Overview

Financial modeling with heavy-tailed stable distributions

John P. Nolan*

The aim of this article was to give an accessible introduction to stable distributions for financial modeling. There is a real need to use better models for financial returns because the normal (or bell curve/Gaussian) model does not capture the large fluctuations seen in real assets. Stable laws are a class of heavy-tailed probability distributions that can model large fluctuations and allow more general dependence structures. © 2013 Wiley Periodicals, Inc.

How to cite this article:
WIREs Comput Stat 2013. doi: 10.1002/wics.1286

Keywords: stable distribution; heavy-tailed models; financial mathematics; robust methods

UNCORRECTED PROOFS

INTRODUCTION

The fluctuations in many financial time series are not normal. The consequences of this are significant: underestimating extreme fluctuations in asset returns causes real hardship to people the world over. Unfortunately, most of the financial world still uses a model based on a normal distribution, even coining phrases like six sigma events to signify fluctuations that should never happen in the lifetime of the earth, yet they have occurred multiple times. One financial expert wryly commented 'We seem to have a once-in-a-lifetime crisis every three or four years'. (Leslie Rahl, founder of Capital Market Risk Advisors, quoted in Ref 1, p. 211.) The simplicity and familiarity of the normal distribution, which is characterized by a mean and a variance, make it an attractive model for practitioners. Yet it does not capture the large fluctuations seen in real-life returns.

In this article, we describe one model for financial returns that explicitly incorporates heavy tails: stable distributions. These are a four parameter family of models that generalize the normal model, allowing both skewness and heavy tails. We do not claim that stable laws perfectly describe real-world returns; no distribution is exact. An old quote relevant here is: 'Essentially, all models are wrong, but some are useful' (Ref 2, p. 424). The key question is what we want to use a model for: if one wants to model the average behavior of an asset, wants a simple model, and is not concerned about extremes, the normal model may be appropriate. Models based on stable laws give another choice: they can describe real data well over most of its range, give a tractable model for compounding returns, and can capture skewness and heavy tails.

Many people have advocated the use of stable laws in finance, starting with Mandelbrot.[3] This idea has been pursued by others, including Samuelson[4] and Rachev and Mittnik.[5] In the past, the lack of efficient numerical methods have made it impractical to use such models in practice. With recent progress in software and increased computational power, it is now worth another look at this class of models.

We note that there are other classes of models that have been proposed for financial returns: generalized t-distributions, generalized hyperbolic, generalized inverse Gaussian, geometric stable, tempered stable, etc. While these models can give a good fit to data sets, they lack all the features described above. A different approach is to use extreme value theory as described in Embrechts et al.[6] and McNeil et al.[7] The discussion of these other methods is beyond

*Correspondence to: jpnolan@american.edu

Department of Mathematics and Statistics, American University, Washington, DC, USA

AQ1 Conflict of interest: •The authors have declared no conflicts of interest for this article.

 © 2013 Wiley Periodicals, Inc.

our scope here, so we focus only on the basics of stable laws and illustrate their use in modeling returns.

UNIVARIATE STABLE LAWS

The theory of stable distributions comes from the pioneering work of Paul Lévy in the 1930s, where he examined what sort of limits can arise when normalizing sums of independent terms. For this reason, these distributions are sometimes called *Lévy stable* laws. Here is the basic definition: a random variable (rv) X is *stable* if for X_1 and X_2 independent copies of X and any positive constants a and b,

$$aX_1 + bX_2 \stackrel{d}{=} cX + d \tag{1}$$

for some positive c and $d \in \mathbb{R}$. ($\stackrel{d}{=}$ denotes 'equal in distribution'). The rv X is *strictly stable* if Eq. (1) holds with $d = 0$ for all choices of a and b. It is *symmetric stable* if it is stable and symmetrically distributed around 0, i.e. $X \stackrel{d}{=} -X$.

For $X_1, \ldots, X_n$ independent and identically distributed as X in Eq. (1), iterating that equation shows that there exist constants $c_n > 0$ and d_n, so that

$$X_1 + \cdots + X_n \stackrel{d}{=} c_n X + d_n. \tag{2}$$

This equation generalizes the familiar property of normal random variables: sums of normal terms are normal. In words, sums of i.i.d. stable terms are stable; this 'stability under addition' property is the reason of the use of the word stable. We started with Eq. (1) and derived Eq. (2), it can be shown that it is possible to reverse this, so either condition can be taken as a definition of stability.

This abstract definition does not specify what the possible distributions are for stable laws. Paul Lévy[8] showed that their characteristic functions (Fourier transform) must have a special form. We will describe two parameterizations here, which we call the 0-parameterization and the 1-parameterization. (There are multiple parameterizations in the mathematical literature: Hall[9] describes a tangled history of meanings of the skewness parameter; Zolotarev[10] has forms A, B, C, C′, E, and M; Samorodnitsky and Taqqu[11] uses the 1-parameterization. To document these and other parameterizations, Nolan[12] lists 11 different parameterizations, numbering them from 0 to 10.)

Four parameters are required to specify a stable law: the *index of stability* α is in the interval (0,2], the skewness β is in the interval $[-1,1]$, the scale parameter γ is any positive number, and the location parameter δ is any number. The notation $S(\alpha,\beta,\gamma,\delta;k)$ will be used to specify a stable distribution with $k = 0$ or $k = 1$ for the two parameterizations.

A random variable X is $S(\alpha,\beta,\gamma,\delta_0;0)$ if it has characteristic function

$$E\exp(iuX) = \begin{cases} \exp\left(-\gamma^\alpha |u|^\alpha \left[1 + i\beta \left(\tan \frac{\pi\alpha}{2}\right)(\operatorname{sign} u) \times \left(|\gamma u|^{1-\alpha} - 1\right)\right] + i\delta u\right) & \alpha \neq 1 \\ \exp\left(-\gamma |u| \left[1 + i\beta \frac{2}{\pi} (\operatorname{sign} u) \times \log(\gamma |u|)\right] + i\delta u\right) & \alpha = 1. \end{cases} \tag{3}$$

A random variable X is $S(\alpha,\beta,\gamma,\delta_1;1)$ if it has characteristic function

$$E\exp(iuX) = \begin{cases} \exp\left(-\gamma^\alpha |u|^\alpha \left[1 - i\beta \left(\tan \frac{\pi\alpha}{2}\right) \times (\operatorname{sign} u)\right] + i\delta u\right) & \alpha \neq 1 \\ \exp\left(-\gamma |u| \left[1 + i\beta \frac{2}{\pi} \times (\operatorname{sign} u) \log |u|\right] + i\delta u\right) & \alpha = 1. \end{cases} \tag{4}$$

Here sign u is the sign of the number u: it is $+1$ if $u > 0$, -1 if $u < 0$, and 0 if $u = 0$, and $x \cdot \log x$ is always interpreted as 0 at $x = 0$. The only difference in the two parameterizations is in the meaning of the location parameter. If $\beta = 0$, then these two parameterizations are identical, it is only when $\beta \neq 0$ that the asymmetry factor (the imaginary term in brackets) becomes an issue, and in this case the laws are shifts of each other: $\delta_0 = \delta_1 + \beta\gamma \tan \frac{\pi\alpha}{2}$ when $\alpha \neq 1$ and $\delta_0 = \delta_1 + \beta \frac{2}{\pi} \gamma \log \gamma$ when $\alpha = 1$.

