Computational aspects of risk estimation in volatile markets: A survey

Stoyan V. Stoyanov
EDHEC Business School
e-mail: stoyan.stoyanov@edhec-risk.com

Svetlozar T. Rachev*
University of Karlsruhe, KIT, Germany
University of California Santa Barbara, USA, and
FinAnalytica USA
e-mail: rachev@kit.edu

Frank J. Fabozzi
Yale School of Management
e-mail: frank.fabozzi@yale.edu

CONTACT AUTHOR

Abstract

Portfolio risk estimation requires appropriate modeling of fat-tails and asymmetries in dependence in combination with a true downside risk measure. In this survey, we discuss computational aspects of a Monte-Carlo based framework for risk estimation and risk capital allocation. We review different probabilistic approaches focusing on practical aspects of statistical estimation and scenario generation. We discuss value-at-risk and conditional value-at-risk and comment on the implications of using a fat-tailed framework for the reliability of risk estimates.

*Prof Rachev gratefully acknowledges research support by grants from Division of Mathematical, Life and Physical Sciences, College of Letters and Science, University of California, Santa Barbara, the Deutschen Forschungsgemeinschaft and the Deutscher Akademischer Austausch Dienst.

1 Introduction

The quality of portfolio risk estimates depends on assumptions about the behavior of risk drivers such as stock returns, exchange rates, and interest rates. Traditional approaches are based either on the historical method or on a normal (Gaussian) distribution for risk driver returns. Neither of them, however, captures adequately the unusual behavior of risk drivers. The historical method assumes that future returns are exact replicas of past returns and the normal distribution approach is greatly limited by its inability to produce extreme returns with a realistic probability. Because of this property, the normal distribution is said to have thin tails while empirical studies indicate that asset returns are generally fat-tailed.

The non-normality of assets returns is explored in many studies and many alternative approaches have been suggested. Among the well known ones are Student's t distribution, generalized hyperbolic distributions (see Hurst et al. (1997), Bibby and Sorensen (2003), and Platen and Rendek (2007)), stable Paretian distributions (see Rachev and Mittnik (2000)), and extreme value distributions (see Embrechts et al. (2004)). At least some of their forms are subordinated normal models and thus provide a practical and tractable framework. Rachev et al. (2005) provide an introduction to heavy-tailed models in finance.

Apart from realistic assumptions for returns' distributions, a reliable risk model requires an appropriate downside risk measure. The industry standard value-at-risk (VaR) has significant deficiencies and alternatives have been suggested in the academic literature. For example, an axiomatic approach towards construction of risk measures gave rise to the family of coherent risk measures which contains superior alternatives to VaR, such as conditional value-at-risk (CVaR), also known as average value-at-risk (see, for example, Rachev et al. (2008) and the references therein).

The plan of the survey is as follows. We start with a brief description of the architecture of a Monte-Carlo based portfolio risk management system. We proceed with a discussion of computational aspects of CVaR estimation assuming fat-tailed distributions for asset returns. We compare the approach based on extreme value theory, which represents a model only for the tail of the return distribution, to approaches based on explicit fat-tailed assumptions for the entire distribution. We proceed with remarks on modeling joint dependence and a discussion of VaR and CVaR, closed-form expressions under certain parametric assumptions, and risk-budgeting. Finally, we discuss the stochastic instability of risk estimation in a Monte-Carlo based framework and provide a particular result for CVaR with fat-tailed scenarios.

2 Generic structure of a portfolio risk management system

The architecture of any portfolio risk management system has three key components as shown in Figure 1. The most basic structure is the database and the corresponding database layer. Historical data, information about portfolio positions, and different user settings are stored in it. Apart from simple storage of data, it is responsible for data retrieval, when requested from the middle layer, and also for maintenance of data.

The middle layer is the business logic layer. It is the heart of the risk management system and is responsible for carrying out mathematical calculations, such as model parameter estimation, scenarios generation, and calculation of portfolio risk statistics. It requests information from the database layer when necessary and also submits queries with changes to the database initiated by the user or by regular batch jobs. Those changes may concern historical data updates or they may be a result of, for example, portfolio rebalancing decisions.

Finally, the top layer is the user interface layer which has two goals. First, it collects the user input and sends it to the business layer for processing. Second, it receives the results from calculations performed by the business layer and presents them to the user.

A portfolio risk management system may also communicate with external systems on a regular basis. This communication may involve, for example, regular updates of the historical data.

From a modeling point of view, all operations are performed by the business logic layer. When a portfolio risk calculation is requested by the user, the following abstract steps are performed:

1. The business logic layer analyzes the requested calculation and sends historical data queries to the database layer. Those queries concern only the risk drivers relevant for the portfolio specified by the user. For example, if there is only one stock option position in the portfolio, the historical data request will concern the underlying stock, the relevant yield curve, and the relevant implied volatility surface.

2. Model parameters are fitted to the historical data. This step usually contains several sub-steps depending on the complexity of the multivariate mathematical model and we discuss it in detail in Section 3.

3. Scenarios for the relevant risk drivers are generated jointly from the fitted model. This step is crucial for the Monte Carlo method. Each vector of scenarios represents one possible state of the world. In the simple one-option portfolio example, one state of the world is represented by a vector containing a price for the underlying stock, a value

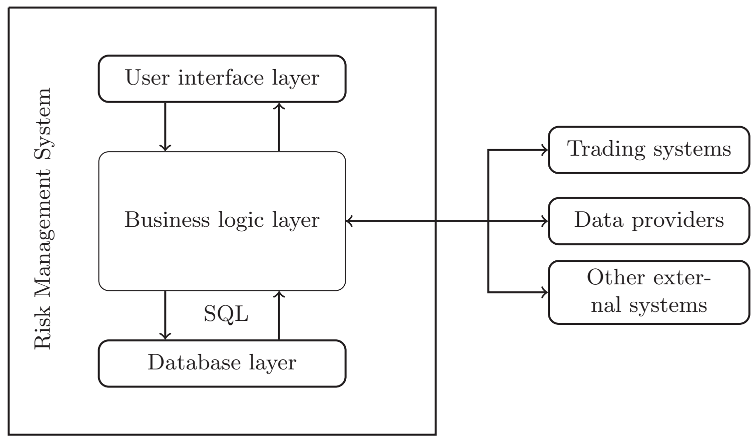

Figure 1: The basic components of a risk system.

for the interest rate, and a value for the implied volatility generated jointly from the fitted multivariate model.

4. Portfolio positions are evaluated in each state of the world. As a result, we obtain samples from the joint distribution of the positions. Scenarios at the portfolio level are calculated by summing up the corresponding position level scenarios.

5. Risk statistics are calculated from the portfolio and position level scenarios. The risk statistics are visualized in a tabular or a graphical format by the user interface layer.

