarithm of the right hand tail probability and the logarithm of the probability density are interpolated as functions of α and x. Accuracy appeared to be satisfactory with 40 Chebyshev-spaced nodes over α and 60 Chebyshev-spaced nodes over x.

Approximate quantiles are found by using the approximation for Regions 3 and its dual if the left hand tail probability is the smaller, and by using the approximation for Regions 5 and 6 if the right hand tail probability is the smaller.

Regions 5 and 6

A good approximation in these regions is given in Zolotarev (1986) Theorem (2.5.6). We need to compute

$$\eta = \tan(\tfrac{\pi\alpha}{2})$$

and

$$x = y + \eta y^{1-\alpha}.$$

The density and right hand tail probability at x are approximately

$$f_\alpha(x) \approx \frac{1}{\pi x} \sum_{n=1}^{\infty} A_n(\alpha) y^{-\alpha n}$$

$$1 - F_\alpha(x) \approx \frac{1}{\pi\alpha} \sum_{n=1}^{\infty} \frac{1}{n} \left(A_n(\alpha) + (1-\alpha) A_{n-1}(\alpha)\right) y^{-\alpha n}$$

where $A_0 = 0$ and, using "Im" to stand for "imaginary part",

$$A_n(\alpha) = \text{Im} \sum_{k=1}^{n} \frac{\Gamma(\alpha k + n - k + 1)}{\Gamma(k+1)\Gamma(n-k+1)} (-\eta)^{n-k} e^{i\pi\alpha k/2} (\eta e^{i\pi/2} - 1)^k.$$

In particular, note that

$$A_1(\alpha) = \text{Im}\, \frac{\Gamma(\alpha+1)}{\Gamma(2)\Gamma(1)} e^{i\pi\alpha/2} (\tan(\tfrac{\pi}{2}\alpha) e^{i\pi/2} - 1)$$

$$= \Gamma(\alpha+1)\ \text{Im}\ \left\{\cos(\tfrac{\pi}{2}\alpha) - i\sin(\tfrac{\pi}{2}\alpha)\right\} \left\{i\tan(\tfrac{\pi}{2}\alpha) - 1\right\} = 2\Gamma(\alpha+1)\sin(\tfrac{\pi}{2}\alpha). \qquad (3)$$

For the purpose of interpolation, the density and the right tail probability at x can be expressed as quantities which depend on α times $y^{-1-\alpha}$ times a polynomial in $y^{-\alpha}$ which may be taken to be unity at infinity (i.e. $x = 0$).

This for the A (or S1) parametrization. For the M (or S0) parametrization, the value of x for given y is

$$x = y + \eta y^{1-\alpha} - \eta = y + \tan(\tfrac{\pi}{2}\alpha) \left[y^{1-\alpha} - 1\right]$$
$$= y + \text{expm1}\left(\varepsilon \log(y)\right) / \tan(\tfrac{\pi}{2}\varepsilon)$$

where $\varepsilon = 1 - \alpha$. The limit as $\varepsilon \to 0$ (i.e. as $\alpha \to 1$) is $x = y + \frac{2}{\pi}\log(y)$.

Zolotarev (1986) indicates that this approximation is intended to be applied when $\alpha < 1$, so interpolation in terms of α and $y^{-1/\alpha}$ can be expected to be satisfactory. Numerical work has indicated that such interpolation also works well when $\alpha > 1$.

Interpolation in Region 5 was done using 40 Chebyshev-spaced nodes over α and 20 Chebyshev-spaced nodes over $y^{-1/\alpha}$. Interpolation in Region 6 was done using 17 Chebyshev-spaced nodes over α and 20 Chebyshev-spaced nodes over $y^{-1/\alpha}$.

Approximations to quantiles can be found by truncating the series in equation (**??**) and using the known value for $A_1(\alpha)$.

$$1 - F_\alpha(x) \approx \frac{1}{\pi\alpha} 2\Gamma(\alpha+1)\sin(\tfrac{\pi}{2}\alpha) y^{-\alpha}$$

Values for y can be substituted into equation (3) to find values for x.