If one is primarily interested in a simple form for the characteristic function and nice algebraic properties, the 1-parameterization is favored. Because it is simpler to use when proving mathematical properties of stable distributions, it is the most common parameterization in the literature. The main practical disadvantage of the 1-parameterization is that the location of the mode is unbounded in any neighborhood of $\alpha = 1$: if $X \sim S(\alpha,\beta,\gamma,\delta;1)$ and $\beta > 0$, then the mode of X tends to $+\infty$ as $\alpha \uparrow 1$ and tends to $-\infty$ as $\alpha \downarrow 1$, see Figure 1 below. So the 1-parameterization does not have the intuitive properties desirable in applications (continuity of the distributions as the parameters vary, a scale and location family, etc.). We recommend using the 0-parameterization for numerical work and statistical inference with stable distributions: it has the simplest form for the characteristic function that is continuous in all parameters. It lets α and β determine the shape of the distribution, while γ and δ determine scale and location in the standard way: if $X \sim S(\alpha,\beta,\gamma,\delta;0)$, then $(X - \delta)/\gamma \sim S(\alpha,\beta,1,0;0)$. This is not true for the 1-parameterization when $\alpha = 1$.

UNCORRECTED PROOFS

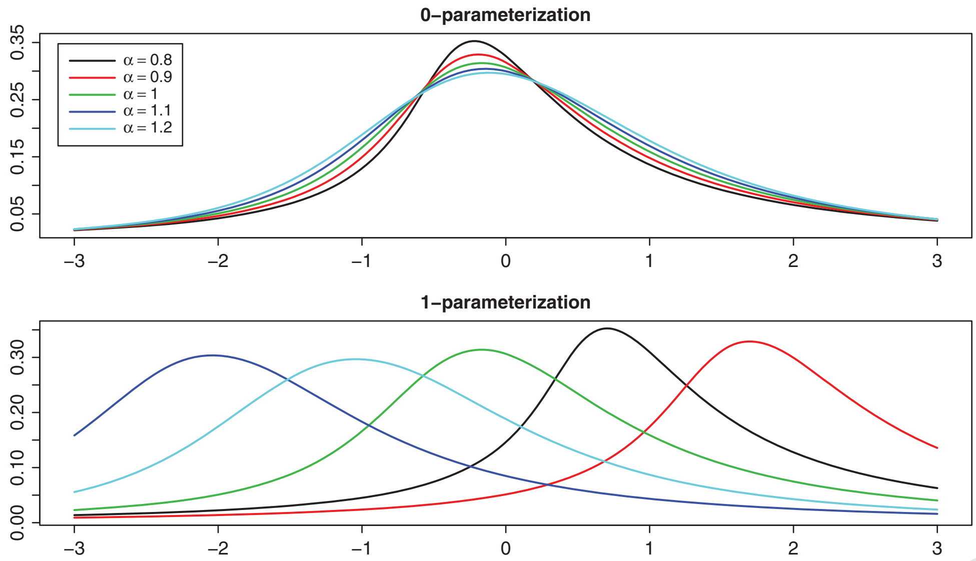

FIGURE 1 | G•raphs of standardized $S(\alpha, \beta = 0.3, 1, 0; 0)$ (0-parameterization, top panel) and $S(\alpha, \beta = 0.3, 1, 0; 1)$ (1-parameterization, bottom panel) densities for a range of α values. The distributions are skewed right because $\beta > 0$. Note that the shapes are the same in both plots for a given (α,β) pair, but the different parameterizations have different shifts. With the 1-parameterization, arbitrarily small changes in α or β can have a large effect on the location of the mode.

Properties of Stable Laws

We summarize some basic properties of $X \sim S(\alpha,\beta,\gamma,\delta;1)$ without proof.

- If $\beta = 0$, then a stable distribution is symmetric.
- Reflection property: $-X \sim S(\alpha, -\beta, \gamma, -\delta; 1)$.
- All stable laws have densities $f(x)$ that are smooth and unimodal.
- In most cases the support of X is the whole real line; the exceptions are when ($\alpha < 1$ and $\beta = 1$), in which case the support is $[\delta, +\infty)$, or ($\alpha < 1$ and $\beta = -1$), in which case the support is $(-\infty, \delta]$.
- Tail behavior. If $\alpha < 2$ and $-1 < \beta \le 1$, then the density $f(x)$ and cumulative distribution function (CDF) $F(x)$ have an asymptotic power law: as $x \to \infty$,

$$1 - F(x) = P(X > x) \sim \gamma^{\alpha} c_{\alpha} (1 + \beta) x^{-\alpha} \quad (5)$$

$$f(x|\alpha, \beta, \gamma, \delta; 0) \sim \alpha \gamma^{\alpha} c_{\alpha} (1 + \beta) x^{-(\alpha+1)}$$

where $c_{\alpha} = \sin \frac{\pi\alpha}{2} \Gamma(\alpha)/\pi$. Using the reflection property, the lower tail properties are similar.Owing to the similarity of the tail behavior to a Pareto distribution (an exact power law), the phrase *stable Paretian distribution* is sometimes used in the non-Gaussian case. For all $\alpha < 2$ and $-1 < \beta < 1$, both tail probabilities and densities are asymptotically power laws. When $\beta = -1$, the right tail of the distribution is not asymptotically a power law; likewise when $\beta = 1$, the left tail of the distribution is not asymptotically a power law. These are not exact relations, only asymptotic ones, and the point at which these approximations are accurate is not known exactly; Fofack and Nolan[13] give some numerical information on this question. The answer is messy: for α near 2, an α-stable law is close to a normal law, and one has to go to a very high quantile to see the power law behavior.
- Fractional moments. When $\alpha < 2$, $E|X|^p$ is finite for $0 < p < \alpha$, but infinite for $p \ge \alpha$. This is a consequence of the power law tail behavior. In particular, for $\alpha < 2$, the (population) variance is infinite and for $\alpha \le 1$, the (population) mean is undefined.
- Generalized Central Limit Theorem. Let $X_1, X_2, \ldots$ be independent identically distributed random variables. The classical Central Limit says that if we start with any distribution with a finite mean μ and standard deviation σ, normalized sums of such terms converge to a normal law:

$$\frac{(X_1 + \cdots + X_n - n\mu)}{n^{1/2}} \xrightarrow{\text{d}} N\left(0, \sigma^2\right).$$

There is a more general result called the Generalized Central Limit Theorem (GCLT) that

 © 2013 Wiley Periodicals, Inc.

UNCORRECTED PROOFS

applies when the summands do not have a finite variance. The simplest version is that when $P(X_i > x) \sim c^+ x^{-\alpha}$ and $P(X_i < -x) \sim c^- |x|^{-\alpha}$ as $x \to \infty$ with $0 < \alpha < 2$ and $c^+ + c^- > 0$. Set $\beta = (c^+ - c^-)/(c^+ + c^-)$, then there are constants b_n and γ such that

$$\frac{(X_1 + \cdots + X_n - b_n)}{n^{1/\alpha}} \xrightarrow{d} S(\alpha, \beta, \gamma, 0; 1)\,.$$

(If $\alpha > 1$, then we may take $b_n = n\mu$.) Note that the normalization factor is different: in the classical case, the scaling is by $n^{1/2}$ whereas in the stable case, the scaling is by the larger factor $n^{1/\alpha}$. There is a more precise statement of the GCLT using the concept of regular variation which can be found in Feller.[14] In fact, stable laws are the only possible nontrivial limits that can arise as limits of normalized sums of i.i.d. terms.