The abstract steps described above are independent of the multivariate probabilistic assumption and the particular risk measure. They are generic for every Monte-Carlo based portfolio risk management system. Therefore, from a computational viewpoint, in any such system there is a trade-off between speed and accuracy. The risk drivers scenario matrix has the dimension equal to the product of the number of risk drivers and the number of scenarios. The accuracy of the final risk numbers depends on the number of scenarios. The larger the number is, the closer the generated samples will be to the theoretical distributions and, therefore, the smaller the numerical error becomes. More scenarios, however, indicate more states of the world involving more evaluations of portfolio positions which means longer calculations.

Generally, there is no universal solution working uniformly well across portfolios of different sizes and different multivariate models. One-dimensional simulation studies in Stoyanov and Rachev (2008*a*) and Stoyanov and Rachev (2008*b*) suggest that fat-tailed models require at least 10,000 scenarios which seems to be a good starting point for experiments in higher-dimensional cases.

Recent developments in the field of computer technology offer a way of pushing up the performance limit. Contemporary computer systems have multi-core processors and tendencies are for the number of cores to increase in the future. This implies that algorithms allowing for distributed calculations can benefit an enormous speed up through splitting the work among the cores.

Monte Carlo methods are inherently very well suited for distributed calculation. In the five-step algorithm above, all states of the world are generated independently which implies that subsequent calculations concerning different states of the world can be processed in parallel. In this fashion, the most computationally demanding step, which is the evaluation of portfolio positions, can be processed by the CPU cores simultaneously.

3 Fat-tailed and asymmetric models for assets returns

Reliable risk management is impossible without specifying realistic models for assets returns. Using inappropriate models may result in underestimation of portfolio risk and may lead to wrong decisions.

The distributional modeling of financial variables has several aspects. First, there should be a realistic model for the returns of individual financial variables. That is, we should employ realistic marginal models for the returns of individual assets. Second, the model should capture properly the dependence among the individual variables. Therefore, we need an appropriate multivariate model with the above two building blocks correctly specified.

3.1 One-dimensional models

The cornerstone theories in finance such as the mean-variance model for portfolio selection and asset pricing models rest upon the assumption that asset returns follow a normal distribution. Yet, there is little, if any, credible empirical evidence that supports this assumption for financial assets traded in most markets throughout the world. Moreover, the evidence is clear that the distribution of financial returns is heavy-tailed and, possibly, skewed. A number of researchers have analyzed the consequences of relaxing the normality assumption and developed generalizations of prevalent concepts in financial theory that can accommodate heavy-tailed returns (see Rachev and Mittnik (2000) and Rachev (2003) and references therein).

Mandelbrot (1963) strongly rejected normality as a distributional model for asset returns, conjecturing that financial returns behave like non-Gaussian stable returns. To distinguish between Gaussian and non-Gaussian stable distributions, the latter are commonly referred to as “stable Paretian” distributions or “Levy stable” distributions.

While there have been several studies in the 1960s that have extended Mandelbrot’s investigation of financial return processes, probably, the most notable is Fama (1963, 1965). His work and others led to a consolidation of the stable Paretian hypothesis. In the 1970s, however, closer empirical scrutiny of the “stability” of fitted stable Paretian distributions also produced evidence that was not consistent with the stable Paretian hypothesis. Specifically, it was often reported that fitted characteristic exponents (or tail-indices) did not remain constant under temporal aggregation. Partly in response to these empirical “inconsistencies,” various alternatives to the stable law were proposed in the literature, including fat-tailed distributions being in the domain of attraction of a stable Paretian law, finite mixtures of normal distributions, the Student’s t distribution, the hyperbolic distribution (see Bibby and Sorensen (2003)), and tempered stable distributions (see Bianchi et al. (2010) and Kim et al. (2010)).

Mandelbrot’s stable Paretian hypothesis involves more than simply fitting marginal asset return distributions. Stable Paretian laws describe the fundamental building blocks (e.g., innovations) that drive asset return processes. In addition to describing these building blocks, a complete model should be rich enough to encompass relevant stylized facts, such as

- non-Gaussian, heavy-tailed, and skewed distributions
- volatility clustering (ARCH-effects)
- temporal dependence of the tail behavior
- short- and long-range dependence

There exists another approach for building a fat-tailed one-dimensional model which is based on extreme value theory (EVT). EVT has been applied for a long time in areas other than finance for modeling the frequency of occurrence of extreme events. Examples include extreme temperatures, floods, winds, ocean waves, and other natural phenomena. From a general perspective, extreme value distributions represent distributional limits for properly normalized maxima of random independent quantities with equal distributions, and therefore can be applied in finance as well. The interpretation is straightforward: we can use them, for example, to describe the behavior of a large portfolio of independent losses. In contrast to the other methods, EVT provides a framework for modeling only the tails of the return distribution. Thus, the remaining part of the return distribution should be modeled by other methods.

In the remainder of this section, we describe in detail several fat-tailed models and compare them to a common EVT-based approach.

3.1.1 Stable distributions

The class of stable distributions is defined by means of their characteristic functions. With very few exceptions, no closed-form expressions are known for their densities and cumulative distribution functions (c.d.f.). A random variable X is said to have a stable distribution if its characteristic function $\varphi_X(t) = Ee^{itX}$ has the following form

$$\varphi_X(t) = \begin{cases} \exp\{-\sigma^\alpha |t|^\alpha (1 - i\beta \frac{t}{|t|} \tan(\frac{\pi\alpha}{2})) + i\mu t\}, & \alpha \neq 1 \\ \exp\{-\sigma |t| (1 + i\beta \frac{2}{\pi} \frac{t}{|t|} \ln(|t|)) + i\mu t\}, & \alpha = 1 \end{cases} \tag{1}$$

where $\frac{t}{|t|} = 0$ if $t = 0$. The formula in (1) implies that they are described by four parameters: α, called the index of stability, which determines the tail weight or density's kurtosis with $0 < \alpha \leq 2$, β, called the skewness parameter, which determines the density's skewness with $-1 \leq \beta \leq 1$, $\sigma > 0$ which is a scale parameter, and $\mu \in \mathbb{R}$ which is a location parameter. Stable distributions allow for skewed distributions when $\beta \neq 0$ and when $\beta = 0$, the distribution is symmetric around μ. Stable Paretian laws have fat tails, meaning that extreme events have high probability relative to a normal distribution when $\alpha < 2$. The tail behavior of non-Gaussian stable distributions is described by the following asymptotic relation

$$P(X > \lambda) \sim \lambda^{-\alpha}$$

which indicates that the tail decays like a power function. The Gaussian distribution is a stable distribution with $\alpha = 2$. (For more details on the properties of stable distributions, see Samorodnitsky, Taqqu (1994).) Of the four parameters, α and β are most important as they identify two fundamental properties that are atypical of the normal distribution — heavy tails and asymmetry.