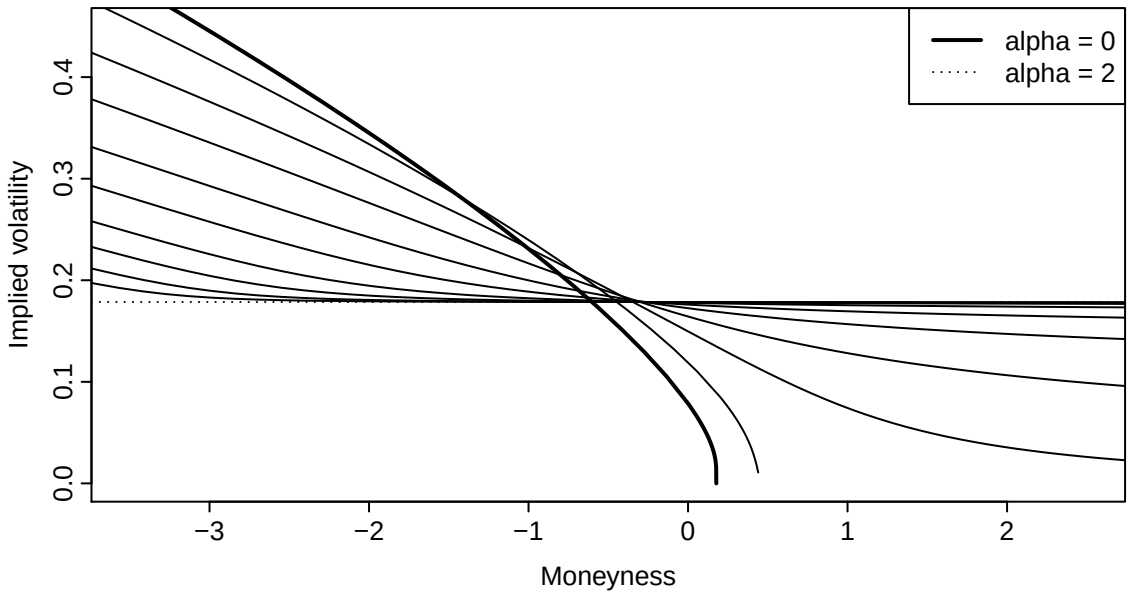

Figure 3: **Implied volatility for option prices based on finite moment log-stable models with mean 1 and standard deviation 0.18 for many different values of the parameter α.**

Region 7

As $\alpha \to 2$, it appears that

$$f_2'(x) = \lim_{\alpha \to 2} \frac{\partial f_\alpha(x)}{\partial \alpha}$$

and

$$F_2'(x) = \lim_{\alpha \to 2} \frac{\partial F_\alpha(x)}{\partial \alpha}$$

are bounded and are smooth functions of x. Interpolation in this region was done using 17 Chebyshev-spaced nodes over $2 - \alpha$ and 100 Chebyshev-spaced nodes over x. This could probably be made computationally faster if good numerical procedures (such as continued fractions) were available for $f_2'(x)$ and $F_2'(x)$.

In this region, approximate quantiles were found using the fact that the distribution for $\alpha = 2$ is normal with variance 2.

4 Values of options

I have written code for valuing options by numerical integration of the distribution function. This provides high precision and is sufficiently rapid for current purposes. Greater speed could be achieved by using interpolation over a three-dimensional table. Other methods that might be useful for achieving greater speed are numerical inversion of Fourier transforms, finding rational approximations at least for the most commonly-used regions of the parameter space (such as α near to 2 and coefficient of variation in the range 0.01 to 0.3), and finding interpolation formulae for the distribution functions of finite moment log-stable distributions that can be integrated term by term.