Calculation and Estimation

There are only a few special cases where there are closed form expressions for stable densities $f(x)$. These cases are: (a) Gaussian/normal distributions ($\alpha = 2$, $\beta = 0$), (b) Cauchy distributions ($\alpha = 1$, $\beta = 0$), and (c) Lévy distribution ($\alpha = 1/2$, $\beta = 1$). The only case where there is a closed form expression for the CDF is the Cauchy case. In all other cases, including the CDF for Gaussian laws, numerical procedures are needed to calculate densities and CDFs. Using results of Zolotarev,[10] Nolan[15] describes and implements algorithms to numerically compute stable densities, CDFs, and quantiles when $\alpha < 2$. In addition, the method of Chambers et al.[16] gives an algorithm to simulate. So there are now reliable programs to compute these quantities, making it practical to apply these models to real problems. Figure 1 shows some plots of stable densities.

Many of the standard parameter estimation techniques do not work for stable data. For example, the regular method of moments does not work: to estimate the four parameters one would normally compute EX, EX^2, EX^3, and EX^4 and then try to solve for α, β γ, and δ. But this will not work, since most (or all) of these moments do not exist. (More precisely, higher order population moments do not exist. While the sample moments do exist, but their behavior is erratic: for example, $(1/n)\sum_{i=1}^{n} x_i^2$ will diverge as $n \to \infty$. Large samples do not help with this approach!). Since there are no closed analytic forms for stable densities, the likelihood cannot be written explicitly, making it impossible to analytically solve for maximum likelihood estimators. As a result, there are multiple nonstandard techniques for estimating the stable parameters, some of them ingenious. Four basic methods are the following.

- Tail estimators. This method uses the tail behavior, Eq. (5), to estimate α. Different methods have been proposed for doing this, ranging from plotting extremes on a log–log scale and estimating slope, to the Hill estimator and generalizations. Unfortunately, these do not work very well with stable laws because the when the power law occurs is a complicated function of the parameters and unless one has a very large data set, it is unlikely that the tail will be exactly a power law.
- Fractional moments. When X is strictly stable, there are expressions for fractional moments $E|X|^p$, for $-1 < p < \alpha$. One can use these for a generalized method of moments: compute sample fractional moments, set them equal to the expressions in term of the parameters, and solve for the parameters. Nikias and Shao[17] used this approach in signal processing case when the distribution is symmetric.
- Quantile matching. Fama and Roll[18] noticed certain patterns in tabulated quantiles x_p(=p-th quantile of a distribution) of symmetric stable laws that could be used to estimate α and the scale. For example, the interquartile range $x_{0.75} - x_{0.25}$ is a monotonic function of the scale γ, and the ratio $(x_{0.95} - x_{0.05})/(x_{0.75} - x_{0.25})$ is a monotonic function of the index α. McCulloch[19] generalized this to the nonsymmetric case, using other functions of quantiles to estimate the location δ and skewness β, giving a way to estimate all four stable parameters from a handful of sample quantiles.
- Empirical characteristic functions. Koutrouvelis[20] used the fact that there is an explicit formula (4) for the characteristic function $\phi(u)$. One can compute the sample/empirical characteristic function $\widehat{\phi}(u_i)$ on a grid of u_i values for a data set and then use regression to estimate the parameters. Kogon and Williams[21] simplified this method by using the continuous parameterization, Eq. (3), and centering and scaling the data to avoid numerical difficulties.
- Numerical maximum likelihood estimation. DuMouchel[22] gave an approximate numerical maximum likelihood method and showed that (away from the boundaries $\alpha = 2$ and/or $\beta = \pm 1$,

UNCORRECTED PROOFS

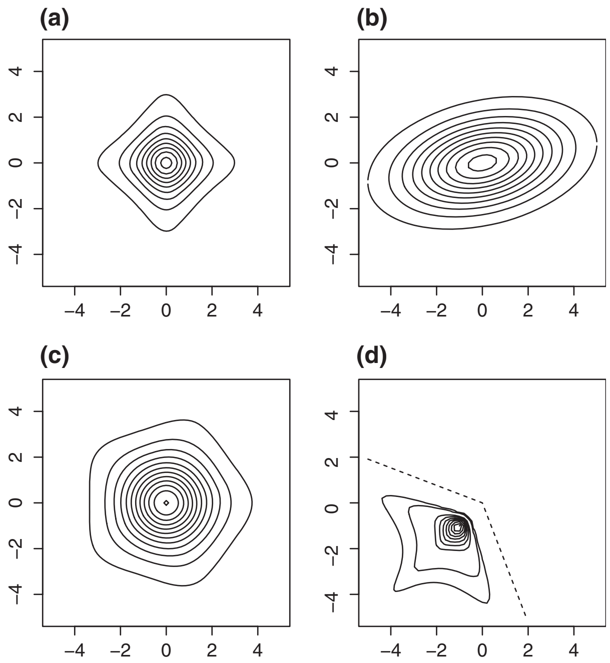

FIGURE 2 | Contour lines of some bivariate stable distributions. (a) Independent components with $\alpha = 1.1$, (b) elliptical contours with $\alpha = 1.7$, (c) discrete spectral measure with 5-point masses and $\alpha = 1.3$, and (d) discrete spectral measure with 3-point masses and $\alpha = 0.75$.

the resulting estimators are asymptotically normal. The present author implemented this in[23] and computed tables that can be used for confidence interval estimates. Further work using a precomputed approximation to stable densities has made this method significantly faster.

See Ref 23 for a summary of these methods and a detailed description of the numerical maximum likelihood approach. Simulations show that the efficiency of the estimate procedures are in reverse order of that listed above.

We end this section with a mention of regression with stable error terms. Nolan and Ojeda[24] describe a procedure for estimating the coefficients for problems of the form

$$Y_i = a_1 x_{i,1} + a_2 x_{i,2} + \cdots + a_m x_{i,m} + \epsilon_i, \quad i = 1, \ldots, n$$

where the error terms ϵ_i are i.i.d. $S(\alpha,\beta,\gamma,\delta;0)$. This gives a robust method of estimating regression coefficients when the error terms are heavy tailed. For example, if one wants to estimate the CAPM β for a volatile asset in relation to a broadly based index, this method is more robust than ordinary least squares.

MODELING FINANCIAL RETURNS

If S_t is the price of an asset at time t, then $X_t := \log(S_t/S_{t-1})$ is the (log) return. (More precisely, this is the log of price changes as it does not include possible dividends or other income.) For many real assets, plots of the returns show a unimodal, mound-shaped distribution. We will show below that some of these distributions have heavier tails than a normal model, and we use a stable distribution to describe the returns.

If $X_t \sim S(\alpha,\beta,\gamma,\delta;1)$ stands for the one period return, then property (2) shows that the multi-period return $X_1 + \cdots + X_n$ is also α-stable. The exact relationship, derived using Eq. (4) is

$$X_1 + \cdots + X_n \sim S\left(\alpha, \beta, n^{1/\alpha}\gamma, n\delta; 1\right).$$

This gives an exact formula for the multi-period return, a very useful fact (•Box 1). AQ4

MULTIVARIATE STABLE LAWS

The definition of stability is exactly the same, with $\mathbf{X}$ a d-dimensional vector in Eq. (1). The surprise here is that there are many possible dependence structures possible in the stable case. One can have an elliptically contoured case, but many other unexpected types of dependence are possible, see the contours of the bivariate stable densities in Figure 2. Feldheim[26] showed that every multivariate stable vector has a characteristic function of the form

$$\begin{aligned}\phi(\mathbf{u}) &= E\exp(i\mathbf{u}\cdot\mathbf{X}) \\ &= \exp\left(\int_{\mathbb{S}} \omega_\alpha(\mathbf{u}\cdot\mathbf{s})\,\Lambda(ds) + i\mathbf{u}\cdot\ \delta\right),\end{aligned}$$

where Λ is a finite measure on the unit sphere $\mathbb{S} = \{|\mathbf{x}| = 1\}$, δ is a shift vector in $\mathbb{R}^d$, and

$$\begin{aligned}\omega_\alpha(t) &= -\log E\exp(itZ) \\ &= \begin{cases} |t|^\alpha[1 + i\tan\frac{\pi\alpha}{2}(\operatorname{sign} t)] & \alpha \neq 1 \\ |t|[1 + i\frac{2}{\pi}(\operatorname{sign} t)\log|t|] & \alpha = 1, \end{cases}\end{aligned}$$

is minus the exponent of the characteristic function of a univariate $Z \sim S(\alpha, \beta = 1, \gamma = 1, \delta = 0; 1)$ Hence every multivariate stable law is characterized by an index of stability α, a spectral measure Λ on the sphere, and a shift δ.