Rachev et al. (2006) studied the daily return distribution of 382 U.S. stocks in the framework of two probability models — the homoskedastic independent, identically distributed model and the conditional heteroskedastic ARMA-GARCH model. In both models, the Gaussian hypothesis is strongly rejected in favor of the stable Paretian hypothesis which better explains the tails and the central part of the return distribution. The companies in the study are the constituents of the S&P 500 with complete history in the 12-year time period from January 1, 1992 to December 12, 2003. The estimated parameters suggest a significant heavy-tail and asymmetry in the residual which cannot be accounted for by a normal distribution.

Even though there is much empirical evidence in favor of the stable hypothesis, it is a theoretical fact that stable distributions with $\alpha < 2$ have an infinite second moment. Thus, if we model the return distribution of a stock with such a model, we assume it has an infinite volatility. This property creates problems in derivatives pricing models and, in order to avoid it, modifications to stable distributions have been proposed such as smoothly truncated stable laws, see Rachev et al. (2005). More general models in this direction applied to option pricing include tempered stable distributions, see Kim et al. (2008).

Generally, stable and tempered stable distributions are difficult to apply in practice because, apart from a few exceptions, there are no closed-form expressions of their density and distribution functions. There are, however, efficient numerical techniques which can be employed to construct approximations of densities and distribution functions, see Kim et al. (2008) for additional details.

3.1.2 Generalized hyperbolic distributions

A random variable X is said to have a one-dimensional generalized hyperbolic distribution if its density function is given by

$$f_X(x) = C \times \frac{K_{\lambda-1/2}\left(\sqrt{\left(\chi + \frac{(x-\mu)^2}{\sigma^2}\right)\left(\psi + \frac{\gamma^2}{\sigma^2}\right)}\right) e^{\frac{\gamma(x-\mu)}{\sigma^2}}}{\left(\sqrt{\left(\chi + \frac{(x-\mu)^2}{\sigma^2}\right)\left(\psi + \frac{\gamma^2}{\sigma^2}\right)}\right)^{1/2-\lambda}} \tag{2}$$

where

$$C = \frac{(\sqrt{\psi\chi})^{-\lambda}\psi^{\lambda}(\psi + \frac{\gamma^2}{\sigma^2})^{1/2-\lambda}}{\sqrt{2\pi}\sigma K_{\lambda}(\sqrt{\psi\chi})}$$

K_λ denotes a modified Bessel function of the third kind with index $\lambda \in \mathbb{R}$, and $x \in \mathbb{R}$. Not all the parameters have interpretations and other parametrizations are also used, see McNeil et al. (2005) for additional remarks. In this parametrization, $\mu \in \mathbb{R}$ is a location parameter, $\sigma > 0$ is a scale parameter, and $\gamma \in \mathbb{R}$ is a skewness parameter. If $\gamma = 0$, then the distribution is symmetric around μ. The parameters λ, $\psi > 0$, γ, and σ control the tail behavior which is given by the following asymptotic relation for the density

$$f_X(x) \sim |x|^{\lambda-1}e^{-\alpha|x|+\beta x} \tag{3}$$

where $\alpha = \sqrt{\frac{\psi+\gamma^2/\sigma^2}{\sigma^2}}$ and $\beta = \gamma/\sigma^2$. The parameter χ is a positive parameter but does not have an intuitive interpretation.

Note that in contrast to the power tail decay of the tails of stable distributions, generalized hyperbolic laws have a faster tail decay which is dominated by the exponential function in (3). The tail decay, however, is slower than that of the normal distribution, which makes them a good choice for modeling asset returns.

The application of generalized hyperbolic distributions in the field of finance has a long history, see, for example, Bibby and Sorensen (2003) and McNeil et al. (2005). They are infinitely divisible and can be used for derivative pricing and also as a building block in time series models.

Even though a special function appears in the definition in (2), numerical work with generalized hyperbolic distributions is facilitated by the closed-form expression of the density function. Random number generators can be constructed using the normal mean-variance mixture representation

$$X \stackrel{d}{=} \mu + W\gamma + Z\sqrt{W}$$

where $Z \in N(0, \sigma^2)$ has a normal distribution, $W \in N^{-}(\lambda, \chi, \psi)$ has a generalized inverse Gaussian distribution and is independent of Z, μ and γ are real-valued parameters, and $\stackrel{d}{=}$ denotes equality in distribution. The Box-Muller algorithm is the standard approach for generation of normally distributed random numbers and a rejection method can be employed for the generalized inverse Gaussian distribution.

With respect to parameter estimation, the classical maximum likelihood (ML) method or the expectation maximization (EM) algorithm can be employed, see McNeil et al. (2005) for additional details.

3.1.3 The EVT-based approach

EVT originated in areas other than finance. It studies the limit behavior of properly normalized maxima of independent and identically distributed (iid) random variables which in financial applications can be assumed to describe portfolio losses. There are two approaches to EVT-based modeling. The block maxima method, leading to a generalized extreme value distribution (GEV), divides the data into consecutive blocks and focuses on the series of the maxima (minima) in these blocks. The peaks-over-threshold (POT), leading to a generalized Pareto distribution (GPD), models those events in the data that exceed a high threshold. We discuss first the block-maxima method and then the POT method.

According to the theory, the limit behavior of properly normalized maxima of iid random variables is described by the GEV distribution given by

$$H_\xi(x) = \begin{cases} \exp(-(1+\xi x)^{-1/\xi}), & \xi \neq 0 \\ \exp(-e^{-x}), & \xi = 0 \end{cases} \tag{4}$$

where $1+\xi x > 0$. The parameter ξ is a shape parameter – when $\xi > 0$, H_ξ is a Frećhet distribution, when $\xi = 0$, H_ξ is a Gumbel distribution, and when $\xi < 0$, H_ξ is a Weibull distribution. The GEV distribution can be extended with a scale and a location parameter $H_{\xi,\mu,\sigma}(x) := H_\xi((x-\mu)/\sigma)$.

The block of maxima method is used to fit the parameters of the GEV distribution. In practice, it works in the following way. The historical data representing financial losses of a portfolio or an asset are divided into k blocks of size n. The maximum loss is calculated for each block. Thus, we obtain k such losses. According to ETV, this sample of k losses is asymptotically described by $H_{\xi,\mu,\sigma}$ and, therefore, we can resort to the ML method to fit the three parameters.

There are two practical problems with the block of maxima method. First, the choice of k and n is important and there is a trade-off between the two parameters because kn equals the initial sample size. While for daily financial time series, n recommended to be three months, six months, or one year, there is no general rule of thumb or any formal approach which could suggest a good choice.

Second, one needs a very large initial sample in order to have a reliable statistical estimation. In academic examples, using 20-30 years of daily returns is common, see, for example, McNeil et al. (2005). However, from a practical viewpoint, it is arguable that observations so far back in the past have any relevance to the present market conditions.