Figure 3 shows the values of options on scales which are often used by options traders. The horizontal axis gives the moneyness. This is the number of standard deviations on a logarithmic scale by which the strike price exceeds the current or spot price. The vertical axis gives the implied volatility. This is the volatility such that the Black-Scholes models gives option values derived from the finite moment log-stable model. The non-bold continuous lines are for finite moment log-stable distributions with mean 1, standard deviation 0.18, and probabilities 0.01, 0.003, 0.001, 0.0003, 0.0001, 0.00003, 0.00001, 0.000003 and 0.000001 of being less than 0.01. The corresponding values for α are 0.687, 1.197, 1.527, 1.769, 1.897, 1.964, 1.9873, 1.99609 and 1.99869. The distribution corresponding to $\alpha = 0$ is discrete with probability 0.03138 at zero and probability 0.96861 at 1.0324. The log-normal distribution corresponding to $\alpha = 2$ has volatility 0.1786 which does not vary with moneyness.

The lines for values of α greater than 1.0 all have shapes consistent with what is called a "volatility smile" or a "volatility smirk".

Estimating the parameter α

The parameter α cannot be estimated accurately from small amounts of data. Figure 4 shows the maximum likelihood estimates of α, $\hat{\alpha}$, from 200 simulations. In each simulation, the data were 100 independent numbers from a finite moment log-stable distribution with mean 1.001, standard deviation 0.025 and $\alpha = 1.8$. These data are in some ways like a set of 100 weekly returns from a share after adjustment for overall market trends, and are referred to as returns. However, weekly returns from real markets are heteroskedastic and have variable correlation structures. Kring, Rachev, H(̈o)chstötter and Fabozzi (2009) found that stable distributions with $\alpha \approx 1.6$, $\beta \approx 0$ were good fits to daily log-returns for 29 stocks that were in the DAX30 index from May 2002 to March 2006.

Here, estimates of α, the mean and the standard deviation were found by maximum likelihood. There were 11 cases where $\hat{\alpha}$ was 2. We can see from Figure 4 that there is a fairly strong relationship between the smallest of the 100 returns and $\hat{\alpha}$.

If a player in financial markets were to estimate α for a large number of companies' shares using two years of data on weekly returns for each company, then the $\hat{\alpha}$ values would depend largely on whether there has been any substantial fall in that company's share price. This variation in would $\hat{\alpha}$ values result in out-of-the-money put options having much smaller prices for companies which had had a substantial fall in share price and for companies which had not had a substantial fall in share price. It seems likely to me that the real values of α would be much less variable over companies that the values of $\hat{\alpha}$.

One alternative way to estimate α would be to calculate the mean and standard deviation of daily or weekly returns, and to subjectively estimate the probability of bankruptcy or the probability of a drop of, say, 30% in price. The parameters of a finite moment log-stable distribution could be fitted to these three pieces of information.

Another way to estimate α would be to use data on the market prices for options, and to find the finite moment log-stable distribution which best fits those prices.