 © 2013 Wiley Periodicals, Inc.

UNCORRECTED PROOFS

BOX 1

A SINGLE ASSET EXAMPLE

Adjusted closing prices for IBM for 10 years were obtained from Yahoo Finance. The period is January 1, 2003–December 31, 2012, giving $n = 2517$ trading days (no weekends or holidays). The price and returns are shown in the following figure.

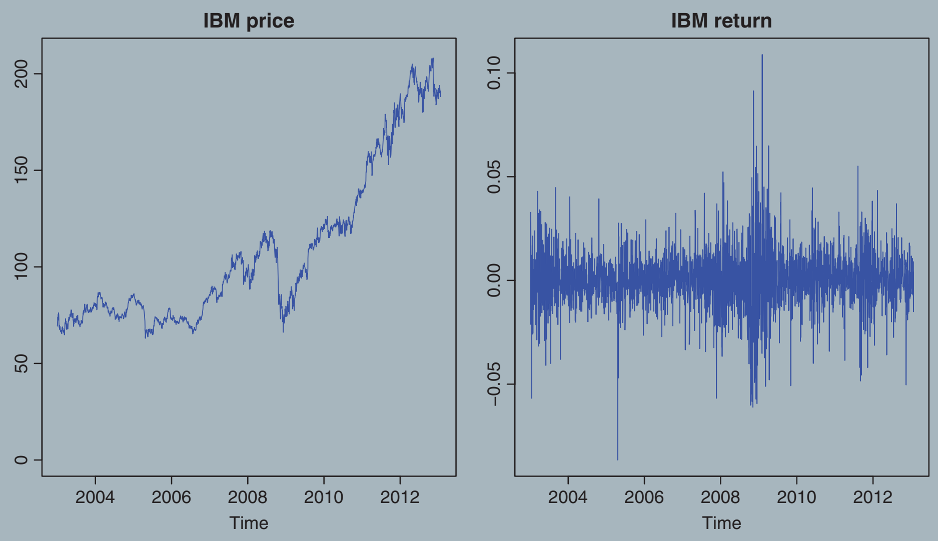

Price and return for IBM-adjusted closing price, 2003–2012.

The simplest computational approach is to analyze the returns directly. Taking all the returns and using numerical maximum likelihood estimation to estimate parameters and giving 95% confidence intervals yields: $\alpha = 1.614873 \pm 0.061525$, $\beta = 0.00000 \pm 0.14376$, $\gamma = 0.007331 \pm 0.000291$, $\delta = 0.000499 \pm 0.000506$, the following figure shows some graphical diagnostics, comparing the observed data, a stable model, and a normal model. The density plot on the left compares the kernel smoothed data density with the stable and normal fits. The plot on the right is a transformed empirical CDF plot that we find useful: in the middle (25th to 75th percentile), it is just a plot of the empirical CDF; on both tails it is a log–log plot. This approach pulls in the extremes horizontally and stretches the vertical axis, allowing one to view the whole distribution. Advantages of this plot are that it is nonparametric, any power decay on the tails show up as straight lines, and one can superimpose different models to compare to the data.

This is a very large class of distributions which cannot be parameterized by a finite number of parameters. To use multivariate stable laws in practice, one has to restrict the type of spectral measure. We describe three accessible classes.

- Independent components. Here the spectral measure is concentrated on the points where the coordinate axes intersect the sphere. The independence makes it easy to work with, simulate and compute densities and CDFs based on the univariate case.
- Discrete spectral measures. Here the spectral measures Λ is discrete, with point mass λ_j at locations $\mathbf{s}_j$. It was shown by Byczkowski et al.[27] that this is a dense class in the sense that for any spectral measure Λ_1, there is a discrete measure Λ_2 with a finite number of point masses such that $|f_1(\mathbf{x}) - f_2(\mathbf{x})|$ (the difference in the corresponding density functions) is uniformly small over all $\mathbf{x}$.
- Elliptical contours. In this case, the joint characteristic function is of the form

$$\phi(\mathbf{u}) = \exp\left(\left(\mathbf{u}^T Q \mathbf{u}\right)^{\alpha/2} + i\mathbf{u}\cdot\ \delta\right),$$

where Q is a $d \times d$ positive definite shape matrix and δ is a shift vector. A major advantage of this class is that it is computationally accessible and that joint dependence is characterized by the set of pairwise parameters, so $d(d-1)/2$ values are needed, just like in the Gaussian case.

BOX 1 CONTINUED.