In contrast to the block-maxima method, the POT method is based on a model for exceedances over a high threshold which, in a financial context, means losses larger than a given high level. It is a model for the tail of the return distribution and not for the body. The distributional model for exceedances over thresholds is GPD given by

$$G_{\xi,\beta}(x) = \begin{cases} 1-(1+\xi x/\beta)^{-1/\xi}, & \xi \neq 0 \\ 1-\exp(-x/\beta), & \xi = 0 \end{cases} \tag{5}$$

where $\beta > 0$, and $x \geq 0$ when $\xi \geq 0$ and $0 \leq x \leq -\beta/\xi$ when $\xi < 0$. The parameters ξ and β are the shape and the scale parameter respectively. When $\xi > 0$, then $G_{\xi,\beta}$ is the distribution function of a Pareto distribution which has a power tail decay. If $\xi = 0$, then GPD turns into an exponential distribution and, finally, if $\xi < 0$, GPD is a short-tailed distribution.

As a consequence of the theoretical model, the parameters of GPD can be fitted using only information from the respective tail. There are two challenges stemming from this restriction: (1) we need to know where the body of the distribution ends and where the tail begins and (2) we need an extremely large sample in order to get a sufficient number of observations from the tail, which is an issue shared with the block-maxima method.

In practice, the high threshold is determined on an ad-hoc basis through visual identification methods, such as the mean-excess plot, or the Hill plot, see Embrechts et al. (2004). Having chosen a threshold, the parameters of GPD can fitted using the ML method, for example. The estimators, and consequentially all portfolio risk statistics based on them, are however sensitive to the choice of the high threshold.

Apart from the ML method, the Hill estimator is widely used. It is defined by

$$\hat{\xi} = \frac{1}{k}\sum_{j=1}^{k}(\log X^*_{n-j+1} - \log X^*_{n-k}) \tag{6}$$

where $X^*_1 \leq \ldots \leq X^*_n$ denote the order statistics of the sample $X_1, \ldots, X_n$, see Embrechts et al. (2004). Thus, it is based on the largest k upper order statistics. The parameter k plays the role of the high threshold in the POT method. Basically, a smaller k leads to a smaller bias but a larger variance of $\hat{\xi}$ and, as a result, there is a bias-variance trade-off that has to be taken into account when choosing k. While the Hill estimator is simple and easy to implement, there are many studies demonstrating that it performs well only if the sample is produced from a Pareto distribution. A simulation study in Rachev and Mittnik (2000) demonstrates that even 100,000 scenarios from a stable distribution prove insufficient for a proper estimation of the tail index α. There are other examples demonstrating that even mild deviations from an exact Pareto tail, such as a logarithmic perturbation of the tail

$$P(X > \lambda) \sim x^{-\alpha}/\log x,$$

may lead to a wrong estimate of the tail index, see Drees et al. (2000).

Finally, we can conclude that the estimation difficulties of EVT-based models arise because the theory is based on the asymptotic behavior of tail events. In effect, very large samples are needed which makes the approach impractical for implementation in a risk system requiring a high degree of automation. Nevertheless, EVT-based may be useful in stress-testing experiments in which the tail behavior can be manually modified to explore the potential effect on portfolio risk statistics.

3.2 Multivariate models

For the purposes of portfolio risk estimation, constructing one-dimensional models for financial instruments is incomplete. Failure to account for the dependencies among financial instruments is inadequate for the analysis.

There are two ways to build a complete multivariate model. It is possible to hypothesize a multivariate distribution directly (i.e., the dependence between stock returns as well as their marginal behavior). Assumptions of this type include, for example, the multivariate normal, the multivariate Student's t, the more general multivartiate elliptical or hyperbolic families, and the multivariate stable. Sometimes in analyzing dependence, an explicit assumption is not made, and the covariance matrix is very often relied on. Although an explicit multivariate assumption is not present, it should be kept in mind that this is consistent with the mutivariate normal hypothesis. More importantly, the covariance matrix can describe only linear dependencies and this is a basic limitation.

Since the turn of the century, a second approach has become popular. One can specify separately the one-dimensional hypotheses and the dependence structure through a function called copula. This is a more general and more appealing method because one is free to choose separately different parametric models for the individual variables and a parametric copula function. For more information, see Embrechts et al. (2002) and Embrechts et al. (2003).

From a mathematical viewpoint, a copula function C is nothing more than a probability distribution function on the d-dimensional hypercube

$$C(u_1, \ldots, u_d), \quad u_i \in [0,1], \quad i = 1, 2, \ldots, d$$

where $C(u_i) = u_i,\ i = 1, 2, \ldots, d$. It is known that for any multivariate cumulative distribution function:

$$F(x_1, \ldots, x_d) = P(X_1 \leq x_1, \ldots, X_d \leq x_d)$$

there exists a copula C such that

$$F(x_1, \ldots, x_d) = C(F_1(x_1), \ldots, F_d(x_d))$$

where the $F_i(x_i)$ are the marginal distributions of $F(x_1, \ldots, x_d)$, and conversely for any copula C the right-hand-side of the above equation defines a multivariate distribution function $F(x_1, \ldots, x_d)$. See, for example, Bradley and Taqqu (2003), Sklar (1996), and Embrechts et al. (2003).

A possible approach for choosing a flexible copula model is to adopt the copula of a parametric multivariate distribution. In this way, the copula itself will have a parametric form. There are many multivariate laws discussed in the literature that can be used for this purpose. One such example is the Gaussian copula (i.e., the copula of a multivariate normal distribution). This copula is easy to work with but it has one major drawback: It implies that extreme events are asymptotically independent. Thus, the probability of joint occurrence of large in absolute value negative returns of two stocks is significantly underestimated. An alternative to the Gaussian copula is the Student's t copula (i.e., the copula of the multivariate Student's t distribution). It models better the probability of joint extreme events but it has the disadvantage that it is symmetric. Thus, the probability of joint occurrence of very large returns is the same as the probability of joint occurrence of very small returns. This deficiency is not present in the skewed hypebolic copula which is the copula of the multivariate hyperbolic distribution defined by means of the following stochastic representation,

$$X = \mu + \gamma W + Z\sqrt{W}$$

where $W \in N^{-}(\lambda, \chi, \psi)$ has a generalized inverse Gaussian distribution, Z is multivariate normal random variable, $Z \in N_d(0, \Sigma)$, W and Z are independent, and the constants μ and γ are such that the sign of a given component of γ controls the asymmetry of the corresponding component of X, and μ is a location parameter contributing to the mean of X. The hyperbolic copula has the following features which make it a flexible and attractive model:

- It has a parametric form which makes the copula an attractive model in higher dimensions.
- The underlying stochastic representation facilitates scenario generation from the copula.
- It can describe tail dependence if it is present in the data.
- It can describe asymmetric dependence, if present in the data.