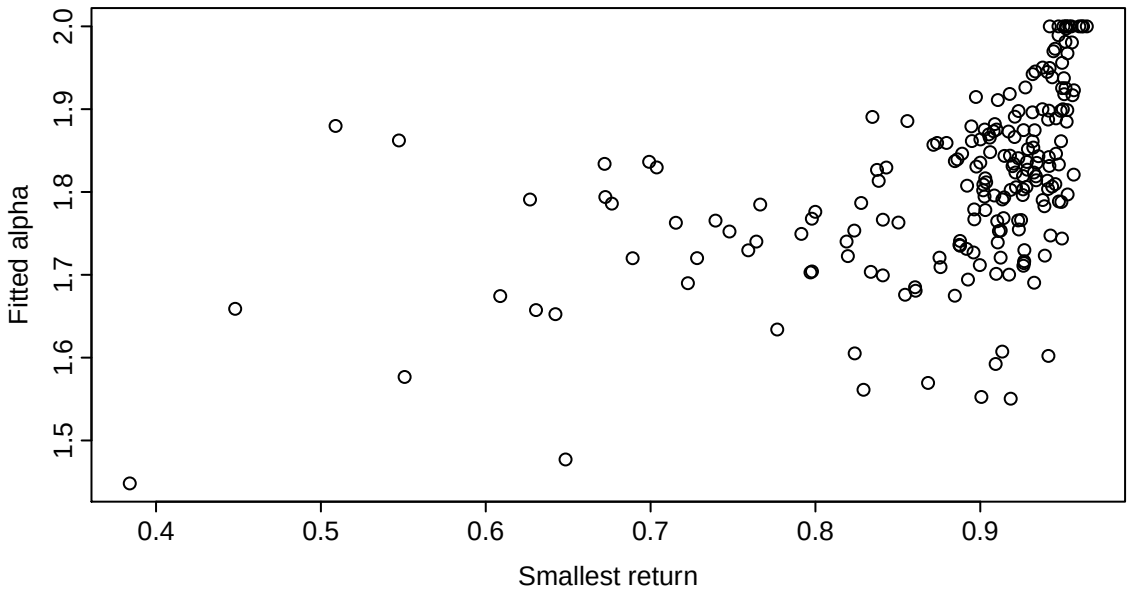

Figure 4: **Estimated values for α for 200 simulated data set each with 100 returns which were from a finite moment log-stable distribution with mean 1.001, standard deviation 0.0.025 and $\alpha = 1.8$.**

Hedging

One interesting aspect of finite moment log-stable models is that attempts to hedge risk are expected to be much less effective under finite moment log-stable models with α substantially less than 2 than under a log-normal model (finite moment log-stable with $\alpha = 2$).

Consider a portfolio which is long one share and short some call options which are at-the-money (meaning that the strike price is equal to the current price) and have one year to expiry. The number of call options is chosen so that the derivative of the value of the portfolio with respect to share price is zero. In Table 1, the second row gives the partial derivative of the value of the call options with respect to current price. This partial derivative is often referred to as "delta". The number of call options in the portfolio is $1/\delta$. It varies only a small amount with the parameter α over the range from $\alpha = 1.5$ to $\alpha = 2$ shown in Table 1.

The expected value of this portfolio is zero for any future time. In order to evaluate this expectation we must take the reduction in the time-to-expiry of the options into account However, the variance of value of the portfolio increases with time into the future. This variance has been calculated by numerical integration over the possible asset prices at that time. The numbers in the body of Table 1 are these variances divided by the length of the time period.

We can see that for the log-normal model ($\alpha = 2$) the variance per unit time is smaller for small time intervals than for large time intervals. The variance per unit time is approximately proportional to the length of the time interval. Therefore this model predicts that the risk of a portfolio can be substantially reduced by frequent adjustment of the hedging ratio. This ability to reduce risk by dynamic hedging decreases as α is reduced from 2. For instance, for $\alpha = 1.9$ the minimum variance per unit time of the

α	1.5	1.6	1.7	1.8	1.9	1.99	1.999	1.9999	2.
δ	0.31994	0.32855	0.33254	0.33354	0.33255	0.33052	0.33027	0.33025	0.33024
Time									
1.000	0.78503	0.74703	0.72040	0.70978	0.71997	0.75095	0.75533	0.75578	0.75583
0.500	0.67526	0.59424	0.51429	0.43951	0.37443	0.32823	0.32441	0.32404	0.32399
0.200	0.62942	0.52913	0.42467	0.31951	0.21756	0.13215	0.12408	0.12328	0.12319
0.100	0.61621	0.51008	0.39806	0.28334	0.16953	0.07123	0.06175	0.06081	0.06070
0.010	0.60509	0.49393	0.37533	0.25220	0.12784	0.01795	0.00719	0.00611	0.00599
0.001	0.60402	0.49236	0.37312	0.24915	0.12374	0.01270	0.00181	0.00072	0.00060

Table 1: Variance per unit time

value of the portfolio per unit time is about 0.12. This is substantially smaller than the variance per unit time of 0.72 that the portfolio is exposed to if the hedging ratio is not dynamically adjusted; but it is much larger than the variance per unit time over a period of 0.01 years according to the log-normal model. Intuitively, this means that much of the risk of the portfolio can be reduced by dynamic hedging, but a component of the risk cannot be eliminated.

5 Discussion

The main point of this paper is to show that use of finite moment log-stable distributions is computationally practical. The issues that have been most critical in making the computations fast enough to be practical are the concentration on maximally-skew stable distributions, so that only two-dimensional rather than three-dimensional interpolation is required, and the use of different forms of interpolation in different parts of the parameter space. The difficulties of dealing with the many different parametrizations for stable distributions and the complicated formulae for moments for the log-stable distributions (and the further complication that different computational algorithms are often appropriate for a single mathematical relationship in different regions of the parameter space) are able to be handled by computer software, rather than requiring users to deal with this complexity. I have done this in `R` by using objects of a class which I have called `stableParameters`, but other software solutions could be developed.

When making use of finite moment log-stable distributions for modelling financial risk, the qualitative feature that I believe matters most is the extent to which risk can be hedged. This comes out as being qualitatively different from when log-normal distributions are used. I recommend that the new methods be used for stress testing the portfolios of organizations with large, partly-hedged portfolios; such as banks, insurance companies and hedge funds; possibly by the regulators of those organizations.

This paper has not discussed the fact that volatility varies over time, due to changes in market sentiment and variations in the rate at which price-sensitive information is expected to arrive. It should be reasonably straight-forward to consider such complexity using finite moment log-stable distributions.—They have only one extra parameter beyond the log-normal model. Such work is beyond the scope of this paper.

The interpolation approach suggested in this paper might be simplified by making use of formulae for derivatives with respect to α at $\alpha = 2$, and could be made more

computationally efficient in several other ways. The interpolation approach could be extended to three variables in order to deal with the distribution function and density of general stable distributions. Similarly, interpolation formulae in three variables might be developed for values of options (or the implied volatility).

References

Kring, S., Rachev, S.T., H$\ddot{(}$o)chstötter, M. and Fabozzi, F.J. (2009). Estimation of α-stable sub-Gaussian distributions for asset returns. Risk assessment: Decisions in Banking and Finance. Edited by Georg Bol, Svetlozar T. Rachev and Reinhold Würth.

Carr, P. and Wu, L. (2003). The finite moment log-stable process and option pricing. Journal of Finance, American Finance Association, Vol. 58(2), 753–778.

Chambers, J.M., Mallows, C.L. and Stuck, B.W. (1976). A method for simulating stable random variables. *Journal of the American Statistical Association* **71,** 340–344.

Feller, W. (1966). An Introduction to Probability Theory and Its Applications. Volume II. Wiley.

Hall, P. (1980). A comedy of errors: the canonical form for a stable characteristic function. *Bull. London Math. Soc.* **13,** 23–27.

Holt, D.R. and Crow, E.L. (1973). Tables and graphs of the the stable probability density functions. *Journal of Research of the National Bureau of Standards—B. Mathematical Sciences,* **77B,** 143–197.

Lèvy, P. (1925). Calcul des Probabilités. Gauthier-Villars, Paris.

Samorodnitsky, G. and Taqqu, M.S. (1994). Stable non-Gaussian random processes : Stochastic models with infinite variance. Chapman & Hall, New York.

Wilkinson, J.H. (1963). Rounding Errors in Algebraic Processes. Prentice-Hall.

Zolotarev, V.M. (1986). *One-Dimensional Stable Distributions.* Amer. Math. Soc. Transl. of Math. Monographs, Vol. 65. Amer Math. Soc., Providence, RI. (Original Russian version was published in 1983.)