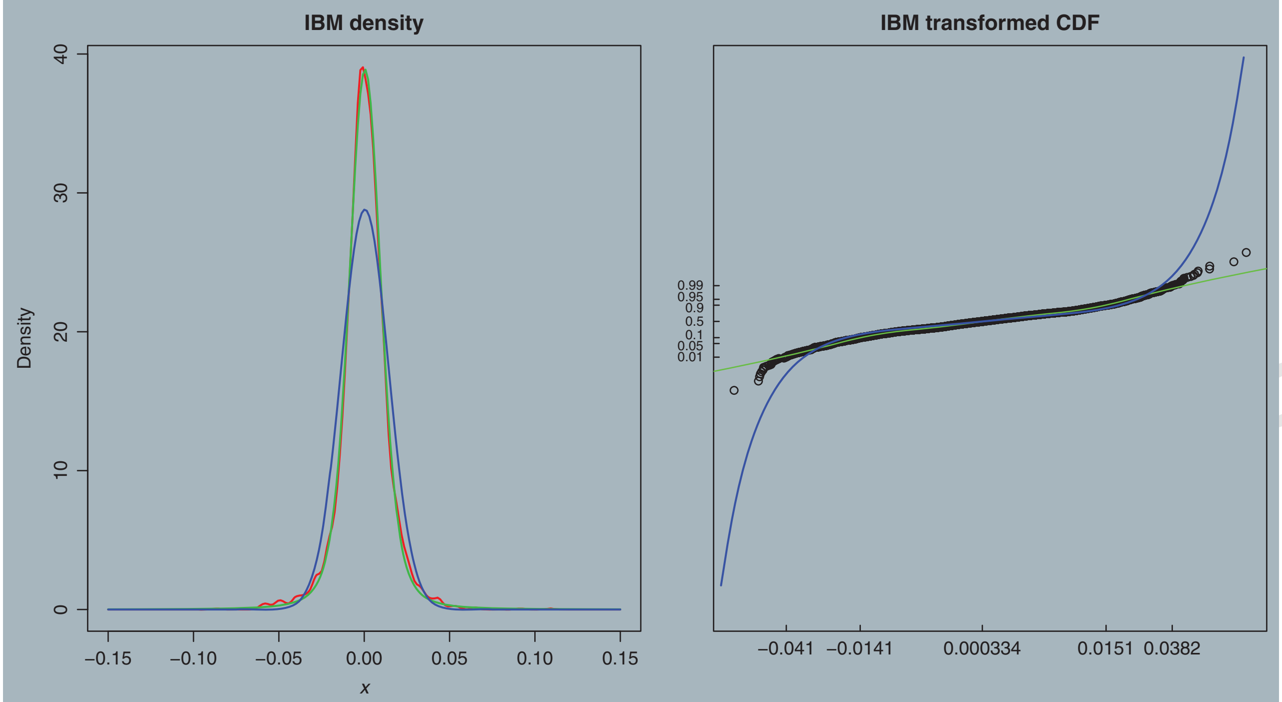

Comparison of the observed returns for IBM to a normal and a stable model. On the left is a plot the three densities: (a) a smoothed density plot in red, (b) a stable $S(\alpha = 1.614873, \beta = 0, \gamma = 0.007331, \delta = 0.000499)$ model in green, and (c) a normal $N(\mu = 0.000400, \sigma^2 = 0.000192)$ in blue. On the right are transformed cumulative distribution function plots for the returns (black circles), the stable model (green curve), and the normal model (blue curve).

In the center and midrange, the stable model describes the observed data much better than a normal model. Less obvious is that in the tails, the normal model significantly underestimates the tails of the data. The stable model tends to overestimate the extreme tails. In contrast, a normal model does a poor job over almost the whole range—the density is too low near the middle, too high in the midrange, and too low on the tails. While most of the attention to stable models has focused on the tails, it is most striking how the stable model approximates the data very well over most of the range, say from 1st to 99th percentile.

A look at the returns shows that they are not stationary; the financial crisis is clearly visible in the years 2008–2009. To deal with the changing volatility, we will apply a GARCH filter to the returns. For further information on estimation of stable-GARCH models, and in particular how they can be used in a multivariate framework suitable for portfolio optimization and risk prediction, see Ref 25 and the references therein. Using the function GARCH in R package TSERIES with a GARCH(1,1) model on the log returns yields the residuals shown in the following figure. These filtered returns are plausibly stationary. When maximum likelihood estimation is used, the estimated stable parameters along with 95% confidence intervals are: $\alpha = 1.825890 \pm 0.051727$, $\beta = -0.081668 \pm 0.253500$, $\gamma = 0.621666 \pm 0.021484$, $\delta = 0.043795 \pm 0.041979$. Note that the index of stability is now 1.82, up from 1.61, and the skewness is still close to 0. The differences in scale γ and location δ are because the GARCH filter changes the scale.

We briefly state some of the basic properties of multivariate stable laws. Sums of independent stable random vectors are stable, all univariate projections $\mathbf{u} \cdot \mathbf{X} = \sum_k u_k X_k$ are univariate stable laws, the support of a stable law is generally the whole space, but like in the one dimensional case, there are exceptions when $\alpha < 1$ and the spectral measure is one-sided. Plot (d) in Figure 2 shows an example where the support of the distribution is the cone bounded by the dashed lines.

 © 2013 Wiley Periodicals, Inc.

BOX 1 CONTINUED.

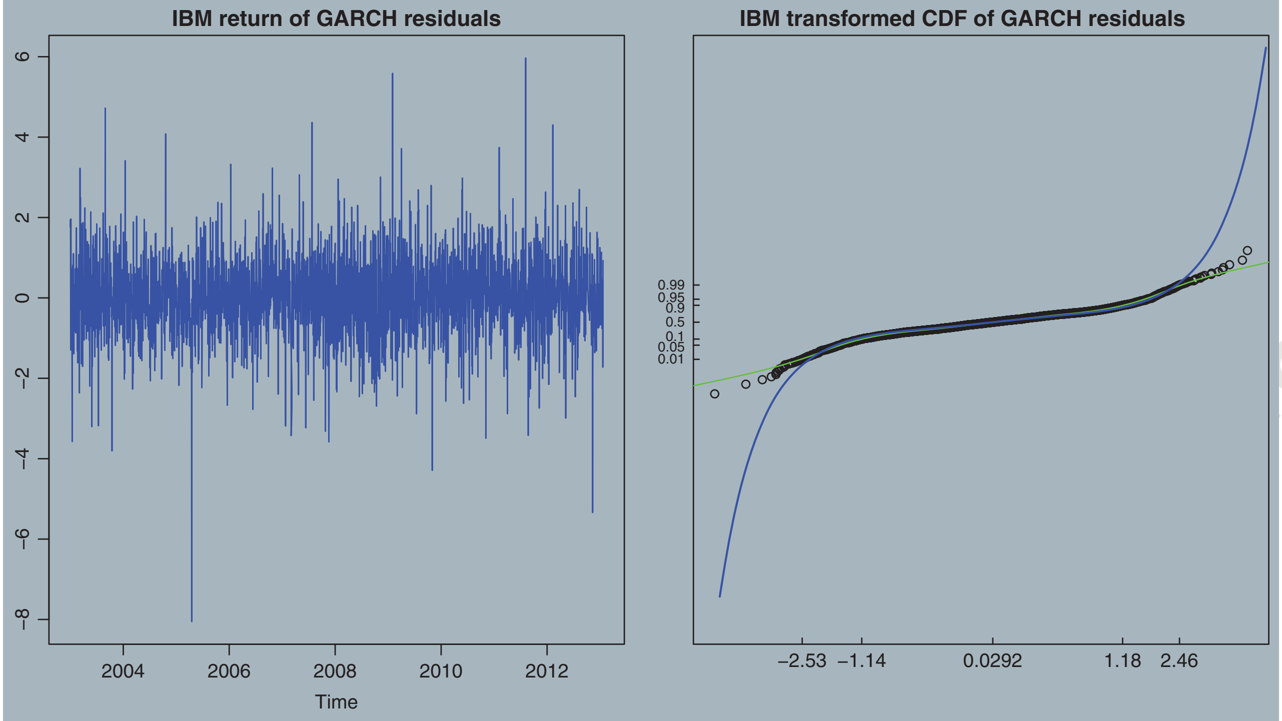

On the left are the residual returns for IBM stock after a GARCH(1,1) filter. On the right are transformed cumulative distribution function plots for the residuals from the GARCH model (black circles), the stable model (green curve), and the normal model (blue curve).

For the raw returns, the transformed CDF plot of the above figure shows that the stable model has heavier tails than the extremes of the data, while the normal model completely underestimates the tail probabilities. The GARCH residuals in the above figure shows that in addition to accounting for most of the changing volatility, the second stable fit captures the tails better. In both cases, the normal model completely misses the heavy tails of the data. Widespread underestimation of tails risks made investors, financial firms, and regulators sanguine about the market in the period leading up to the financial meltdown in 2008.

To be jointly stable, there has to be one α for which every component is univariate α-stable. In finance and other applications, there may be different components that have different tail behavior. In this case, one can fit each component with a univariate stable model, each having possibly different α's. A joint distribution can be constructed using copulas or vines, see McNeil et al.[7] or Kurowicka and Joe.[28] If multiple components have the same or similar index of stability, then it may make sense to use a joint stable model for those components, and then build a higher dimensional distribution out of these. An unfortunate consequence of these procedures is that the full distribution is generally not jointly stable.