The skewed Student's t copula is a special case of the hyperbolic copula. For additional information and a case study for the German equity market, see Sun et al. (2008).

4 Risk measurement

An important activity in financial institutions is to recognize the sources of risk, manage them, and control them. A quantitative approach is feasible only if risk can be quantified. In this way, we can measure the risk contribution of portfolio constituents and then make re-allocation decisions, calculate portfolio risk break-downs by market, geography, risk driver type, or optimize portfolio risk subject to certain constraints. Functionals suitable for risk measurement cannot be arbitrary.

From a historical perspective, Markowitz (1952) introduced standard deviation as a proxy for risk. However, standard deviation is symmetric, thereby penalizing both profits and losses, and, therefore, it is more appropriate for a measure of uncertainty rather than a measure of risk. In spite of the disadvantages of this approach, pointed out in numerous studies, it is still widely used by practitioners.

A risk measure which has been widely accepted since the 1990s is value-at-risk (VaR). It was approved by regulators as a valid approach for calculation of capital reserves needed to cover market risk. Even though approved by regulators and widely used in practice, VaR has major shortcomings. In order to overcome them, axiomatic approaches were developed spawning entire families of risk measures, such as spectral risk measures, the larger family of coherent risk measures, and distortion risk measures. In the remainder of this section, we discuss VaR and conditional value-of-risk (CVaR) which is a coherent risk measure suggested in the academic literature as a superior alternative to VaR.

4.1 VaR and CVaR

Value-at-risk (VaR) at a confidence level $1-\epsilon$ is defined as the negative of the ϵ-quantile of the return distribution,

$$VaR_\epsilon(X) = -F^{-1}(\epsilon), \tag{7}$$

where F^{-1} is the inverse distribution function of X. It has been widely adopted as a risk measure. However, it is not very informative which we illustrate in the following example. Suppose that X and Y are two random variables describing the return distribution of two financial instruments. If at a given confidence level $VaR_\epsilon(X) = VaR_\epsilon(Y) = q_\epsilon$, can we state that the two financial instruments are equally risky? The answer is negative because while we know that losses larger than q_ϵ for both financial instruments will occur with the same probability ϵ, we are not sure about the magnitude of these losses.

Not only is VaR non-informative about extreme losses but it also fails to satisfy an important property directly related to diversification. The VaR of a portfolio of two positions may be larger than the sum of the VaRs of these positions,

$$VaR_\epsilon(X+Y) > VaR_\epsilon(X) + VaR_\epsilon(Y),$$

in which X and Y stand for the random payoff of the positions. As a consequence, portfolio managers may choose to make the irrational decision to concentrate the portfolio in a few positions which can be dangerous.

A risk measure which is more informative than VaR about extreme losses and can always identify diversification opportunities is CVaR. It is defined as the average VaR beyond the VaR at the corresponding confidence level,

$$CVaR_\epsilon(X) := \frac{1}{\epsilon}\int_0^\epsilon VaR_p(X)dp. \tag{8}$$

Apart from the definition in (8), CVaR can be represented through a minimization formula,

$$CVaR_\epsilon(X) = \min_{\theta\in\mathbb{R}}\left(\theta + \frac{1}{\epsilon}E(-X-\theta)_+\right) \tag{9}$$

where $(x)_+ = \max(x,0)$ and X describes the return distribution. It turns out that this formula has an important application in optimal portfolio problems based on CVaR as a risk measure. Equation (9) was first studied by Pflug (2000). A proof that equation (8) is indeed the CVaR can be found in Rockafellar and Uryasev (2002). For a geometric interpretation of (9), see Rachev et al. (2008).

4.1.1 Closed-form expressions of CVaR

Under a parametric assumption, the calculation of VaR is numerically relatively straightforward. From the definition in (7), it follows that we only need to know the inverse distribution function of the assumed distribution. Even if F^{-1} is not available in closed-form, numerical algorithms are usually readily available in the statistical toolboxes of software tools such as MATLAB or R.

Calculating CVaR is more involved due to the fact that the numerical calculation of the integral in the definition in (8) is not always simple because of the unbounded range of integration. For some continuous distributions, however, it is possible to calculate explicitly the CVaR through the definition. We provide closed-form expressions for the normal distribution and Student's t distribution.

1. *The Normal distribution*

   Suppose that X is distributed according to a normal distribution with standard deviation σ_X and mathematical expectation EX. The CVaR of X at tail probability ϵ equals

$$CVaR_\epsilon(X) = \frac{\sigma_X}{\epsilon\sqrt{2\pi}} \exp\left(-\frac{(VaR_\epsilon(Y))^2}{2}\right) - EX \qquad (10)$$

   where Y has the standard normal distribution, $Y \in N(0,1)$.

2. *The Student's t distribution*

   Suppose that X has Student's t distribution with ν degrees of freedom, $X \in t(\nu)$. The CVaR of X at tail probability ϵ equals

$$CVaR_\epsilon(X) = \begin{cases} \dfrac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right)} \dfrac{\sqrt{\nu}}{(\nu-1)\epsilon\sqrt{\pi}} \left(1 + \dfrac{(VaR_\epsilon(X))^2}{\nu}\right)^{\frac{1-\nu}{2}} & , \ \nu > 1 \\ \infty & , \ \nu = 1 \end{cases}$$

Note that equation (10) can be represented in a more compact way,

$$CVaR_\epsilon(X) = \sigma_X C_\epsilon - EX, \qquad (11)$$

where C_ϵ is a constant which depends only on the tail probability ϵ. Therefore, the CVaR of the normal distribution has the same structure as the normal VaR — the difference between the properly scaled standard deviation and the mathematical expectation. In effect, similar to the normal VaR, the normal CVaR properties are dictated by the standard deviation. Even though CVaR is focused on the extreme losses only, due to the limitations of the normal assumption, it is symmetric. Exactly the same conclusion holds for the CVaR of the Student's t distribution. The true merits of CVaR become apparent if the underlying distributional model is skewed.

It turns out that it is possible to arrive at formulae for the CVaR of stable distributions and skewed Student's t distributions. The expressions are more complicated even though they are suitable for numerical work. They involve numerical integration but this is not a severe restriction because the integrands are nicely behaved functions. The calculations for the case of stable distributions can be found in Stoyanov et al. (2006). In this section, we only provide the result.

Suppose that the random variable X has a stable distribution with tail exponent α, skewness parameter β, scale parameter σ, and location parameter μ, $X \in S_\alpha(\sigma, \beta, \mu)$. If $\alpha \leq 1$, then $CVaR_\epsilon(X) = \infty$. The reason is that stable distributions with $\alpha \leq 1$ have infinite mathematical expectation and the CVaR is unbounded.

