Phi-alpha Optimal Portfolios and Extreme Risk Management

R. Douglas Martin, Professor of Statistics, University of Washington and Executive Chairman, FinAnalytica.
Svetlozar (Zari) Rachev, Professor of Statistics and Economics, U.C. Santa Barbara, University of Karlsruhe, (Germany, and Chief Scientist, FinAnalytica.
Frederic Siboulet, Executive Vice President of Sales & Marketing, FinAnalytica.

Abstract

We introduce a practical alternative to Gaussian risk factor distributions based on Svetlozar Rachev's extensive work on Stable distributions in Finance (see Rachev and Mittnik, 2000), and called the ϕ_α Paradigm. In contrast to Normal distributions, Stable distributions capture the fat tails and the asymmetries of real-world risk factor distributions. In addition, we make use of copulas, a generalization of overly restrictive linear correlation models, to account for the dependencies between risk factors during extreme events, and multivariate ARCH-type processes with stable innovations to account for joint volatility clustering. We demonstrate that the application of these techniques results in more accurate modeling of extreme risk event probabilities, and consequently delivers more accurate risk measures for both trading and risk management. Using these superior models, VaR becomes a much more reliable measure of downside risk. More importantly Stable Expected Tail Loss (SETL) can be accurately calculated and used as a more informative risk measure for both market and credit portfolios. Along with being a superior risk measure, SETL enables an elegant approach to portfolio optimization via convex optimization that can be solved using standard scalable linear programming software. We show that this approach, called ϕ_α portfolio optimization, yields superior risk adjusted returns relative to Markowitz portfolios. We introduce a new indicator of investment performance π_α that quantifies the excess mean return of ϕ_α optimal portfolios compared to Markowitz's mean-variance optimal portfolios. Finally, we introduce a novel investment performance measurement tool: the Stable Tail Adjusted Return Ratio (STARR), which is a generalization of the Sharpe ratio in ϕ_α.

"When anyone asks me how I can describe my experience of nearly forty years at sea, I merely say uneventful. Of course there have been winter gales and storms and fog and the like, but in all my experience, I have never been in an accident of any sort worth speaking about. I have seen but one vessel in distress in all my years at sea (. . .) I never saw a wreck and have never been wrecked, nor was I ever in any predicament that threatened to end in disaster of any sort."

E.J. Smith, Captain, 1907, RMS Titanic

1 High Market Volatility Demands New Solutions

Paul Wilmott likes to recount the ritual by which he questions his students on the likelihood of Black Monday 1987. Under the commonly accepted Gaussian risk factor distribution assumption, they calculate that there should be no such event in the entire existence of the universe and beyond!

The last two decades have witnessed a considerable increase in the volatility of financial markets—dramatically so in the last few years. Extreme events are the corollary of that increased volatility.

Legacy risk systems have done a reasonable job at managing *ordinary* financial events. However up to now, very few institutions or vendors have demonstrated the systematic ability to deal with the *unusual* event, the one that should *almost never* happen. Therefore, one can reasonably question the soundness of some of the current risk management practices and tools used in Wall Street as far as extreme risk is concerned.

The two principal approaches to modeling asset returns are based either on Historical or on Normal (Gaussian) distribution. Neither

approach adequately captures unusual asset price and return behaviors. The *Historical* model is bounded by the scope of the available observations and the *Normal* model inherently cannot produce atypical returns. The financial industry is beleaguered with both under-optimized portfolios with often-shabby ex-post risk-adjusted returns, as well as deceptive aggregate risk indicators (e.g. VaR) that lead to substantial unexpected losses.

The inadequacy of the Normal distribution is well recognized by the risk management community. Yet up to now, no consistent and comprehensive alternative probability models had adequately addressed unusual returns. To quote one major vendor:

"It has often been argued that the true distributions returns (even after standardizing by the volatility) imply a larger probability of extreme returns than that implied from the Normal distribution. Although we could try to specify a distribution that fits returns better, it would be a daunting task, especially if we consider that the new distribution would have to provide a good fit across all asset classes." (Technical Manual, RMG, 2001)

In response to the challenge, we use Stable risk-factor distributions and generalized risk-factor dependencies, thereby creating a *paradigm shift* to consistent and uniform use of the most viable class of non-Normal probability models in finance. This approach leads to distinctly improved financial risk management and portfolio optimization solutions for highly volatile markets with extreme events.

2 The Stable Distribution Framework

2.1 Stable Distributions
In spite of wide-spread awareness that most risk factor distributions are heavy-tailed, to date, risk management systems have essentially relied either on historical, or on univariate and multivariate Normal (or Gaussian) distributions for Monte Carlo scenario generation. Unfortunately, historical scenarios only capture conditions actually observed in the past, and in effect use empirical probabilities that are zero outside the range of the observed data, a clearly undesirable feature. On the other hand Gaussian Monte Carlo scenarios have probability densities that converge to zero too quickly (exponentially fast) to accurately model real-world risk factor distributions that generate extreme losses. When such large returns occur separately from the bulk of the data they are often called outliers.

The figure below shows quantile–quantile (qq)-plots of daily returns versus the best-fit Normal distribution of nine randomly selected micro-cap stocks for the two-year period 2000–2001. If the returns were Normally distributed, the quantile points in the qq-plots would all fall
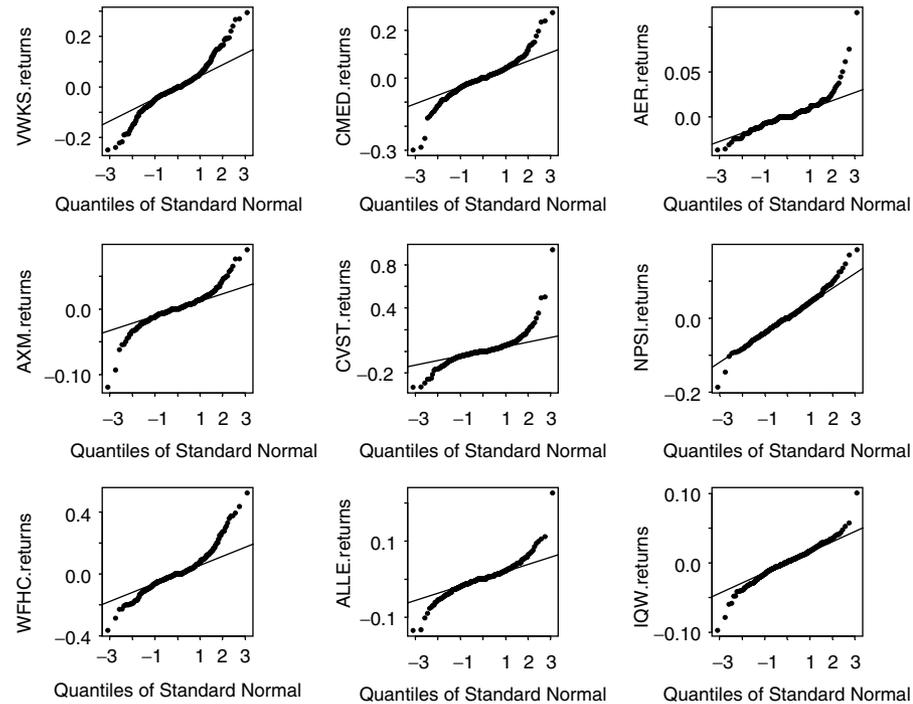
close to a straight line. Instead they all deviate significantly from a straight line (particularly in the tails), reflecting a higher probability of occurrence of extreme values than predicted by the Normal distribution, and showing several outliers.

Such behavior occurs in many asset and risk factor classes, including well-known indices such as the S&P 500, and corporate bond prices. The latter are well known to have quite non-Gaussian distributions that have substantial negative skews to reflect down-grading and default events. For such returns, non-Normal distribution models are required to accurately model the tail behavior and compute probabilities of extreme returns.

Various non-Normal distributions have been proposed for modeling extreme events, including:

- Mixtures of two or more Normal distributions,
- *t*-distributions, hyperbolic distributions, and other scale mixtures of normal distributions
- Gamma distributions,
- Extreme Value distributions,
- Stable non-Gaussian distributions (also known as Lévy-Stable and Pareto-Stable distributions).

Among the above, only Stable distributions have attractive enough mathematical properties to be a viable alternative to Normal distributions in trading, optimization and risk management systems. A major drawback of all alternative models is their lack of stability. Benoit

Mandelbrot (1963) demonstrated that the stability property is highly desirable for asset returns. These advantages are particularly evident in the context of portfolio analysis and risk management.

An attractive feature of Stable models, not shared by other distribution models, is that they allow generation of Gaussian-based financial theories and, thus allow construction of a coherent and general framework for financial modeling. These generalizations are possible only because of specific probabilistic properties that are unique to (Gaussian and non-Gaussian) Stable laws, namely; the Stability property, the Central Limit Theorem, and the Invariance Principle for Stable processes.

Benoit Mandelbrot (1963), then Eugene Fama (1965), provided seminal evidence that Stable distributions are good models for capturing the heavy-tailed (leptokurtic) returns of securities. Many follow-on studies came to the same conclusion, and the overall Stable distributions theory for finance is provided in the definitive work of Rachev and Mittnik (2000).