Multivariate Computations, Simulation, and Estimation

At the current time, there is limited ability to compute densities and probabilities for multivariate stable laws. For discrete spectral measures, Nolan and Rajput[29] gave a program to compute bivariate densities $f(\mathbf{x})$ as in Figure 2. There are integral expressions for stable densities in Abdul-Hamid and Nolan[30] in higher dimensions, but they are complicated and difficult to evaluate numerically, and the difficulties increase as the dimension increases. Modarres and Nolan[31] give an algorithm to simulate stable random vectors with discrete Λ in any dimension. Cheng and Rachev,[32] Rachev and Xin,[33] and Nolan et al.[34] describe ways to estimate a discrete spectral measure. While the

BOX 2

A PORTFOLIO EXAMPLE

We examine a small portfolio with three assets: Ford (F), IBM, and Proctor and Gamble (PG). Adjusted closing prices were obtained for 10 years: January 1, 2003–December 31, 2012. As above, there were 2517 trading days in that period. As in the univariate example, there is changing volatility, most clearly in 2008, so we applied a GARCH(1,1) filter to the data. The residuals and the pairwise scatterplots of the residuals are shown in the following figures.

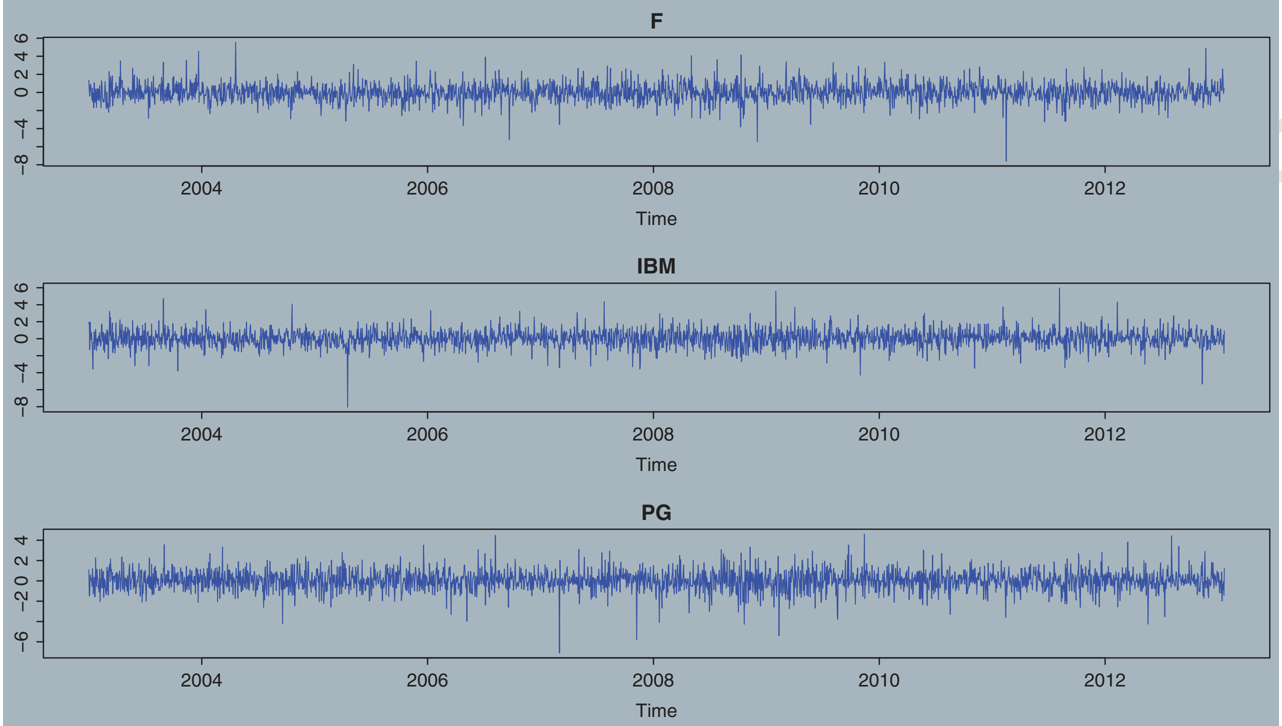

Residuals after GARCH(1,1) filtering of the three stocks over 10-year period 2003–2012: Ford (F), IBM, and Proctor and Gamble (PG).

methods work in higher dimensions, to our knowledge these have only been implemented in two-dimensions.

The elliptical case is much more accessible. Nolan[35] develops algorithms to evaluate the density for dimensions up to $d = 100$. There are also methods for simulating and estimating in arbitrary dimensions, all based on the $d \times d$ shape matrix Q. In finance, joint distributions of returns are frequently somewhat elliptically shaped. One intriguing feature of this class is that unlike the Gaussian case, there can be positive tail dependence when there are elliptical contours and $\alpha < 2$ (Box 2). This is due to Hult and Lindskog[36]; values for the tail dependence coefficient are tabulated in Ref 35.

CONCLUSION

Stable distributions are a flexible class of probability models that can be used in finance. They give a better fit to data over most of the range, with a overestimation of the extreme tails. In contrast, a normal model does a poor job of describing the data over most of the range, and radically underestimates the chance of extreme values. The use of stable models can be valuable for parties that are concerned about extreme losses, e.g., risk adverse investors, reinsurers, and regulators.

© 2013 Wiley Periodicals, Inc.

BOX 2 CONTINUED.

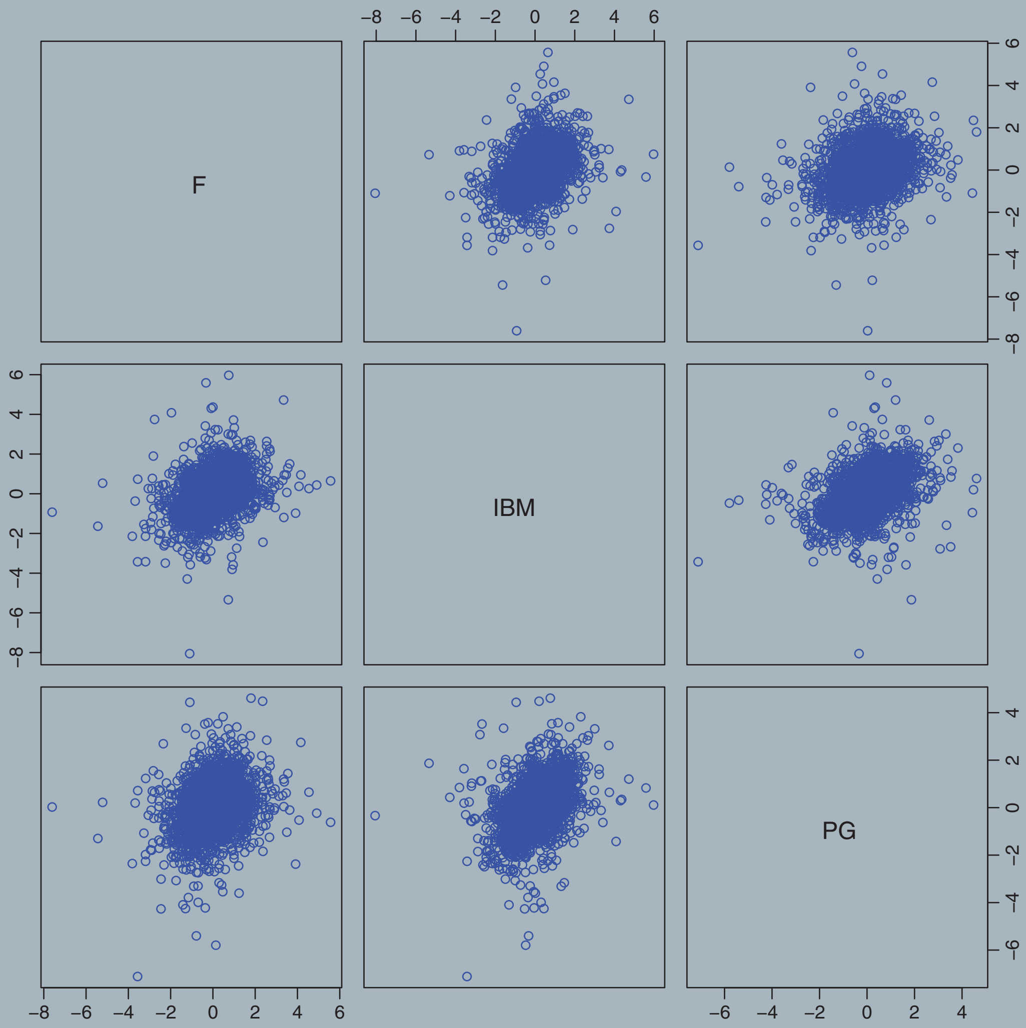

Pairwise scatterplots of the residuals after GARCH(1,1) filter for Ford (F), IBM, and Proctor and Gamble (PG).