If $\alpha > 1$ and $VaR_\epsilon(X) \neq 0$, then the CVaR can be represented as

$$CVaR_\epsilon(X) = \sigma A_{\epsilon,\alpha,\beta} - \mu$$

where the term $A_{\epsilon,\alpha,\beta}$ does not depend on the scale and the location parameters. In fact, this representation is a consequence of the positive homogeneity and the invariance property of CVaR. Concerning the term $A_{\epsilon,\alpha,\beta}$,

$$A_{\epsilon,\alpha,\beta} = \frac{\alpha}{1-\alpha} \frac{|VaR_\epsilon(X)|}{\pi\epsilon} \int_{-\overline{\theta}_0}^{\pi/2} g(\theta) \exp\left(-|VaR_\epsilon(X)|^{\frac{\alpha}{\alpha-1}} v(\theta)\right) d\theta$$

where

$$g(\theta) = \frac{\sin(\alpha(\overline{\theta}_0 + \theta) - 2\theta)}{\sin \alpha(\overline{\theta}_0 + \theta)} - \frac{\alpha \cos^2 \theta}{\sin^2 \alpha(\overline{\theta}_0 + \theta)},$$

$$v(\theta) = \left(\cos \alpha\overline{\theta}_0\right)^{\frac{1}{\alpha-1}} \left(\frac{\cos\theta}{\sin \alpha(\overline{\theta}_0 + \theta)}\right)^{\frac{\alpha}{\alpha-1}} \frac{\cos(\alpha\overline{\theta}_0 + (\alpha-1)\theta)}{\cos\theta},$$

in which $\overline{\theta}_0 = \frac{1}{\alpha} \arctan\left(\overline{\beta} \tan \frac{\pi\alpha}{2}\right)$, $\overline{\beta} = -\text{sign}(VaR_\epsilon(X))\beta$, and $VaR_\epsilon(X)$ is the VaR of the stable distribution at tail probability ϵ.

If $VaR_\epsilon(X) = 0$, then the CVaR admits a very simple expression,

$$CVaR_\epsilon(X) = \frac{2\Gamma\left(\frac{\alpha-1}{\alpha}\right)}{(\pi - 2\theta_0)} \frac{\cos\theta_0}{(\cos \alpha\theta_0)^{1/\alpha}}.$$

in which $\Gamma(x)$ is the gamma function and $\theta_0 = \frac{1}{\alpha} \arctan(\beta \tan \frac{\pi\alpha}{2})$.

A similar result for skewed Student's t distribution is given in Dokov et al. (2008).

4.1.2 Estimation of CVaR

Suppose that we have a sample of observed portfolio returns and we are not aware of their distribution. Provided that we do not impose any distributional model, both the VaR and CVaR of portfolio return can be estimated from the sample of observed portfolio returns. Denote the observed portfolio returns by $r_1, r_2, \ldots, r_n$ at time instants $t_1, t_2, \ldots, t_n$. Denote the sorted sample by $r_{(1)} \leq r_{(2)} \leq, \ldots, \leq r_{(n)}$. Thus, $r_{(1)}$ equals the smallest observed portfolio return and $r_{(n)}$ is the largest. The portfolio VaR can be estimated as the corresponding empirical quantile through an order statistic, $\widehat{VaR}_\epsilon(r) = -r_{(\lfloor n\epsilon \rfloor)}$, where $\lfloor x \rfloor$ stands for the largest integer smaller than x. The portfolio CVaR at tail probability ϵ is estimated according to the formula[1]

$$\widehat{CVaR}_\epsilon(r) = -\frac{1}{\epsilon}\left(\frac{1}{n}\sum_{k=1}^{\lceil n\epsilon \rceil - 1} r_{(k)} + \left(\epsilon - \frac{\lceil n\epsilon \rceil - 1}{n}\right) r_{(\lceil n\epsilon \rceil)}\right) \tag{12}$$

where the notation $\lceil x \rceil$ stands for the smallest integer larger than x. The "hat" above CVaR denotes that the number calculated by equation (12) is an estimate of the true value of the CVaR.

Besides formula (12), there is another method for calculation of CVaR. It is based on the minimization formula (9) in which we replace the mathematical expectation by the sample average,

$$\widehat{CVaR}_\epsilon(r) = \min_{\theta \in \mathbb{R}} \left(\theta + \frac{1}{n\epsilon}\sum_{i=1}^{n} \max(-r_i - \theta, 0)\right). \tag{13}$$

Even though it is not obvious, equations (12) and (13) are completely equivalent.

The minimization formula in equation (13) is appealing because it can be calculated through the methods of linear programming. It can be restated as a linear optimization problem by introducing auxiliary variables $d_1, \ldots, d_n$, one for each observation in the sample,

$$\begin{array}{rl} \displaystyle\min_{\theta, d} & \displaystyle\theta + \frac{1}{n\epsilon}\sum_{k=1}^{n} d_k \\ \text{subject to} & -r_k - \theta \leq d_k,\ k = 1, 2, \ldots, n \\ & d_k \geq 0,\ k = 1, 2, \ldots, n \\ & \theta \in \mathbb{R}. \end{array} \tag{14}$$

The linear problem (14) is obtained from (13) through standard methods in mathematical programming. We summarize the attractive properties of CVaR as below:

[1] This formula is a simple consequence of the definition of CVaR for discrete distributions. A detailed derivation is provided by Rockafellar and Uryasev (2002).

- CVaR gives an informed view of losses beyond VaR and is, therefore, better suited for risk management in a fat-tailed world.
- CVaR is a convex function of portfolio weights, and is therefore attractive to optimize portfolios (see Rockafellar and Uryasev (2002)).
- CVaR is sub-additive and satisfies a set of intuitively appealing properties of coherent risk measures by (see Artzner et al. (1998)).
- CVaR 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 Rachev et al. (2007)).

Even though CVaR is not widely adopted, we expect it to become an accepted risk measure as portfolio and risk managers become more familiar with its attractive properties. For portfolio optimization, we recommend the use of heavy-tailed distributions and CVaR, and limiting the use of historical, normal or stable VaR to required regulatory reporting purposes only. Finally, organizations should consider the advantages of CVaR with heavy-tailed distributions for risk assessment purposes and non-regulatory reporting purposes.

4.2 Risk budgeting with CVaR

The concept of CVaR allows for scenario-based risk decomposition which is a concept similar to the standard deviation based percentage contribution to risk (PCTR). The practical issue is to identify the contribution of each position to portfolio risk and since CVaR is a tail risk measure, percentage contribution to CVaR allows one to build a framework for tail risk budgeting. The approach largely depends on one of the coherence axioms given in Artzner et al. (1998), which is the positive homogeneity property

$$CVaR_\epsilon(aX) = aCVaR_\epsilon(X), \quad a > 0.$$

Euler's formula is valid for such functions. According to it, the risk measure can be expressed in terms of a weighted average of the partial derivatives with respect to portfolio weights,

$$CVaR_\epsilon(w'X) = \sum_i w_i \frac{\partial CVaR_\epsilon(w'X)}{\partial w_i}$$

where w is a vector of weights, X is a random vector describing the multivariate return of all financial instruments in the portfolio, and $w'X$ is the portfolio return. The left hand-side of the equation equals total portfolio risk and if we divide both sides by it, we obtain the tail risk decomposition,

$$
\begin{aligned}
1 &= \sum_i \frac{w_i}{CVaR_\epsilon(w'X)} \frac{\partial CVaR_\epsilon(w'X)}{\partial w_i} \\
&= \sum_i p_i.
\end{aligned} \tag{15}
$$

In order to compute the percentage contribution to risk of the i-th position, the i-th summand p_i in (15), we have to calculate first the partial derivative. It turns out that the derivative can be expressed as a conditional expectation,

$$
\frac{\partial CVaR_\epsilon(w'X)}{\partial w_i} = -E(X_i | w'X < -VaR_\epsilon(w'X)).
$$

when X is an absolutely continuous random variable, see Zhang and Rachev (2006) and the references therein. The conditional expectation can be computed through the Monte Carlo method.

4.3 VaR and CVaR variability

There are two sources of variability in any risk model based on the Monte Carlo method. First, there is the variability of the statistical estimators which can be intuitively described in the following way. The parameter estimates of the assumed probabilistic model depend on the input sample. Thus, a change in the historical data will result in different parameter estimates which, in turn, will change the portfolio risk numbers.

Consider, for example, the closed-form expression of CVaR given in (10). Even though this expression is based on the normal distribution, the reasoning is generic. A small change in the historical data will result in a different standard deviation σ_X and a different mean EX and, since the closed-form expression is explicitly dependent on σ_X and EX, CVaR may change as a consequence.

There are various ways to investigate the relative impact of the statistical estimators variability on portfolio risk. The most straightforward and universal one, but also very computationally demanding one, is the non-parametric bootstrap method which is a popular statistical method. The algorithm is simple and can be used if there is a closed-form expression for the risk measure or if it can be numerically approximated: (1) obtain bootstrapped samples from the initial historical data, (2) estimate the distribution parameters for each of the bootstrapped samples, and (3) calculate the risk measure for each set of distribution parameters.

Another, more simple, approach can be employed if the risk measure is a smooth function of the distribution parameters. Assume that small changes in the input sample result in small changes in the parameter estimates. In this case, we can calculate the derivatives of the risk measure with respect to the distribution parameters. The size of the derivatives determines the relative sensitivity of the risk measure for a small unit change of the corresponding parameters.

The second source of variability is inherent in the Monte Carlo method. Basically, Monte Carlo scenarios are used to calculate portfolio risk when it cannot be represented in a suitable way as a function of the distribution parameters. We hypothesize a parametric model for the multivariate distribution of financial returns, we fit the model, and then we generate a large number of scenarios. From the generated scenarios, we compute scenarios for portfolio return. Finally, employing formula (12) for example, we can calculate portfolio CVaR at a specified tail probability ϵ. In a similar fashion, we can calculate portfolio VaR using the generated portfolio return scenarios.

We can regard the generated scenarios as a sample from the fitted model and thus the computed CVaR in the end appears as an estimate of the true CVaR. The larger the sample, the closer the estimated CVaR is to the true value. If we regenerate the scenarios, the portfolio CVaR number will change and it will fluctuate around the true value. In the remaining part of this section, we discuss the asymptotic distribution of the estimator in (12) which we can use to determine approximately the variance of (12) when the number of scenarios is large. We do not consider the asymptotic theory for VaR because it is trivially obtained from the asymptotic theory of sample quantiles available, for example, in van der Vaart (1998).

Before proceeding to a more formal result, let us check what intuition may suggest. If we look at equation (12), we notice that the leading term is the average of the smallest observations in the sample. The fact that we average observations reminds us of the central limit theorem (CLT) and the fact that by averaging the smallest observations in the sample suggests that the variability should be influenced by the behavior of the left tail of the portfolio return distribution. Basically, a result based on the CLT would state that the distribution of the CVaR estimator becomes more and more normal as we increase the sample size. Applicability of the CLT however depends on certain conditions such as finite variance which guarantee certain regularity of the random numbers. If this regularity is not present, the smallest numbers in a sample may vary quite a lot as they are not naturally bounded in any respect. Therefore, for heavy-tailed distributions we can expect that the CLT may not hold and the distribution of the estimator in such cases may not be normal at all.

The formal result in Stoyanov and Rachev (2008*b*) confirms these observations. Taking advantage of the generalized CLT, we can demonstrate that

Theorem 1. *Suppose that X is random variable with distribution func-*

tion $F(x)$ which satisfies the following conditions

- $x^{\alpha}F(x) = L(x)$ *is slowly varying at infinity, i.e.* $\lim_{x\to\infty} L(tx)/L(x) = 1$, $\forall t > 0$.
- $\int_{-\infty}^{0} x dF(x) < \infty$
- $F(x)$ *is differentiable at* $x = q_\epsilon$ *where* q_ϵ *is the* ϵ*-quantile of* X.

Then, there exist c_n $n = 1, 2, ...,$ *such that for any* $0 < \epsilon < 1$,

$$c_n^{-1}\left(\widehat{CVaR}_\epsilon(X) - CVaR_\epsilon(X)\right) \xrightarrow{w} S_{\alpha^*}(1, 1, 0) \tag{16}$$

in which $\xrightarrow{w}$ *denotes weak limit,* $1 < \alpha^* = \min(\alpha, 2)$, *and* $c_n = n^{1/\alpha^*}L_0(n)/\epsilon$ *where* L_0 *is a function slowly varying at infinity and* $\widehat{CVaR}_\epsilon(X)$ *is computed from a sample of independent copies of* X *according to equation* (12).

This theorem implies that the limit distribution of the CVaR estimator in (12) is necessarily a stable distribution totally skewed to the left. In the context of the theorem, we can think of X as a random variable describing portfolio return. If the index α governing the left tail of X is $\alpha \geq 2$, then the above result reduces to the classical CLT as in this case $\alpha^* = 2$ and the limit distribution is normal. This case is considered in detail in Stoyanov and Rachev (2008*a*).

5 Summary

In this paper, we provid a review of the essential components of a model for portfolio risk estimation in volatile markets and discussed related computational aspects. We started with a description of the generic components of the architecture of a Monte-Carlo based risk management system. Any model implemented in a risk management solution adequate for volatile markets has to be capable of describing well the marginal distribution phenomena of the returns series such as fat-tails, skewness, and volatility clustering. Second, the model has to capture the dependence structure which can be done through a copula function. Finally, the risk model has to incorporate an appropriate risk measure. We considered the CVaR risk measure which has a practical meaning and appealing properties. It allows for building a risk budgeting framework based on Monte Carlo scenarios produced from a fat-tailed probabilistic model and is more suitable than the widely accepted VaR. Finally, we discussed the variability of the sample CVaR estimator in a Monte-Carlo based framework with fat-tailed scenarios.

References

Artzner, P., F. Delbaen, J.-M. Eber and D. Heath (1998), 'Coherent measures of risk', *Math. Fin.* **6**, 203–228.

Bianchi, M. L., S. T. Rachev, Y. S. Kim and F. J. Fabozzi (2010), 'Tempered infinitely divisible distributions and processes', *Theory of Probability and Its Applications (TVP), Society for Industrial and Applied Mathematics (SIAM), forthcoming* .

Bibby, B. M. and M. Sorensen (2003), 'Hyperbolic processes in finance', *Handbook of heavy-tailed distributions in finance, S. Rachev (ed)* pp. 212–248.

Bradley, B. and M. S. Taqqu (2003), 'Financial risk and heavy tails', *In Handbook of Heavy-Tailed Distributions in Finance, S. T. Rachev, ed. Elsevier, Amsterdam* pp. 35–103.

Dokov, S., S. V. Stoyanov and S. T. Rachev (2008), 'Computing VaR and AVaR of Skewed t distribution', *Journal of Applied Functional Analysis* **3**, 189–209.

Drees, H., L. de Haan and S. Resnick (2000), 'How to make a hill plot', *Annals of Statistics* **28, (1)**, 254–274.

Embrechts, P., A. McNeil and D. Straumann (2002), 'Correlation and dependence in risk management: properties and pitfalls', *In Risk management: Value at risk and beyond, edited by M. Dempster, published by Cambridge University Press, Cambridge* **10, (3)**, 341–352.

Embrechts, P., C. Klüppelberg and T. Mikosch (2004), *Modeling extremal events for insurance and finance*, Springer.

Embrechts, P., F. Lindskog and A. McNeil (2003), 'Modelling dependence with copulas and applications to risk management', *In: Handbook of heavy-tailed distributions in finance, ed. S. Rachev* pp. 329–384.

Fama, E. (1963), 'Mandelbrot and the stable Paretian hypothesis', *Journal of Business* **36**, 420–429.

Fama, E. (1965), 'The behavior of stock market pricess', *Journal of Business* **38**, 34–105.

Hurst, S. H., E. Platen and S. T. Rachev (1997), 'Subordinated market index models: A comparison', *Financial Engineering and the Japanese Markets* **4**, 97–124.

Kim, Y. S., S. T. Rachev, M. L. Bianchi and F. J. Fabozzi (2008), 'Financial market models with Lévy processes and time-varying volatility', *Journal of Banking and Finance* **32,7**, 1363–1378.

Kim, Y. S., S. T. Rachev, M. L. Bianchi and F. J. Fabozzi (2010), 'Tempered stable and tempered infinitely divisible garch models', *Journal of Banking and Finance, forthcoming* .

Mandelbrot, B. (1963), 'The variation of certain speculative prices', *Journal of Business* **26**, 394–419.

Markowitz, H. M. (1952), 'Portfolio selection', *Journal of Finance* **7, (1)**, 77–91.

McNeil, A., R. Frey and P. Embrechts (2005), *Quantitative risk management: Concepts, techniques and tools*, Princeton University Press.

Pflug, G. (2000), 'Some remarks on the value-at-risk and the conditional value-at-risk', *In: Uryasev, S. (Ed.), Probabilistic Constrained Optimization: Methodology and Applications. Kluwer Academic Publishers, Dordrecht* .

Platen, E. and R. Rendek (2007), 'Empirical evidence on Student-t log-returns of diversified world stock indices', *Research Paper Series 194, Quantitative Finance Research Centre, University of Technology, Sydney* .

Rachev, S. T. (2003), *Handbook of heavy-tailed distributions in finance*, editor, Rachev, Elsevier/North-Holland.

Rachev, S. T., C. Menn and F. J. Fabozzi (2005), *Fat tails and skewed asset returns distributions*, Wiley.

Rachev, S. T., D. Martin, B. Racheva-Iotova and S. V. Stoyanov (2007), 'Stable ETL optimal portfolios and extreme risk management', *in Risk Assessment: Decisions in Banking and Finance, G. Bol, S. T. Rachev and R. Wuerth (eds), Springer-Physika Verlag* .

Rachev, S. T. and S. Mittnik (2000), *Stable Paretian Models in Finance*, John Wiley & Sons, Series in Financial Economics.

Rachev, S. T., S. V. Stoyanov, A. Biglova and F. J. Fabozzi (2006), 'An empirical examination of daily stock return distributions for U.S. stocks', *in Data Analysis and Decision Making, Springer* .

Rachev, S. T., S. V. Stoyanov and F. J. Fabozzi (2008), *Advanced stochastic models, risk assessment, and portfolio optimization: The ideal risk, uncertainty, and performance measures*, Wiley.

Rockafellar, R. T. and S. Uryasev (2002), 'Conditional value-at-risk for general loss distributions', *Journal of Banking and Finance* **26, (7)**, 1443–1471.

Sklar, A. (1996), 'Random variables, distribution functions and copulas — a personal look backward and forward.', *In: Rüschendorf, L., Schweizer, B., Taylor, M. D. (Eds.), Distributions with fixed marginals and related topics. Institute of Mathematical Statistics, Hayward, CA* pp. 1–14.

Stoyanov, S. V., G. Samorodnitsky, S. T. Rachev and S. Ortobelli (2006), 'Computing the portfolio conditional value-at-risk in the α-stable case', *Probability and Mathematical Statistics* **26**, 1–22.

Stoyanov, S. V. and S. T. Rachev (2008*a*), 'Asymptotic distribution of the sample average value-at-risk', *Journal of Computational Analysis and Applications* **10**, 465–483.

Stoyanov, S. V. and S. T. Rachev (2008*b*), 'Asymptotic distribution of the sample average value-at-risk in the case of heavy-tailed returns', *Journal of Applied Functional Analysis* **3**, 443–461.

Sun, W., S. T. Rachev, S. V. Stoyanov and F. J. Fabozzi (2008), 'Multivariate skewed Student's t copula in the analysis of nonlinear and asymmetric dependence in the german equity market', *Studies in Nonlinear Dynamics and Econometrics* **12, (2)**, Article 3.

van der Vaart, A. W. (1998), *Asymptotic statistics*, Cambridge University Press.

Zhang, Y. and S. T. Rachev (2006), 'Risk attribution and portfolio performance measurement', *Journal of Applied Functional Analysis* **4, (1)**, 373–402.