But in spite the convincing evidence, Stable distributions have seen virtually no use in capital markets. There have been several barriers to the application of stable models, both conceptual and technical:

- Except for three special cases, described below, Stable distributions have no closed form expressions for their probability densities.
- Except for Normal distributions, which are a limiting case of Stable distributions (with $\alpha = 2$ and $\beta = 0$), Stable distributions have infinite variance and only a mean value for $\alpha > 1$.
- Without a general expression for stable probability densities, one cannot directly implement maximum likelihood methods for fitting these densities, even in the case of a single (univariate) set of returns.

The availability of practical techniques for fitting univariate and multivariate stable distributions to asset and risk factor returns has been *the* barrier to the progress of Stable distributions in finance. Only the recent development of advanced numerical methods has removed this obstacle. These patented methods form the foundation of the **Cognity**™ market & credit risk management and portfolio optimization solution (see further comments in section 4.7).

Univariate Stable Distributions

A stable distribution for a random risk factor X is defined by its characteristic function:

$$F(t) = E\left(e^{itX}\right) = \int e^{itx} f_{\mu,\sigma}(x)dx,$$

where

$$f_{\mu,\sigma}(x) = \frac{1}{\sigma} f\left(\frac{x - \mu}{\sigma}\right)$$

is any probability density function in a location-scale family for X:

$$\log F(t) = \begin{cases} -\sigma^\alpha |t|^\alpha \left(1 - i\beta \operatorname{sgn}(t) \tan\left(\frac{\pi\alpha}{2}\right)\right) + i\mu t, & \alpha \neq 1 \\ -\sigma |t| \left(1 - i\beta \frac{2}{\pi} \operatorname{sgn}(t) \log |t|\right) + i\mu t, & \alpha = 1 \end{cases}$$

A stable distribution is therefore determined by the four key parameters:

1. α determines density's kurtosis with $0 < \alpha \leq 2$ (e.g. tail weight)
2. β determines density's skewness with $-1 \leq \beta \leq 1$
3. σ is a scale parameter (in the Gaussian case, $\alpha = 2$ and $2\sigma^2$ is the variance)
4. μ is a location parameter (μ is the mean if $1 < \alpha \leq 2$)

Symmetric PDFs
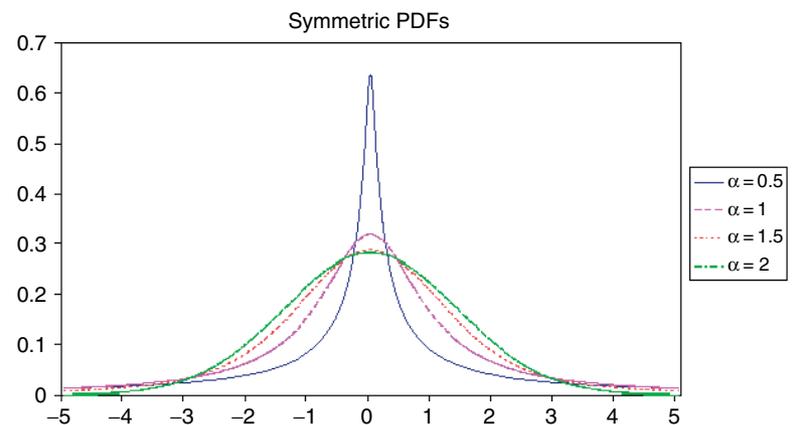
Stable distributions for risk factors allow for skewed distributions when $\beta \neq 0$ and fat tails relative to the Gaussian distribution when $\alpha < 2$. The graph above shows the effect of α on tail thickness of the density as well as peakedness at the origin relative to the Normal distribution (collectively the "kurtosis" of the density), for the case of $\beta = 0$, $\mu = 0$, and $\sigma = 1$. As the values of α decrease the distribution exhibits fatter tails and more peakedness at the origin.

The case of $\alpha = 2$ and $\beta = 0$ and with the reparameterization in scale, $\tilde{\sigma} = \sqrt{2}\sigma$, yields the Gaussian distribution, whose density is given by:

$$f_{\mu,\tilde{\sigma}}(x) = \frac{1}{\sqrt{2\pi}\tilde{\sigma}} e^{-\frac{(x - \mu)^2}{2\tilde{\sigma}^2}}.$$

Skewed PDFs ($\alpha = 1.5$)
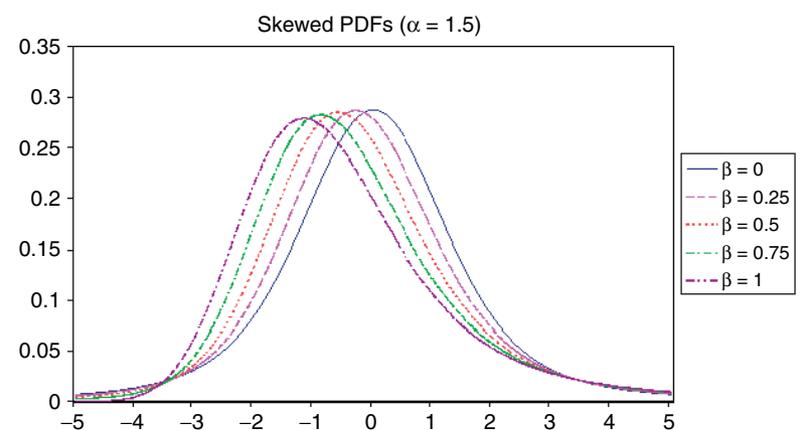
The case $\alpha = 1$ and $\beta = 0$ yields the Cauchy distribution with much fatter tails than the Gaussian, and is given by:

$$f_{\mu,\sigma}(x) = \frac{1}{\pi \cdot \sigma} \left(1 + \left(\frac{x - \mu}{\sigma} \right)^2 \right)^{-1}$$

The figure below illustrates the influence of β on the skewness of the density for $\alpha = 1.5$, $\mu = 0$ and $\sigma = 1$. Increasing (decreasing) values of β result in skewness to the right (left).

Fitting Stable and Normal Distributions: DJIA Example

Aside from the Gaussian, Cauchy, and one other special case of stable distribution for a positive random variable with $\alpha = 0.5$, there is no closed form expression for the probability density of a Stable random variable.

Thus one is not able to directly estimate the parameters of a Stable distribution by the method of maximum likelihood. To estimate the four parameters of the stable laws, the *Cognity*™ solution uses a special patent-pending version of the FFT (Fast Fourier Transform) approach to numerically calculate the densities with high accuracy, and then applies MLE (Maximum Likelihood Estimation) to estimate the parameters.

The results from applying the *Cognity*™ Stable distribution modeling to the DJIA daily returns from January 1, 1990 to February 14, 2003 is displayed in the figure below. In both cases a GARCH model has been used to account for the clustering of volatility.

The figure shows the left-hand tail detail of the resulting stable density, along with that of a Normal density fitted using the sample mean and sample standard deviation, and that of a non-parametric kernel density estimate (labeled "Empirical" in the plot legend). The parameter estimates are:

- Stable parameters $\hat{\alpha} = 1.699$, $\hat{\beta} = -.120$, $\hat{\mu} = .0002$, and $\hat{\sigma} = .006$,
- Normal density parameter estimates $\hat{\mu} = .0003$, and $\hat{\sigma} = .010$.
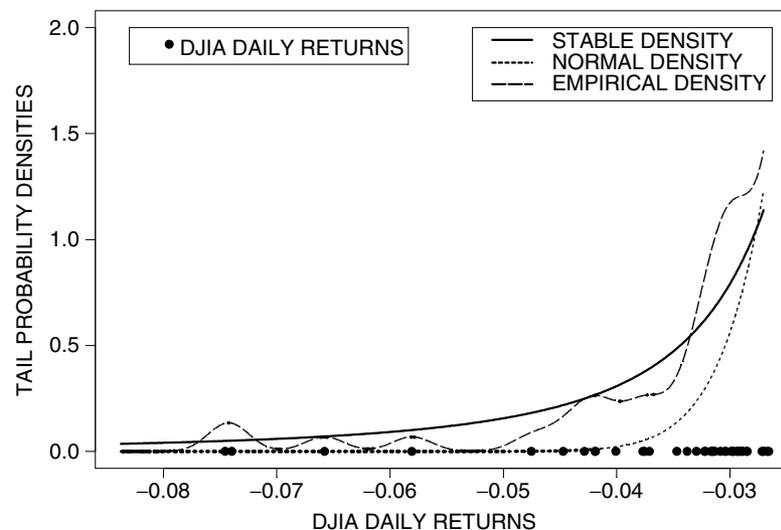
Note that the Stable density tail behavior is reasonably consistent with the Empirical non-parametric density estimate, indicating the existence of some extreme returns. At the same time it is clear from the figure that the tail of the Normal density is much too thin, and will provide inaccurate estimates of tail probabilities for the DJIA returns. The table below shows just how bad the Normal tail probabilities are for several negative returns values.

PROBABILITY (DJIA RETURN < *x*)			
x	−0.04	−0.05	−0.06	−0.07
Stable Fit	0.0066	0.0043	0.0031	0.0023
Normal Fit	0.000056	0.0000007	3.68E-09	7.86E-12

A daily return smaller than −0.04 with the Stable distribution occurs with probability 0.0066, or roughly seven times every four years, whereas such a return with the Normal fit occurs on the order of once every four years.

Similarly, a return smaller than −0.05 with the Stable occurs about once per year and with the Normal fit about once every forty years. Clearly the Normal distribution fit is an exceedingly optimistic predictor of DJIA tail return values.

The figure below displays the central portion of the fitted densities as well as the tails, and shows that the Normal fit is not nearly peaked enough near the origin as compared with the empirical density estimate (even though the GARCH model was applied), while the stable distribution matches the empirical estimate quite well in the center as well as in the tails.
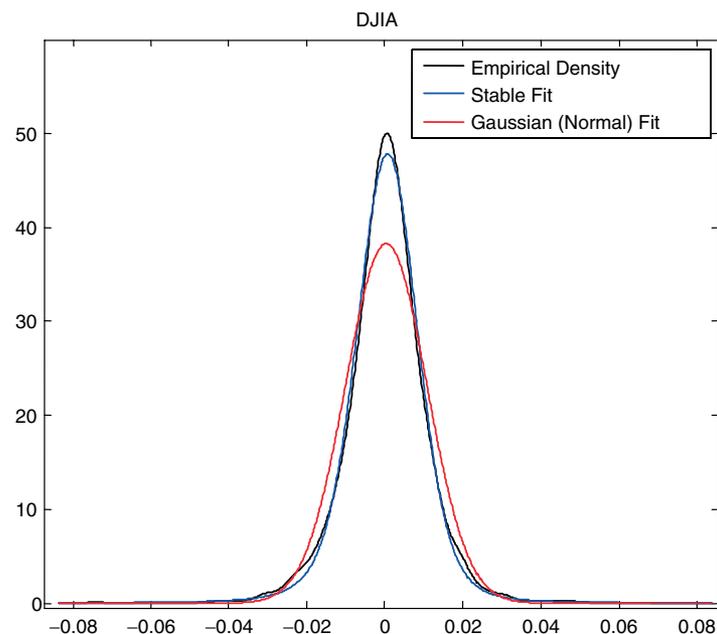
Fitting Stable Distributions: Micro-Caps Example

Noting that micro-cap stock returns are consistently strongly non-normal (see sample of normal qq-plots at the beginning of this section), we fit stable distributions to a random sample of 182 micro-cap daily returns for the two-year period 2000–2001. The results are displayed in the boxplot below.

ESTIMATED ALPHAS OF 182 MICRO-CAP STOCKS

ESTIMATED STABLE PARAMETER ALPHAS
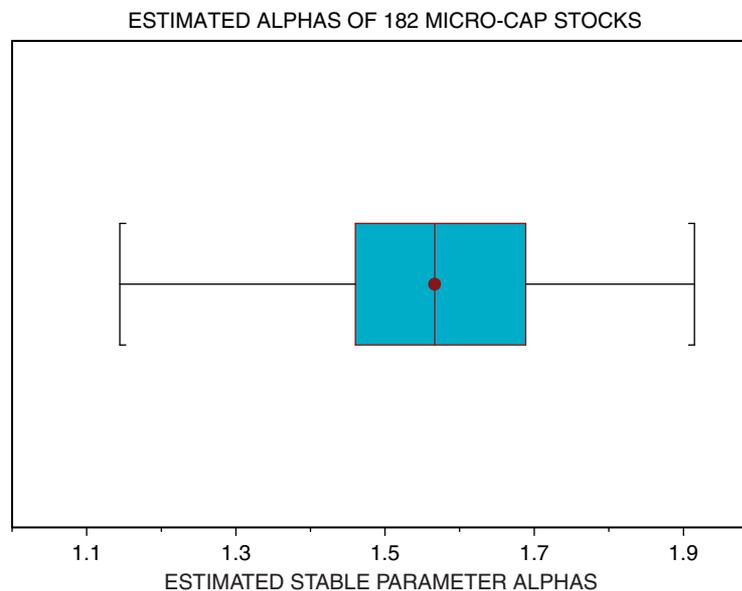
The median of the estimated alphas is 1.57, and the upper and lower quartiles are 1.69 and 1.46 respectively. Somewhat surprisingly, the distribution of the estimated alphas turns out to be quite Normal.

Multivariate Stable Distribution Modeling

Multivariate Stable distribution modeling involves univariate Stable distributions for each risk factor, each with its own parameter estimates $\hat{\alpha}_i$, $\hat{\beta}_i$, $\hat{\mu}_i$, $\hat{\sigma}_i$, $i = 1, 2, \cdots, K$, where K is the number of risk factors, along with a dependency structure.

One way to produce the dependency structure is through a subordinated process approach as follows. First compute a robust mean vector and covariance matrix estimate of the risk factors by trimming a small percentage of the observations on a coordinate-wise basis (to get rid of the outliers, and have a good covariance estimate for the central bulk of the data). Next you generate multivariate normal scenarios with this mean vector and covariance matrix. Then you multiply each of random variable component of the scenarios by a Stable subordinator which is a strictly positive Stable random variable with index $\hat{\alpha}_i/2, i = 1, 2, \cdots, K$. The vector of subordinators is usually independent of the normal scenario vectors, but it can also be dependent. See for example Rachev S. and Mittnik S. (2000), and Rachev S., Schwartz E. and Khindanova I. (2003).

Another very promising approach to building the cross-sectional dependence model is through the use of copulas, an approach that is

quite attractive because it allows for modeling higher correlations during extreme market movements, thereby accurately reflecting lower portfolio diversification at such times. The next section briefly discussion copulas.

2.2 Copula Multivariate Dependence Models

Why Copulas?

Classical correlations and covariances are quite limited measures of dependence, and are only adequate in the case of multivariate Gaussian distributions. A key failure of correlations is that, for non-Gaussian distributions, zero correlation does not imply independence, a phenomenon that arises in the context of time-varying volatilities represented by ARCH and GARH models. The reason we use copulas is that we need more general models of dependence, ones which:

- Are not tied to the elliptical character of the multivariate normal distribution
- Have multivariate contours and corresponding data behavior that reflect the local variation in dependence that is related to the level of returns, in particular, those shapes that correspond to higher dependence for extreme values of two or more of the returns.

What are Copulas?

A copula may be defined as a multivariate cumulative distribution function with uniform marginal distributions:

$$C(u_1, u_2, \cdots, u_n), \ u_i \in [0, 1] \text{ for } i = 1, 2, \cdots, n$$

where

$$C(u_i) = u_i \text{ for } i = 1, 2, \cdots, n.$$

It is known that for any multivariate cumulative distribution function:

$$F(x_1, x_2, \cdots, x_n) = P(X_1 \leq x_1, X_2 \leq x_2, \cdots X_n \leq x_n)$$

there exists a copula C such that

$$F(x_1, x_2, \cdots, x_n) = C(F_1(x_1), F_2(x_2), \cdots, F_n(x_n))$$

where the $F_i(x_i)$ are the marginal distributions of $F(x_1, x_2, \cdots, x_n)$, and conversely for any copula C the right-hand-side of the above equation defines a multivariate distribution function $F(x_1, x_2, \cdots, x_n)$. See for example, Bradley and Taqqu (2001) and Embrects et. al. (2003).

The main idea behind the use of copulas is that one can first specify the marginal distributions in whatever way makes sense, e.g. fitting marginal distribution models to risk factor data, and then specify a copula C to capture the multivariate dependency structure in the best suited manner.

There are many classes of copula, particularly for the special case of bivariate distributions. For more than two risk factors beside the traditional

Gaussian copula, the t-copula is very tractable for implementation and provides a possibility to model dependencies of extreme events. It is defined as:

$$C_{\nu,\mathbf{c}}(u_1, u_2, \cdots, u_n) = \frac{\Gamma((\nu + n)/2)}{\Gamma(\nu/2)\sqrt{|\mathbf{c}|(\nu\pi)^n}} \int_{-\infty}^{t_\nu^{-1}(u_1)} \cdots \int_{-\infty}^{t_\nu^{-1}(u_n)} \left(1 + \frac{\mathbf{s}\,\mathbf{c}^{-1}\mathbf{s}}{\nu}\right) d\mathbf{s}$$

where **c** is a correlation matrix.

A sample of 2000 bivariate simulated risk factors generated by a t-copula with 1.5 degrees of freedom and Normal marginal distributions is displayed in the figure below.

T-COPULA WITH 1.5 DOF AND NORMAL MARGINALS
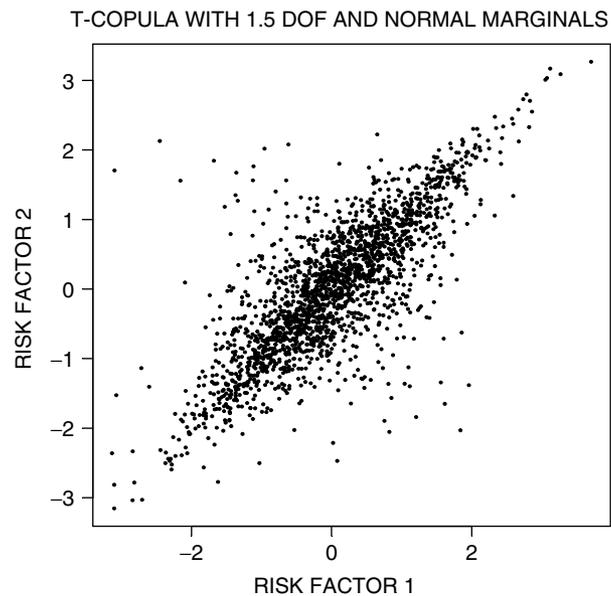
The example illustrates that these two risk factors are somewhat uncorrelated for small to moderately large returns, but are highly correlated for the infrequent occurrence of very large returns. This can be seen by noting that the density contours of points in the scatter plot are somewhat eliptical near the origin, but are nowhere close to eliptical for more extreme events. This situation is in contrast to a Gaussian linear dependency relationship where the density contours are expected to be eliptical.

2.3 Volatility Modeling and Stable Versus Normal VaR

It is well known that risk factors returns exhibit volatility clustering, and that even after adjusting for such clustering the returns will still be non-normal and contain extreme values. There may also be some serial dependency effects to account for. In order to adequately model these collective behaviors we recommend using ARIMA models with an ARCH/GARCH "time-varying" volatility input, where the latter has Stable (non-Gaussian) innovations. This approach is more flexible and accurate than the commonly used simple EWMA (exponentially weighted moving average) volatility model, and provides accurate time-varying

estimates of VaR and Exptected Tail Loss (ETL) risk measures. See 3 section for discussion of ETL vs. VaR that emphasizes the advantages of ETL. However, we stress that those who must use VaR to satisfy regulatory requirements will get much better results with Stable VaR than with Normal VaR, as the following example vividly shows.

Consider the following portfolio of Brady bonds:

- Brazil C 04/14
- Brazil EIB 04/06
- Venezuela DCB Floater 12/07
- Samsung KRW Ord Shares
- Thai Farmers Bank THB

We have run Normal, Historical and Stable 99% (1% tail probability) VaR calculations for one-year of daily data from January 9, 2001 to January 9, 2002. We used a moving window with 250 historical observations for the Normal VaR model, 500 for the historical VaR model and 700 for the Stable VaR model. For each of these cases we used a GARCH(1,1) model for volatility clustering of the risk factors, with Stable innovations. We backtested these VaR calculations by using the VaR values as one-step ahead predictors, and got the results shown in the figure below.
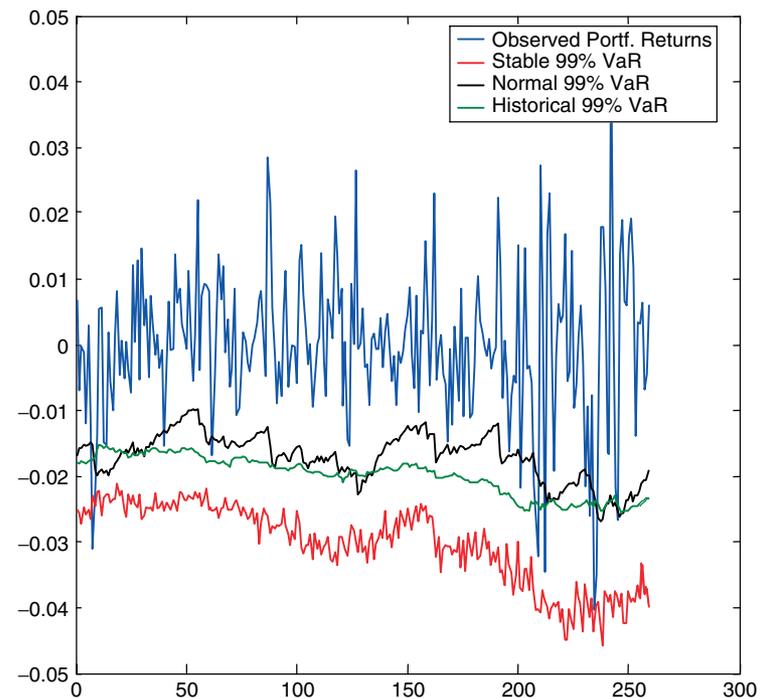
The figure shows: the returns of the Brady bond portfolio (top curve); the Normal+EWMA (a la RiskMetrics) VaR (curve with jumpy behavior, just below the returns); the Historical VaR (the smoother curve mostly below but sometimes crossing the Normal+EWMA VaR); the Stable+GARCH VaR (the bottom curve). The results with regard to exceedances of the 99% VaR, and keeping in mind Basel II guidelines, may be summarized as follows:

- Normal 99% VaR produced 12 exceedances (red zone)
- Historical 99% VaR produced 9 exceedances (on upper edge of yellow zone)
- Stable 99% VaR produced 1 exceedence (solidly in the green zone)

Clearly Stable (+GARCH) 99% VaR produces much better results with regard to Basel II compliance. This comes at the price of higher initial capital reserves, but results in a much safer level of capital reserves and a very clean bill of health with regard to compliance. Note that organizations in the red zone may be fined and will have to increase their capital reserves by 33%, which at some times for some portfolios will result in larger capital reserves than when using the Stable VaR, this in addition to being viewed as having inadequate risk measures. This unfortunate situation is much less likely to happen to the organization using Stable VaR.

3 ETL is the Next Generation Risk Measure

3.1 Why Not Value-at-Risk (VaR)?

There is no doubt that VaR's popularity is in large part due to its simplicity and its ease of calculation for 1 to 5% confidence levels. However, there is a price to be paid for the simplicity of VaR in the form of several limitations:

- VaR does not give any indication of the risk beyond the quantile!
- VaR portfolio optimization is a non-convex, non-smooth problem with multiple local minima that can result in portfolio composition discontinuities. Furthermore it requires complex calculation techniques such as integer programming.
- VaR is not sub-additive; i.e. the VaR of the aggregated portfolio can be larger the sum of the VaR's of the sub-portfolios.
- Finally, and most importantly, VaR can be a very misleading risk indicator: examples exist where an investor, unintentionally or not, decreases portfolio VaR while simultaneously increasing the expected losses beyond the VaR, i.e., by increasing the "tail risk" of the portfolio (see the discussion of ETL below).

In addition to these intrinsic limitations, specific VaR implementations are fraught with further flaws:

- Historical VaR limits the range of the scenarios to data values that have actually been observed, while Normal Monte Carlo tends to seriously underestimate the probability of extreme returns. In either case, the probability functions beyond the sample range are either zero or excessively close to zero.
- Lacking the ability to accurately model extreme returns, practitioners are forced to use stress testing as a palliative to the limitations of traditional VaR models. In doing so, they use a large degree of subjectivity in the design of the stress test and in the selection of past data to use in making a risk assessment.

- The traditional modeling of risk factor dependences cannot account for intraday volatility patterns, long-term volatility patterns, or more importantly unusual extreme volatility. In stressed markets, the simple linear diversification assumptions fail, and atypical short-term concentration patterns that bind all the assets in a bearish spiral emerge.

Yamai and Yoshiba (2002) note in their concluding remarks: "The widespread use of VaR for risk management could lead to market instability. Basak and Shapiro (2001) show that when investors use VaR for their risk management, their optimizing behavior may result in market positions that are subject to extreme loss because VaR provides misleading information regarding the distribution's tail."

## 3.2 ETL and Stable Versus Normal Distributions

Expected Tail Loss (ETL) is simply the average (or expected value) loss for losses larger than VaR. ETL is also known as *Conditional Value-at-Risk* (CVaR), or *Expected Shortfall* (ES).

Usual (1 to 5%) Normal ETL is close to Normal VaR (See VaR by Jorion, 2001 p98):

- For CI = 5%, VaR = 1.645 and ETL = 2.062
- For CI = 1%, VaR = 2.336 and ETL = 2.667

By failing to capture kurtosis, Normal distributions underestimate ETL. The ubiquitous Normal assumption makes ETL difficult to interpret, in spite of ETL's remarkable properties (see below). Unlike Normal distributions, Stable distributions capture leptokurtic tails ("fat tails"). Unlike Normal ETL, Stable ETL provides reliable values. Further, when Stable distributions are used, ETL is generally substantially different from VaR.

The graph below shows the time series of daily returns for the stock OXM from January 2000 to December 2001. Observe the occurrences of extreme values.

OXM Returns
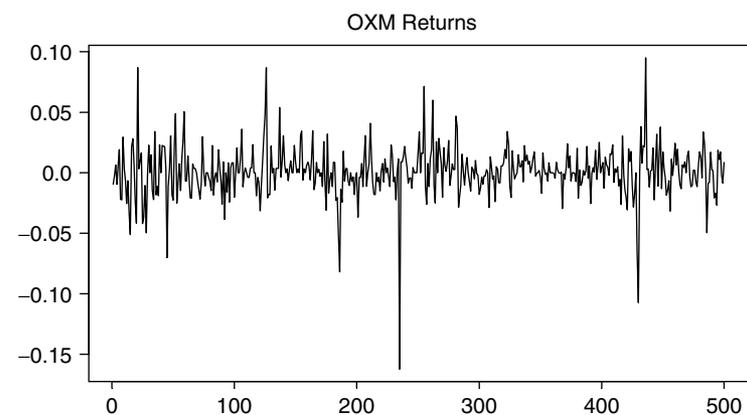
While this series also displays obvious volatility clustering, that deserves to be modeled as described in section 3.3, we shall ignore this aspect for the moment. Rather, here we provide a compelling example of

the difference between ETL and VaR based on a well-fitting stable distribution, as compared with a poor fitting Normal distribution.

The figure below shows a histogram of the OXM returns with a Normal density fitted using the sample mean and sample standard deviation, and a Stable density fitted using maximum-likelihood estimates of the Stable distribution parameters. The Stable density is shown by the solid line and the normal density is shown by the dashed line. The former is obviously a better fit than the latter, when using the histogram of the data values as a reference. The estimated Stable tail thickness index is $\hat{\alpha} = 1.62$. The 1% VaR values for the Normal and Stable fitted densities are .047 and .059 respectively, a ratio of 1.26 which reflects the heavier-tailed nature of the Stable fit.

99% VAR FOR NORMAL AND STABLE DENSITIES


99% ETL FOR NORMAL AND STABLE DENSITIES
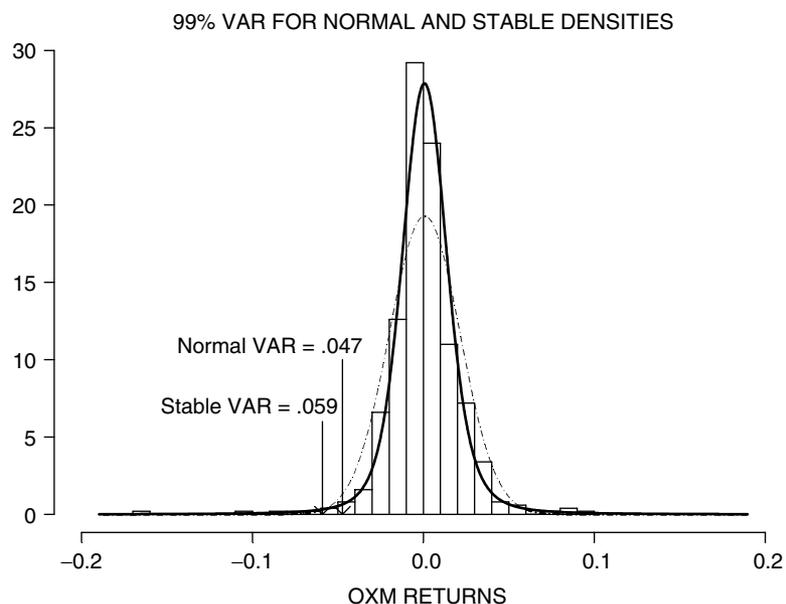
The figure below displays the same histogram and fitted densities with 1% ETL values instead of the 1% VaR values. The 1% ETL values for the Normal and Stable fitted densities are .054 and .174 respectively, a ratio of a little over three-to-one. This larger ratio is due to the Stable density's heavy tail contribution to ETL relative to the Normal density fit.

Unlike VaR, ETL has a number of attractive properties:

- ETL gives an informed view of losses beyond VaR
- ETL is a convex, smooth function of portfolio weights, and is therefore attractive to optimize portfolios (see Uryasev & Rockafellar, 2000). This point is vividly illustrated in the subsection below on ETL and Portfolio Optimization.
- ETL is sub-additive and satisfies a complete set of coherent risk measure properties (see Artzner et. al., 1999)
- ETL is a form of expected loss (i.e. a conditional expected loss) and is a very convenient form for use in scenario-based portfolio optimization.
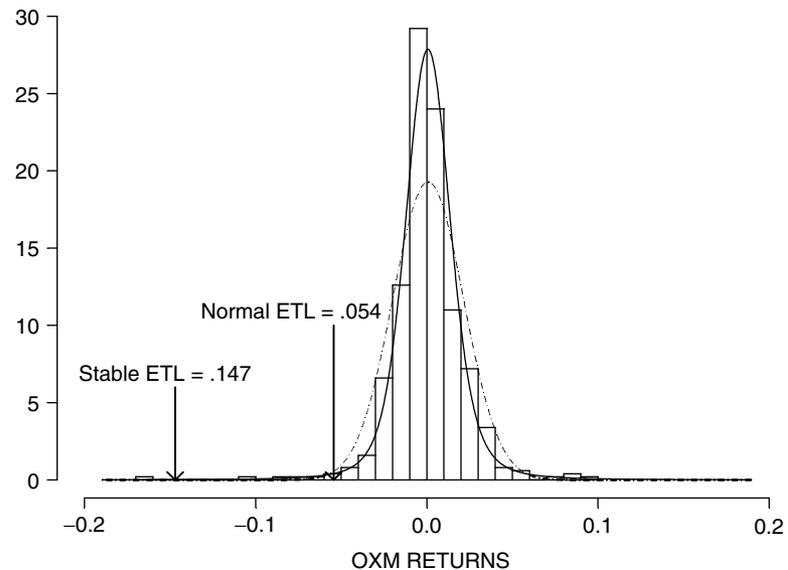
It is also quite a natural risk-adjustment to expected return (see STARR, or Stable Tail Adjusted Return Ratio).

The limitations of current Normal risk factor models and the absence of regulator blessing have held back the widespread use of ETL, in spite of its highly attractive properties.

For portfolio optimization, we recommend the use of Stable ETL, and limiting the use of Historical, Normal or Stable VaR to regulatory reporting purposes. At the same time, organizations should consider the advantages of Stable ETL for risk assessment purposes and non-regulatory reporting purposes.

The following quotation is relevant: "Expected Tail Loss gives an indication of extreme losses, should they occur. Although it has not become a standard in the financial industry, expected tail loss is likely to play a major role, as it currently does in the insurance industry" (Embrechts et. al., 1997).

3.3 Portfolio Optimization and ETL Versus VaR

To the surprise of many, portfolio optimization with ETL turns out to be a smooth, convex problem with a unique solution (Rockafellar and Uryasev, 2000). These properties are in sharp contrast to the non-convex, rough VaR optimization problem.

The contrast between VAR and ETL portfolio optimization surfaces is illustrated with the following pair of figures for a simple two-asset portfolio. The horizontal axes show one of the portfolio weights (from 0 to 100%) and the vertical axes display portfolio VAR and ETL respectively. The data consist of 200 simulated uncorrelated returns.

The VAR objective function is quite rough with respect to varying the portfolio weight(s), while that of the ETL objective function is smooth and convex. One can see that optimizing with ETL is a much more tractable problem than optimizing with VaR.

NON-CONVEX ROUGH VAR SURFACE

CONVEX SMOOTH ETL SURFACE

COGNITY STABLE VAR OPTIMAL PORTFOLIOS
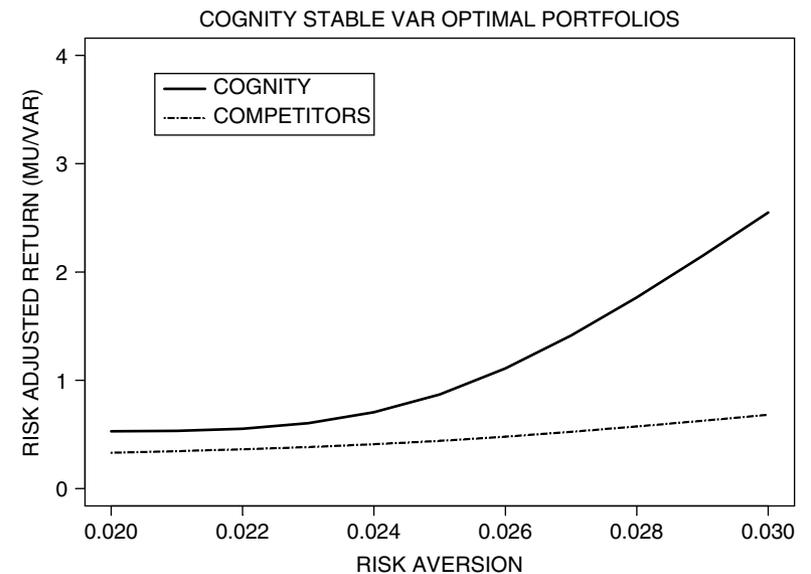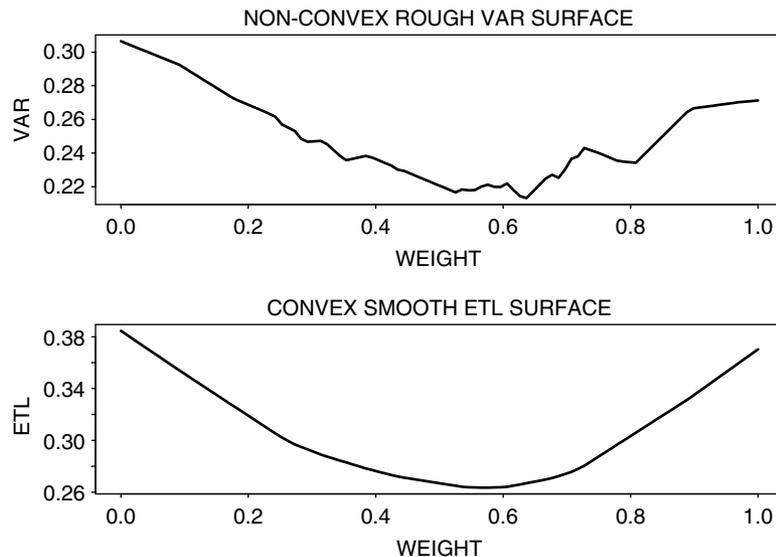
Rockafellar and Uryasev (2000), show that the ETL Optimal Portfolio (ETLOP) weight vector can be obtained based on historical (or scenario) returns data by minimizing a relatively simple convex function (Rockafellar and Uryasev used the term CVaR whereas we use the, less confusing, synonym ETL). Assuming p assets with single period returns $\mathbf{r}_i = (r_{i1}, r_{i2}, \cdots, r_{ip})$, for period i, and a portfolio weight vector $\mathbf{w} = (w_1, w_2, \ldots, w_p)$, the function to be minimized is

$$F(\mathbf{w}, \gamma) = \gamma + \frac{1}{\varepsilon \cdot n} \sum_{i=1}^{n} [\mathbf{w}'\mathbf{r}_i - \gamma]^+.$$

where $[x]^+$ denotes the positive part of x. This function is to be minimized jointly with respect to $\mathbf{w}$ and γ, where ε is the tail probability for which the expected tail loss is computed. Typically $\varepsilon = .05$ or $.01$, but larger values may be useful, as we discuss in section 4.6. The authors further show that this optimization problem can be cast as a LP (linear programming) problem, solvable using any high-quality LP software.

Cognity™ combines this approach along with multivariate Stable scenario generation. The stable scenarios provide accurate and well-behaved estimates of ETL for the optimization problem.

3.4 Stable ETL Leads to Higher Risk Adjusted Returns

ETLOP (Expected Tail Loss Optimal Portfolio) techniques, combined with multivariate Stable distribution modeling, can lead to significant improvements in risk adjusted return as compared to not only Normal VAROP methods but also compared to Normal ETL optimization. In practice, a VAROP is difficult to compute accurately with more than two or three assets.

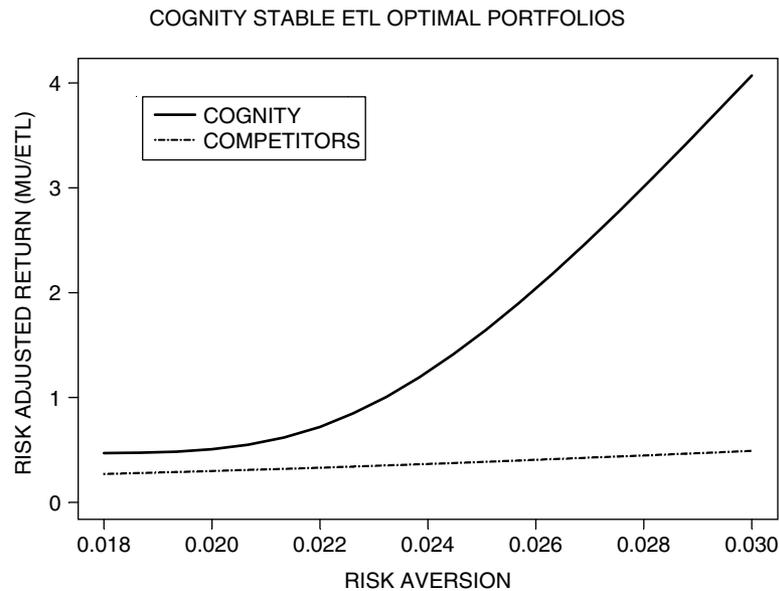The figures below are supplied to illustrate the claim that Stable ETLOP produces consistently better risk-adjusted returns. These figures show the risk adjusted return MU/VAR (mean return divided by VAR) and MU/ETL (mean return divided by ETL) for 1% VAROP and ETLOP, and using a multi-period fixed-mix optimization in all cases.

In this simple example, the portfolio to be optimized consists of two assets, cash and the S&P 500. The example is based on monthly data from February 1965 to December 1999. Since we assume full investment, the VAROP depends only on a single portfolio weight and the optimal weight(s) is found by a simple grid search on the interval 0 to 1. The use of a grid search technique, overcomes the problems with non-convex and non-smooth VAR optimization. In this example the optimizer is maximizing $MU - c \cdot VAR$ and $MU - c \cdot ETL$, where c is the risk aversion (parameter), and with VAR or ETL as the penalty function.

The first plot shows that even using VAROP, one gets a significant relative gain in risk-adjusted return using Stable scenarios when compared to Normal scenarios, and with the relative gain increasing with increasing risk aversion. The reason for the latter behavior is that with Stable distributions the optimization pays more attention to the S&P returns distribution tails, and allocates less investment to the S&P under stable distributions than under Normal distributions as risk aversion increases.

The figure below for the risk-adjusted return for the ETLOP has the same vertical axis range as the previous plot for VAROP. The figure below shows that the use of ETL results in much greater gain under the Stable distribution relative to the Normal than in the case of VAROP.

At every level of risk aversion, the investment in the S&P 500 is even less in ETLOP than in the case of the VAROP. This behavior is to be expected because the ETL approach pays attention to the losses beyond VAR (the expected value of the extreme loss), and which in the Stable case are much greater than in the Normal case.

COGNITY STABLE ETL OPTIMAL PORTFOLIOS

P2) Risk factor returns are associated through models that describe the dependence between the individual factors more accurately than classical correlations. Often these will be copula models.

P3) Risk factor modeling typically includes a ϕ_α-econometric model in the form of multivariate ARIMA-GARCH processes with residuals driven by fractional Stable innovations. The ϕ_α econometric model captures clustering and long-range dependence of the volatility.

4.2 Phi-Alpha (ϕ_α) Optimal Portfolios

A Phi-Alpha (ϕ_α) optimal portfolio is one that minimizes portfolio expected tail loss (ETL) subject to a constraint of achieving expected portfolio returns at least as large as an investor defined level, where both quantities are evaluated in ϕ_α. Alternatively, a ϕ_α optimal portfolio solves the dual problem of maximizing portfolio expected return subject to a constraint that portfolio expected tail loss (ETL) is not greater than an investor defined level, where again both quantities are evaluated in ϕ_α. In order to define the above ETL precisely we use the following quantities:

R_p:	the random return of portfolio p
SER_p:	the ϕ_α expected return of portfolio p
$L_p = -R_p + SER_p$:	the loss of portfolio p relative to its ϕ_α expected return
ε:	a tail probability of the ϕ_α distribution L_p
$SVaR_p(\varepsilon)$:	the ϕ_α Value-at-Risk for portfolio p

The latter is defined by the equation

$$\Pr[L_p > SVaR_p(\varepsilon)] = \varepsilon$$

where the probability is calculated in ϕ_α, that is $SVaR_p(\varepsilon)$ is the ε-quantile of the ϕ_α distribution of L_p. In the value-at-risk literature $(1 - \varepsilon) \times 100\%$ is called the confidence level. Here we prefer to use the simpler, unambiguous term *tail probability*. Now we define the ϕ_α expected tail loss as

$$SETL_p(\varepsilon) = E[L_p \,|\, L_p > SVaR_p(\varepsilon)]$$

where the conditional expectation is also computed in ϕ_α. Note that the ϕ_α expected value of L_p is zero. We use the "S" in SER_p, $SVaR_p(\varepsilon)$ and $SETL_p(\varepsilon)$ as a reminder that Stable distributions are a key aspect of the ϕ_α (but not the only aspect!).

Proponents of (Gaussian) VaR typically use tail probabilities of .01 or .05. When using $SETL_p(\varepsilon)$ risk managers may wish to use other tail probabilities such as .1, .15, .20, .25, or .5. We note that use of different tail probabilities is similar in spirit to using different utility functions. We return to discuss this point further in section 4.xx.

The following assumptions are in force for the ϕ_α investor:

A1) The universe of assets is Q (the set of mandate admissible portfolios)

A2) The investor may borrow or deposit at the risk-free rate r_f without restriction

4 The ϕ_α Paradigm and ϕ_α Optimal Portfolios

4.1 The Phi-Alpha (ϕ_α) Paradigm

Our approach uses multi-dimensional risk factor models based on multivariate Stable process models for risk management and constructing optimal portfolios, and stresses the use of Stable ETL as the risk measure of choice. These Stable distribution models incorporate generalized dependence structure with copulas, and include time varying volatilities based on GARCH models with Stable innovations. Collectively these modeling foundations form the basis of a new, powerful overall basis for investment decisions that we call the *Phi-Alpha (ϕ_α) Paradigm*.

Currently the ϕ_α Paradigm has the following basic components: ϕ_α scenario engines, ϕ_α integrated market risk and credit risk (with integrated operational risk under development), ϕ_α optimal portfolios and efficient frontiers, and ϕ_α derivative pricing. Going forward, additional classes of ϕ_α investment decision models will be developed, such as ϕ_α betas, ϕ_α factor models, and ϕ_α asset liability models. The rich structure of these models will encompass the heavy-tailed distributions of the asset returns, stochastic trends, heteroscedasticity, short- and long-range dependence, and more. We use the term "ϕ_α model" to describe any such model in order to keep in mind the importance of the Stable tail-thickness parameter α in financial investment decisions.

It is essential to keep in mind the following ϕ_α fundamental principles concerning risk factors:

P1) Risk factor returns have Stable distributions where each risk factor i typically has a different Stable tail-index α_i.

A3) The portfolio is optimized under a set of asset allocation constraints λ

A4) The investor seeks an expected return of at least μ (alternatively an ETL risk of at most η).

To simplify the notation we shall let A3 be implicit in the following discussion. At times we shall also suppress the ε when its value is taken as fixed and understood.

The ϕ_α investor's optimal portfolio is

$$\omega_\alpha(\mu|\varepsilon) = \arg \min_{q \in Q} SETL_q(\varepsilon)$$

subject to

$$SER_q \geq \mu.$$

In other words the ϕ_α optimum portfolio ω_α minimizes the ϕ_α expected tail loss (SETL) among all portfolios with ϕ_α mean return (SER) at least μ, for fixed tail probability ε asset allocation constraints λ. Alternatively, the ϕ_α optimum portfolio ω_α solves the dual problem

$$\omega_\alpha(\eta|\varepsilon) = \arg \max_{q \in Q} SER_q$$

subject to

$$SETL_q(\varepsilon) \leq \eta.$$

The ϕ_α *efficient frontier* is given by $\omega_\alpha(\mu|\varepsilon)$ as a function of μ for fixed ε, as indicated in the figure below. If the portfolio includes cash account with risk free rate r_f, then ϕ_α *efficient frontier* will be the ϕ_α capital market line (CML_α) that connects the risk-free rate on the vertical axis with the ϕ_α tangency portfolio (T_α), as indicated in the figure.
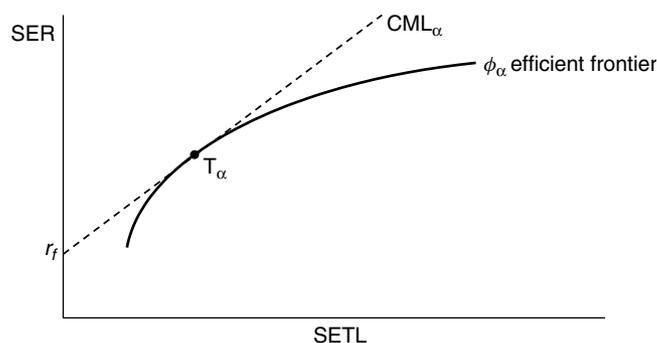
We now have a ϕ_α separation principal analogous to the classical separation principal: The tangency portfolio T_α can be computed without reference to the risk-return preferences of any investor. Then an investor chooses a portfolio along the ϕ_α capital market line CML_α according to his/her risk-return preference.

We note that it is convenient to think of ω_α in two alternative ways: (1) the ϕ_α optimal portfolio, or (2) the vector of ϕ_α optimal portfolio weights.

Keep in mind that in practice when a finite sample of returns one ends up with a ϕ_α efficient frontier, tangency portfolio and capital market line that are estimates of true values for these quantities. Under regularity conditions these estimates will converge to true values as the sample size increases to infinity.

4.3 Markowitz Portfolios are Sub-Optimal

While the ϕ_α investor has a ϕ_α optimal portfolios described above, let's assume that the Markowitz investor is not aware of the ϕ_α Paradigm and constructs a mean-variance optimal portfolio. We assume that the Markowitz investor operates under the same assumptions A1–A4 as the ϕ_α investor. Let ER_q be the expected return and σ_q the standard deviation of the returns of a portfolio q. The Markowitz investor's optimal portfolio is

$$\omega_2(\mu) = \arg \min_{q \in Q} \sigma_q$$

subject to

$$ER_q \geq \mu.$$

The Markowitz optimal portfolio can also be constructed by solving the obvious dual optimization problem of maximizing the expected return for a constrained risk level. One knows that, in the mean-variance paradigm, contrary to the ϕ_α paradigm, the mean-variance optimal portfolio is independent of any ETL tail probability specification.

The subscript 2 is used in ω_2 as a reminder that $\alpha = 2$ you have the limiting Gaussian distribution member of the Stable distribution family, and in that case the Markowitz portfolio is optimal. Alternatively you can think of the subscript 2 as a reminder that the Markowitz optimal portfolio is a second-order optimal portfolio, i.e., an optimal portfolio based on only first and second moments.

Note that the Markowitz investor ends up with a different portfolio, i.e., a different set of portfolio weights with different risk versus return characteristics, than the ϕ_α investor.

The performance of the Markowitz portfolio, like that of the ϕ_α portfolio, is evaluated under the ϕ_α distributional model, i.e., its expected return and expected tail loss are computed under the ϕ_α distributional model. If in fact the distribution of the returns were exactly multivariate Gaussian (which they never are) then the ϕ_α investor and the Markowitz investor would end up with one and the same optimal portfolio. However, when the returns are non-Gaussian ϕ_α returns, the Markowitz portfolio is sub-optimal. This is because the ϕ_α investor constructs his/her optimal portfolio using the ϕ_α distribution model, whereas the Markowitz investor does not. Thus the Markowitz investor's frontier lies below and to the right of the ϕ_α efficient frontier, as shown in the figure below, along with the Markowitz tangency portfolio T_2 and Markowitz capial market line CML_2.
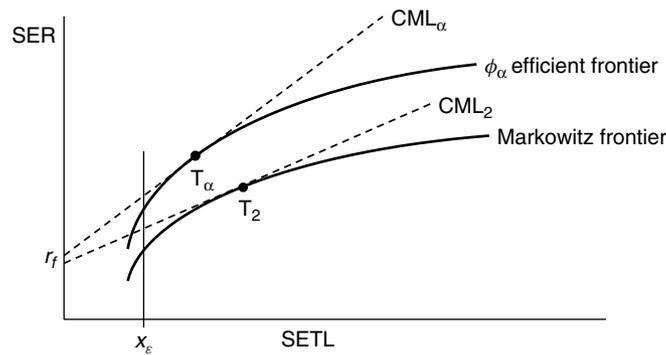
As an example of the performance improvement achievable with the ϕ_α optimal portfolio approach, we computed the ϕ_α efficient frontier and the Markowitz frontier for a portfolio of 47 micro-cap stocks with the smallest alphas from the random selection of 182 micro-caps in section 21. The results are displayed in the figure below. The results are based on 3,000 scenarios from the fitted ϕ_α multivariate distribution model based on two years of daily data during years 2000 and 2001. We note that, as is generally the case, each of the 47 stock returns has its own estimate Stable tail index $\hat{\alpha}_i, \ i = 1, 2, \ldots, 47$.
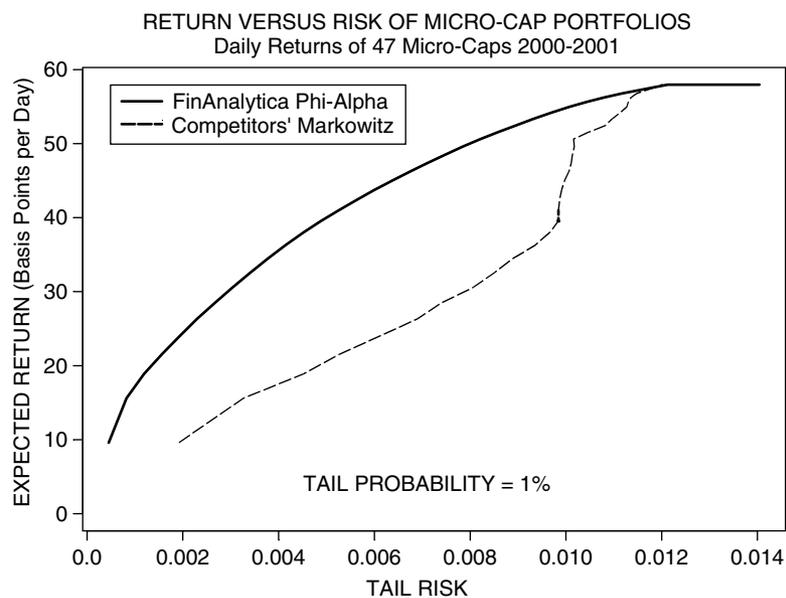
Here we have plotted values of *TailRisk* $= \varepsilon \cdot SETL(\varepsilon)$, for $\varepsilon = .01$, as a natural decision theoretic risk measure, rather than $SETL(\varepsilon)$ itself. We note that over a large range of tail risk the ϕ_α efficient frontier dominates the Markowitz frontier by 14–20 bp daily!

We note that the 47 micro-caps with the smallest alphas used for this example have quite heavy tails as indicated by the boxplot of their estimated alphas shown below.
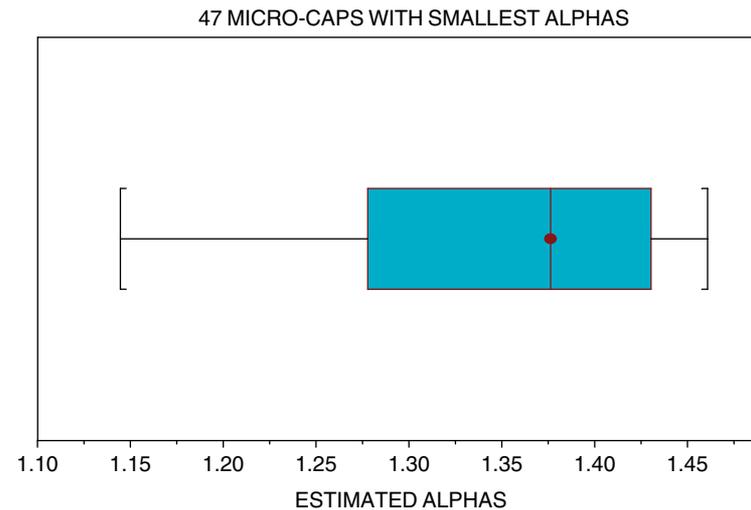
Here the median of the estimated alphas is 1.38, while the upper and lower quartiles are 1.43 and 1.28 respectively. Evidently there is a fair amount of information in the non-Gaussian tails of such micro-caps that can be exploited by the ϕ_α approach.

We also note that the gap between the ϕ_α efficient frontier and the Markowitz mean-variance frontier will decrease as the Stable tail index values α_i get closer to 2, i.e., as the multivariate distribution gets closer to a multivariate Normal distribution. This will be the case for example when moving from micro-cap and small-cap stocks to mid-cap and large cap stocks.

4.4 The Excess Profit Pi-Alpha

Let $x_\varepsilon = SETL(\varepsilon)$ be a fixed value of Stable expected tail loss, or alternatively let $x_\varepsilon = \varepsilon \cdot SETL(\varepsilon)$ be a fixed value of Stable tail risk. Also, let $SER_{\omega_\alpha}(x_\varepsilon)$ be the ϕ_α expected return of the corresponding ϕ_α optimal portfolio, and let $SER_{\omega_2}(x_\varepsilon)$ be the ϕ_α expected return of the corresponding Markowitz mean-variance optimal portfolio. We define the excess profit (Pi-Alpha) of the ϕ_α portfolio in basis points as

$$\pi_\alpha = \pi_\alpha(\varepsilon, x_\varepsilon) = [SER_{\omega_\alpha}(x_\varepsilon) - SER_{\omega_2}(x_\varepsilon)] * 100.$$

In the case of portfolios with no cash component the values of Stable expected return above are the ordinate values where the vertical line at x_ε intersect the two frontiers (see the figure above). When cash is present these expected returns are the values where the vertical line at x_ε intersects the capital market lines CML_α and CML_2.

Based on our current studies, yearly π_α in the range of tens to possibly one or two hundred basis points are achievable, depending upon the portfolio under management.

4.5 New Ratios: From Sharpe to STARI and STARR

The *Sharpe Ratio* for a given portfolio p is defined as follows:

$$SR_p = \frac{ER_p - r_f}{\sigma_p} \qquad (2)$$

where ER_p is the portfolio expected return, σ_p is the portfolio return standard deviation as a measure of portfolio risk, and r_f is the risk-free rate. While the Sharpe ratio is the single most widely used portfolio performance measure, it has several disadvantages due to its use of the standard deviation as risk measure:

- σ_p is a symmetric measure that does not focus on downside risk,
- σ_p is not a coherent measure of risk (see Artzner et. al., 1999),
- the classical estimate of σ_p is a highly unstable measure of risk when the portfolio has a heavy-tailed distribution,
- σ_p and has infinite value for non-Gaussian Stable distributions.

Stable Tail Adjusted Return Indicator

As an alternative performance measure that does not suffer these disadvantages, we propose the *Stable Tail Adjusted Return Indicator* (*STARI*) defined as:

$$STARI_p(\varepsilon) = \frac{SER_p - r_f}{SETL_p(\varepsilon)}. \qquad (3)$$

Referring to the first figure in section 4.3, one sees that the overall maximum *STARI* is attained by the ϕ_α optimal portfolio, and is the slope of the ϕ_α capital market line CML_α. On the other hand the maximum STARI of Markowitz's mean-variance optimal portfolio is equal to the slope of the Markowitz capital market line CML_2. CML_2 is always dominated by CML_α, and CML_2 is equal to CML_α only when the returns distribution is multi-variate normal in which case $\alpha_i = 2$ for all risk factors i.

We conclude that the risk adjusted return of the ϕ_α optimal portfolio ω_α is generally superior to the risk adjusted return of Markowitz's mean variance optimal portfolio ω_2. The ϕ_α paradigm results in improved investment performance.

Stable Tail Adjusted Return Ratio

While *STARI* provides a natural measure of return per unit risk, the numerical values obtained are not in a range familiar to users of the Sharpe ratio, even in the case where the returns are multivariate normal. However, it is easy to rescale *STARI* so that when the returns are normally distributed the rescaled STARI is equal to the Sharpe ratio. We use the term *Stable Tail Adjusted Return Ratio* (*STARR*) for this rescaled ratio, and its formula is

$$STARR_p(\varepsilon) = \frac{SER_p - r_f}{SETL_p(\varepsilon)/NETL^{0,1}(\varepsilon)} \qquad (4)$$

where $NETL^{0,1}(\varepsilon)$ is the *ETL* for a *standard normal distribution* at tail probability ε.

It is easy to check that $STARR_p(\varepsilon)$ coincides with the Sharpe ratio SR_p when the portfolio has a normal distribution. First one easily verifies that

$$NETL^{0,1}(\varepsilon) = \frac{1}{\varepsilon\sqrt{2\pi}} e^{-\frac{(N^{0,1}VaR(\varepsilon))^2}{2}}$$

where $N^{0,1}VaR(\varepsilon) = -\Phi^{-1}(\varepsilon)$ is the VaR of a *standard normal distribution* at tail probability ε and Φ is the standard normal cumulative distribution function, i.e., $N^{0,1}VaR(\varepsilon)$ is the ε-quantile of the *standard normal distribution*. Now suppose that the loss L_p of portfolio p has a normal distribution standard deviation σ_p, recall that the loss has zero expected value, and call the corresponding expected tail loss $NETL_p(\varepsilon)$. Then

1) $NETL_p(\varepsilon) = \sigma_p \cdot NETL_{0,1}(\varepsilon)$ (easy to verify)
2) $SETL_p(\varepsilon) = NETL_p(\varepsilon)$
3) $SER_p = ER_p$ (in any event).

Using *NTARR* to denote the resulting *STARR*, we have

$$NTARR_p = SR_p$$

which is now independent of ε.

4.6 The Choice of Tail Probability

We mentioned earlier that when using $SETL_p(\varepsilon)$ rather than $VaR_p(\varepsilon)$, risk managers and portfolio optimizers may wish to use other values of ε than the conventional VaR values of .01 or .05, for example values such as .1, .15, .2, .25 and .5 may be of interest. The choice of a particular ε amounts to a choice of particular risk measure in a the SETL family of measures, and such a choice is akin to the choice of a utility function. The tail probability parameter ε is at the asset managers disposal to choose according to his/her asset management and risk control objectives.

Note choosing a tail probability ε is not the same as choosing a risk aversion parameter. Maximizing

$$SER_p - c \cdot SETL_p(\varepsilon)$$

for various choices of risk aversion parameter c for a fixed value of ε merely corresponds to choosing different points along the ϕ_α efficient frontier. On the other hand changing ε results in different shapes and locations of the ϕ_α efficient frontier, and corresponding different excess profits π_α relative to a Markowitz portfolio.

It is intuitively clear that increasing ε will decrease the degree to which a ϕ_α optimal portfolio depends on extreme tail losses. Where $\varepsilon = .5$, which may well be of interest to some managers since it uses the average loss below zero of L_p as its penalty function, small to moderate losses are mixed in with extreme losses in determining the optimal portfolio. Our studies to date show that some of the excess profit advantage relative to Markowitz mean-variance optimal portfolios will be given up as ε increases, and that not surprisingly, this effect is most noticeable for portfolios with smaller Stable tail index values.

In summary: the smaller the tail probability ε, i.e. the more concentrated in the tail that the manager calculates risk, the higher the expected excess mean return π_α of the ϕ_α optimal portfolio over the mean-variance optimal portfolio.

It will be interesting to see what values of ε will be used by fund managers of various types and styles in the future.

4.7 The Cognity Implementation of the ϕ_α Paradigm

The ϕ_α Paradigm described in this section has been implemented in the *Cognity*™ Risk Management and Portfolio Optimization product. This product contains separate Market Risk, Credit Risk and Portfolio Optimization modules, with integrated Market and Credit Risk, and implemented in Java based architecture to support Web delivery.
For further details see *www.finanalytica.com*.

email: sales@finanalytica.com
Tel: 001 646 244 4462

REFERENCES

■ Bradley, B. O. and Taqqu, M. S. (2003). "Financial Risk and Heavy Tails", in *Handbook of Heavy Tailed Distributions in Finance*, edited by S. T. Rachev, Elsevier/North-Holland.
■ Embrechts, P., Kluppelberg, C., and Mikosch, T. (1997). *Modelling Extremal Events in Insurance and Finance*, Springer.
■ Embrechts, P., Lindskog, and McNeil, A. (2003). "Modelling Dependence with Copulas and Applications to Risk Management", in *Handbook of Heavy Tailed Distributions in Finance*, Edt. S. T. Rachev, Handbooks in Finance ,Volume 1, Elsevier Science, 329-328-384.
■ Fama, E. (1963). "Mandelbrot and the Stable Paretian Hypothesis", *Journal of Business*, 36, 420–429.
■ Mandelbrot, B. B. (1963). "The Variation in Certain Speculative Prices", *Journal of Business*, 36, 394–419.
■ Rachev, S., Schwartz, E. and Khindanova, I. (2003) "Stable Modeling of Market and Credit Value at Risk". In Handbook of Heavy Tailed Distributions in Finance, Edt. S. T. Rachev, Handbooks in Finance ,Volume 1, Elsevier Science, 249–328.
■ Rachev, Svetlozar, and Mittnik, Stefan (2000). *Stable Paretian Models in Finance*. J. Wiley.
■ Rockafellar, R. T. and Uryasev, S. (2000). "Optimization of Conditional Value-at-Risk", *Journal of Risk*, 3, 21–41.

W