The scatterplots are roughly elliptical and all three components of the residuals have $\alpha \approx 1.86$, so we will estimate a jointly stable three-dimensional elliptical model for the data. The residuals were fit with an elliptical stable model with index $\alpha = 1.86$ shift $\delta = (-0.0291, 0.0400, 0.0386)$, and shape matrix

$$Q = \begin{pmatrix} 0.3889 & 0.1871 & 0.1592 \\ 0.1871 & 0.3894 & 0.2040 \\ 0.1592 & 0.2040 & 0.3791 \end{pmatrix}.$$

ACKNOWLEDGMENTS

The author would like to thank Paul Embrechts for discussions about this topic, and the Forschungsinstitut für Mathematik at ETH Zurich for support during a visit. This publication was supported by an agreement with Cornell University, Operations Research & Information Engineering under W911NF-12-1-0385 from the Army Research Development and Engineering Command.

© 2013 Wiley Periodicals, Inc.

REFERENCES

1. Kaplan PD. *Frontiers of Modern• Asset Allocation*. Hoboken, NJ: John Wiley & Sons; 2012.
2. Box GEP, Draper NR. *Empirical Model Building and Response Surfaces*. New York: John Wiley & Sons; 1987.
3. Mandelbrot BB. The variation of certain speculative prices. *J Bus* 1963, 26:394–419.
4. Samuelson P. Efficient portfolio selection for Pareto-Lévy investments. *J Finance Quant Anal* 1967, 2:107–117.
5. Rachev ST, Mittnik S. *Stable Paretian Models in Finance*. New York: Wiley; 2000.
6. Embrechts P, Klüppelberg C, Mikosch T. *Modelling Extreme Events for Insurance and Finance*. Berlin: Springer-Verlag; 1997.
7. McNeil AJ, Frey R, Embrechts P. *Quantitative Risk Management*. Princeton Series in Finance. Princeton University Press: Princeton, NJ; 2005.
8. Lévy P. *Théorie de l'addition des variables aléatoires*. Paris: Gauthier-Villars; 1954 Original edition appeared in 1937.
9. Hall P. A comedy of errors: the canonical form for a stable characteristic function. *Bull Lond Math Soc* 1981, 13:23–27.
10. Zolotarev VM. *One-Dimensional Stable Distributions*. Translations of mathematical monographs, vol. 65. xxx: American Mathematical Society; 1986 Translation from the original 1983 Russian edition.•
11. Samorodnitsky G, Taqqu MS. *Stable Non-Gaussian Random Processes*. New York: Chapman and Hall; 1994.
12. Nolan JP. *Stable Distributions—Models for Heavy Tailed Data*. Birkhäuser, Boston. In preparation. Chapter 1 available at academic2.american.edu/jpnolan, along with other papers and free software for working with stable laws.
13. Fofack H, Nolan JP. Tail behavior, modes and other characteristics of stable distributions. *Extremes* 1999, 2:39–58.
14. Feller W. *An Introduction to Probability Theory and Its Applications*, vol. 2. 2nd ed. New York: Wiley; 1971.
15. Nolan JP. Numerical calculation of stable densities and distribution functions. *Commun Stat Stochastic Models* 1997, 13:759–774.
16. Chambers JM, Mallows CL, Stuck BW. A method for simulating stable random variables. *J Am Stat Assoc* 1976 Correction in *J Acoust Soc Am* 82 (1987), 704, 71:340–344.
17. Nikias CL, Shao M. *Signal Processing with α-Stable Distributions and Applications*. New York: Wiley; 1995.
18. Fama E, Roll R. Parameter estimates for symmetric stable distributions. *J Am Stat Assoc* 1971, 66:331–338.
19. McCulloch JH. Simple consistent estimators of stable distribution parameters. *Commun Stat Simul Comput* 1986, 15:1109–1136.
20. Koutrouvelis IA. Regression type estimation of the parameters of stable laws. *J Acoust Soc Am* 1980, 75:918–928.
21. Kogon SM, Williams DB. Characteristic function based estimation of stable parameters. In: Adler R, Feldman R, Taqqu M, eds. *A Practical Guide to Heavy Tailed Data*. Boston, MA: Birkhäuser; 1998, 311–338.
22. DuMouchel WH. On the asymptotic normality of the maximum-likelihood estimate when sampling from a stable distribution. *Ann Stat* 1973, 1:948–957.
23. Nolan JP. Maximum likelihood estimation of stable parameters. In: Barndorff-Nielsen OE, Mikosch T, Resnick SI, eds. *Lévy Processes: Theory and Applications*. Boston: Birkhäuser; 2001, 379–400.
24. Nolan JP, Ojeda-Revah D. Linear and nonlinear regression with stable errors. *J Econom* 2013, 172:186–194.
25. Broda SA, Haas M, Krause J, Paolella MS, Steude SC. Stable mixture GARCH models. *J Econom* 2013, 172:292–306.
26. Feldheim E. Etude de la stabilité des lois de probabilité. PhD thesis, Faculté des Sciences de Paris, Paris, France, 1937.
27. Byczkowski T, Nolan JP, Rajput B. Approximation of multidimensional stable densities. *J Multivar Anal* 1993, 46:13–31.
28. Kurowicka D, Joe H, eds. *Dependence Modeling: Vine Copula Handbook*. Hackensack, NJ: World Scientific Publishing Company; 2011.
29. Nolan JP, Rajput B. Calculation of multidimensional stable densities. *Commun Stat Simul* 1995, 24:551–556.
30. Abdul-Hamid H, Nolan JP. Multivariate stable densities as functions of one dimensional projections. *J Multivar Anal* 1998, 67:80–89.
31. Modarres R, Nolan JP. A method for simulating stable random vectors. *Comput Stat* 1994, 9:11–19.
32. Cheng BN, Rachev ST. Multivariate stable securities in financial markets. *Math Finance* 1995, 5:133–153.
33. Rachev ST, Xin H. Test for association of random variables in the domain of attraction of multivariate stable law. *Probab Math Stat* 1993, 14:125–141.
34. Nolan JP, Panorska A, McCulloch JH. Estimation of stable spectral measures. *Math Comput Model* 2001, 34:1113–1122.
35. Nolan JP. Multivariate elliptically contoured stable distributions: theory and estimation. *Comput Stat*. In press.
36. Hult H, Lindskog F. Multivariate extremes, aggregation and dependence in elliptical distributions. *Adv Appl Probab* 2002, 34:587–608.

UNCORRECTED PROOFS

QUERIES TO BE ANSWERED BY AUTHOR

IMPORTANT NOTE: Please mark your corrections and answers to these queries directly onto the proof at the relevant place. DO NOT mark your corrections on this query sheet.

Queries from the Copyeditor:

AQ1. We have provided this standard statement as per journal style. Please modify only if you have any conflict to declare, else retain the same.

AQ2. Please note that as per journal style figure are not numbered inside boxes. So we have renumbered figures and modified the citation text. Kindly confirm.

AQ3. Please confirm the inserted running title.

AQ4. Please confirm the inserted citation for Boxes 1 and 2.

AQ5. Please note that we have renumbered the references to maintain sequential order. Kindly confirm.

AQ6. Please provide place of publication for Ref 10.

UNCORRECTED PROOFS

WILEY-BLACKWELL

USING e-ANNOTATION TOOLS FOR ELECTRONIC PROOF CORRECTION

Required software to e-Annotate PDFs: Adobe Acrobat Professional or Adobe Reader (version 7.0 or above). (Note that this document uses screenshots from Adobe Reader X)
The latest version of Acrobat Reader can be downloaded for free at: http://get.adobe.com/uk/reader/

Once you have Acrobat Reader open on your computer, click on the Comment tab at the right of the toolbar:

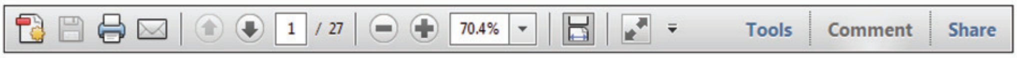

This will open up a panel down the right side of the document. The majority of tools you will use for annotating your proof will be in the Annotations section, pictured opposite. We've picked out some of these tools below:

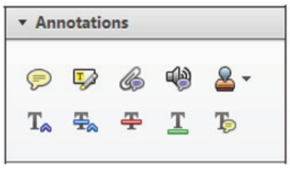

1. Replace (Ins) Tool – for replacing text.

Strikes a line through text and opens up a text box where replacement text can be entered.

How to use it

- Highlight a word or sentence.
- Click on the Replace (Ins) icon in the Annotations section.
- Type the replacement text into the blue box that appears.

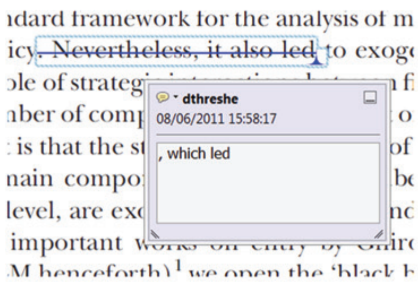

2. Strikethrough (Del) Tool – for deleting text.

Strikes a red line through text that is to be deleted.

How to use it

- Highlight a word or sentence.
- Click on the Strikethrough (Del) icon in the Annotations section.

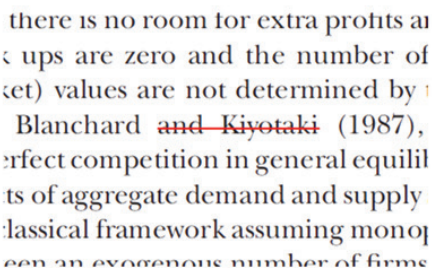

3. Add note to text Tool – for highlighting a section to be changed to bold or italic.

Highlights text in yellow and opens up a text box where comments can be entered.

How to use it

- Highlight the relevant section of text.
- Click on the Add note to text icon in the Annotations section.
- Type instruction on what should be changed regarding the text into the yellow box that appears.

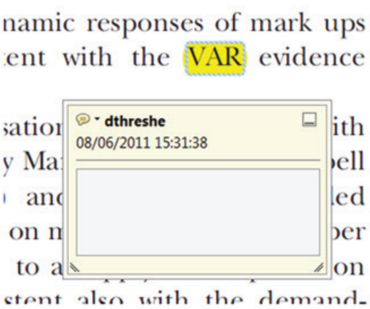

4. Add sticky note Tool – for making notes at specific points in the text.

Marks a point in the proof where a comment needs to be highlighted.

How to use it

- Click on the Add sticky note icon in the Annotations section.
- Click at the point in the proof where the comment should be inserted.
- Type the comment into the yellow box that appears.

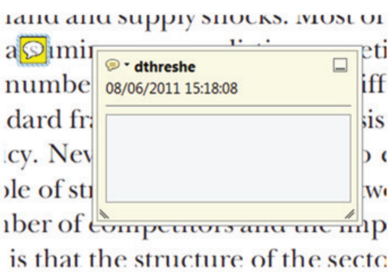

WILEY-BLACKWELL

USING e-ANNOTATION TOOLS FOR ELECTRONIC PROOF CORRECTION

5. Attach File Tool – for inserting large amounts of text or replacement figures.

Inserts an icon linking to the attached file in the appropriate pace in the text.

How to use it

- Click on the Attach File icon in the Annotations section.
- Click on the proof to where you'd like the attached file to be linked.
- Select the file to be attached from your computer or network.
- Select the colour and type of icon that will appear in the proof. Click OK.

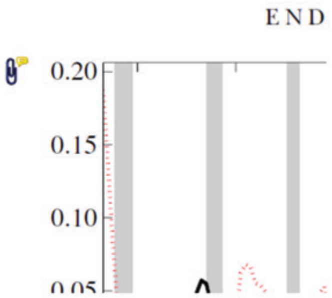

6. Add stamp Tool – for approving a proof if no corrections are required.

Inserts a selected stamp onto an appropriate place in the proof.

How to use it

- Click on the Add stamp icon in the Annotations section.
- Select the stamp you want to use. (The Approved stamp is usually available directly in the menu that appears).
- Click on the proof where you'd like the stamp to appear. (Where a proof is to be approved as it is, this would normally be on the first page).

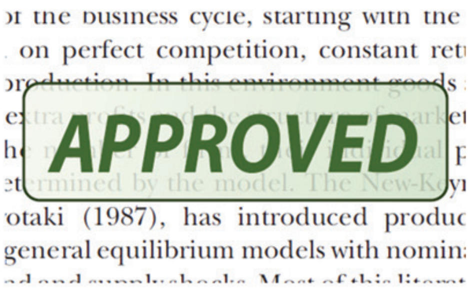

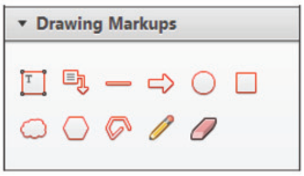

7. Drawing Markups Tools – for drawing shapes, lines and freeform annotations on proofs and commenting on these marks.

Allows shapes, lines and freeform annotations to be drawn on proofs and for comment to be made on these marks..

How to use it

- Click on one of the shapes in the Drawing Markups section.
- Click on the proof at the relevant point and draw the selected shape with the cursor.
- To add a comment to the drawn shape, move the cursor over the shape until an arrowhead appears.
- Double click on the shape and type any text in the red box that appears.

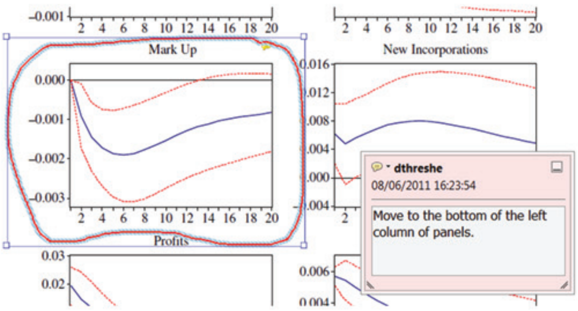

For further information on how to annotate proofs, click on the Help menu to reveal a list of further options: