

Der Open-Access-Publikationsserver der ZBW – Leibniz-Informationzentrum Wirtschaft
The Open Access Publication Server of the ZBW – Leibniz Information Centre for Economics

Stoyanov, Stoyan V.; Rachev, Svetlozar T.; Racheva-Iotova, Boryana; Fabozzi, Frank J.

Working Paper

Fat-tailed models for risk estimation

Working paper series in economics, No. 30

Provided in Cooperation with:

Karlsruhe Institute of Technology (KIT), Department of Economics and Business Engineering

Suggested Citation: Stoyanov, Stoyan V.; Rachev, Svetlozar T.; Racheva-Iotova, Boryana; Fabozzi, Frank J. (2011) : Fat-tailed models for risk estimation, Working paper series in economics, No. 30, <http://nbn-resolving.de/urn:nbn:de:swb:90-232444>

This Version is available at:

<http://hdl.handle.net/10419/45631>

Nutzungsbedingungen:

Die ZBW räumt Ihnen als Nutzerin/Nutzer das unentgeltliche, räumlich unbeschränkte und zeitlich auf die Dauer des Schutzrechts beschränkte einfache Recht ein, das ausgewählte Werk im Rahmen der unter

→ <http://www.econstor.eu/dspace/Nutzungsbedingungen> nachzulesenden vollständigen Nutzungsbedingungen zu vervielfältigen, mit denen die Nutzerin/der Nutzer sich durch die erste Nutzung einverstanden erklärt.

Terms of use:

The ZBW grants you, the user, the non-exclusive right to use the selected work free of charge, territorially unrestricted and within the time limit of the term of the property rights according to the terms specified at

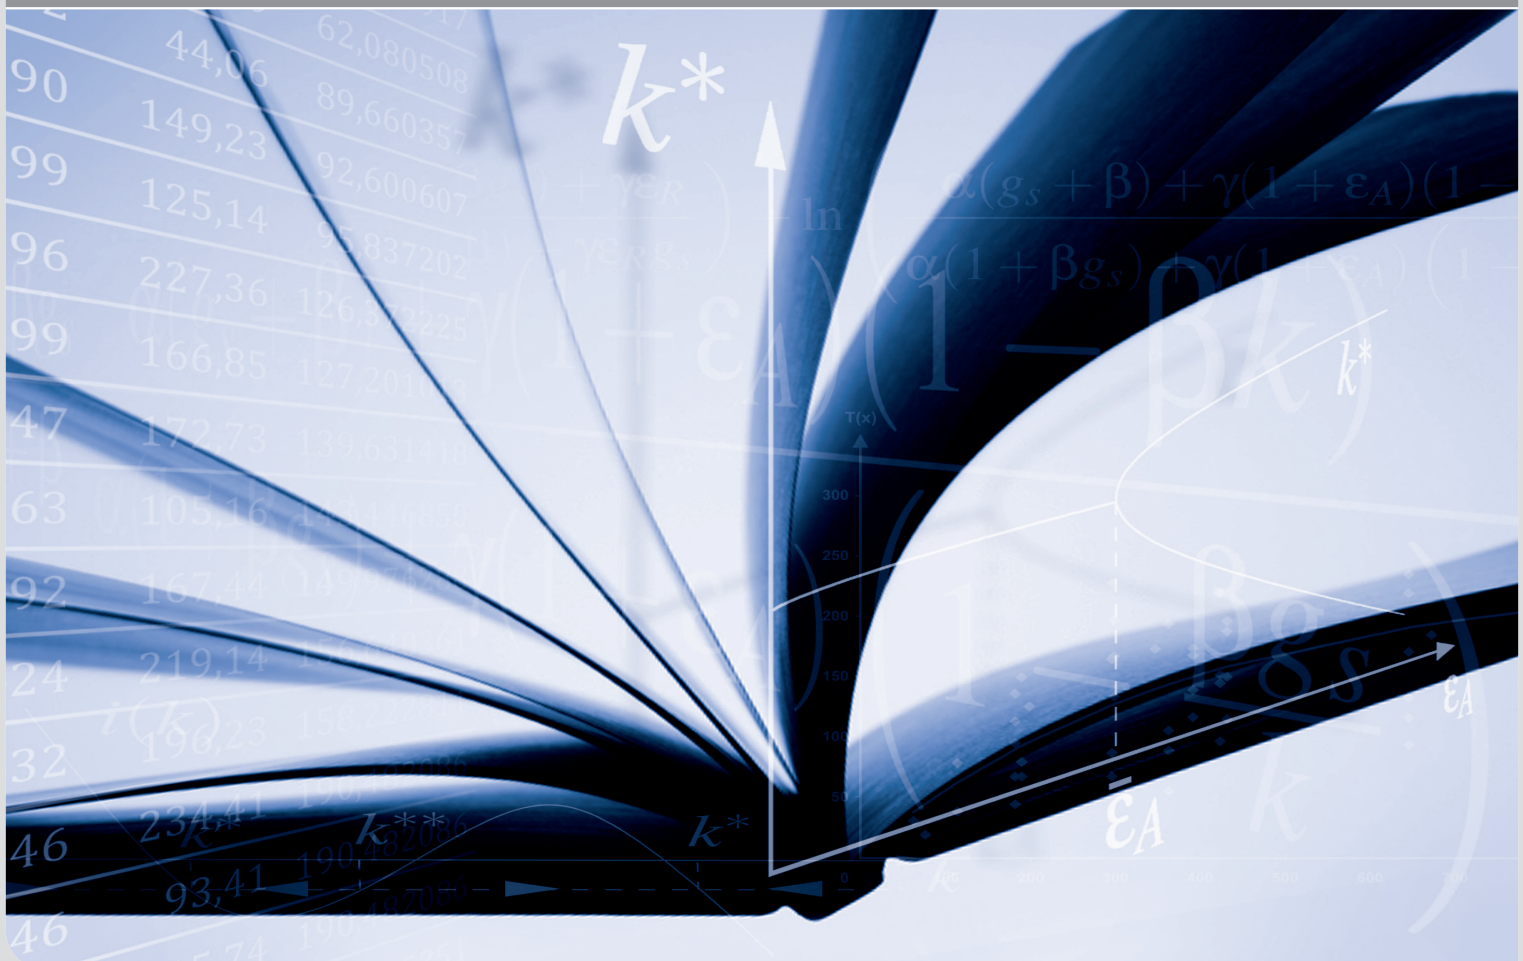
→ <http://www.econstor.eu/dspace/Nutzungsbedingungen>
By the first use of the selected work the user agrees and declares to comply with these terms of use.

Fat-tailed models for risk estimation

by Stoyan V. Stoyanov, Svetlozar T. Rachev,
Boryana Racheva-Iotova, Frank J. Fabozzi

No. 30 | MAY 2011

WORKING PAPER SERIES IN ECONOMICS



Impressum

Karlsruher Institut für Technologie (KIT)
Fakultät für Wirtschaftswissenschaften
Institut für Wirtschaftspolitik und Wirtschaftsforschung (IWW)
Institut für Wirtschaftstheorie und Statistik (ETS)

Schlossbezirk 12
76131 Karlsruhe

KIT – Universität des Landes Baden-Württemberg und
nationales Forschungszentrum in der Helmholtz-Gemeinschaft

Working Paper Series in Economics
No. 30, May 2011

ISSN 2190-9806

econpapers.wiwi.kit.edu

Fat-tailed models for risk estimation

Stoyan V. Stoyanov

EDHEC Business School
EDHEC Risk Institute–Asia

Svetlozar T. Rachev*

Stony Brook University, USA
Karlsruhe Institute of Technology, Germany, and
FinAnalytica USA

Boryana Racheva-Iotova

FinAnalytica USA

Frank J. Fabozzi

Yale School of Management

Abstract

In the post-crisis era, financial institutions seem to be more aware of the risks posed by extreme events. Even though there are attempts to adapt methodologies drawing from the vast academic literature on the topic, there is also skepticism that fat-tailed models are needed. In this paper, we address the common criticism and discuss three popular methods for extreme risk modeling based on full distribution modeling and extreme value theory.

*Dr Rachev is Frey Family Foundation Chair-Professor at Department of Applied Mathematics & Statistics, Stony Brook University. He 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.

This is the submitted version of the following article: Fat-tailed Models for Risk Estimation, Stoyanov, S. V., Rachev, S., Racheva-Iotova, B., Fabozzi, F. J., Journal of Portfolio Management, 37/2, Copyright ©2011, Institutional Investor, Inc., which has been published in final form at: <http://www.iijournals.com/doi/abs/10.3905/jpm.2011.37.2.107>

1 Introduction

The extreme losses which occurred in the financial crisis of 2008 raised the question of whether existing models and practices, largely based on the Gaussian distribution represent an adequate and reliable framework for risk measurement and management. Following the financial crisis, different ideas for methodology improvements have been proffered in white papers such as Sheikh and Qiao (2009) and Dubikovskiy et al. (2010), drawing from the existing literature on modeling of extreme events.

Empirical studies on the properties of asset return distributions date back to the pioneering work of Mandelbrot (1963) and Fama (1963, 1965) who rejected the assumption of normality and suggested the class of stable Paretian distributions as an extension to the Gaussian hypothesis. Their work led to the consolidation of the stable Paretian hypothesis which underwent empirical scrutiny in the 1970s and 1980s. It was reported that even though stock returns exhibit tails thicker than those of the Gaussian distribution, there are a few inconsistencies between observed financial return distributions and those predicted by the stable Paretian distribution. The empirical inconsistencies are related mainly to a failure of the temporal aggregation property, that is, tail thickness does not change with return frequency, and the infinite-variance property of stable laws.

Even though some of the statistical techniques used to derive these conclusions are disputed and the class of stable laws has been suggested as a viable theoretical framework,¹ alternative classes of distributions were suggested. Examples include the Student's t distribution considered by Blattberg and Gonedes (1974), the more general class of hyperbolic distributions considered by Eberlein and Keller (1995), and finite mixture of normals considered by Kon (1984). Generalizing the notion of stability to other probabilistic schemes, Mittnik and Rachev (1993) considered geo-stable, Weibull, and other distributions. More recently, tempered stable distributions have

¹See Rachev and Mittnik (2000) and the references therein.

been suggested as a model for asset returns because they retain some of the attractive theoretical properties of stable Paretian distributions but have a finite variance and can explain the temporal aggregation property.²

Clearly, there is no fundamental theory that can suggest a distributional model for financial returns and the problem remains largely a statistical one. Nevertheless, the extensive body of empirical research carried out since the 1950s leads to the following stylized facts:

- (1) clustering of volatility – large price changes tend to be followed by large price changes and small price changes tend to be followed by small price changes.
- (2) autoregressive behavior – price changes depend on price changes in the past, e.g. positive price changes tend to be followed by positive price changes.
- (3) skewness – there is an asymmetry in the upside and downside potential of price changes.
- (4) fat tails – the probability of extreme profits or losses is much larger than predicted by the normal distribution. Tail thickness varies from asset to asset.
- (5) temporal behavior of tail thickness – the probability of extreme profits or losses can change through time; it is smaller in regular markets and much larger in turbulent markets.
- (6) tail thickness varies across frequencies – high-frequency data tends to be more fat-tailed than lower-frequency data.

A methodology for risk measurement can be based on any statistically acceptable time-series model capable of capturing these stylized facts and an appropriate risk measure.

Apart from the full distribution modeling, there exists another method for modeling extreme events. It is based on the extreme value theory (EVT) developed to model extreme events in nature, such as extreme winds, temperatures, and river levels, and provides a model only for the tails for the

²See Kim et al. (2008) and Samorodnitsky and Grabchak (2010).

distribution. Embrechts et al. (1997) provide applications in the field of insurance and finance.

In the post-crisis era, while on the surface there seems to be universal agreement that financial assets are indeed fat-tailed and that investment managers must take extreme events into account as part of their everyday risk management processes, the attitude in the academic circles is not as uniform. There are constructive efforts to identify factors explaining tail events which can, potentially, be hedged (see Bhansali (2008)), or verify whether the observed dynamics in correlations during market crashes are not artifacts of the implicit assumption of normality (see Campbell et al. (2008)). Yet, in stark contrast, there are papers such as Esch (2010) suggesting fat-tailed models should be ignored in favor of Gaussian-based models on the grounds of parsimony. While from a statistical viewpoint simpler models explaining all stylized facts of the data are preferred to more complex models, representativity should not be sacrificed for simplicity. In other words, models should be as simple as possible but not any simpler.

The paper is organized in the following way. Section 2 discusses whether non-Gaussian models in general are really needed. Sections 3 and 4 focus on the two common techniques for modeling extreme losses. Finally, Section 5 discusses general techniques for model selection.

2 Are non-Gaussian models needed?

Since fat-tailed models challenge well-established classical theories, there has been a natural resistance towards them both in academia and in the practice of finance. The resistance is usually not rooted in doubts whether fat-tails exist in the real world; first, empirical research has firmly established this fact and, second, the fresh memories of the recent financial crisis are food for thought for those in denial. Rather, it is based on pragmatic reasoning which can be summarized in the following two statements: (i) the classical models are preferred because any attempt to employ a non-Gaussian model necessarily results in huge estimation errors which introduce additional uncertainty and outweigh the benefits and (ii) the observed fat tails can be explained through Gaussian-based classical models. In this section, we consider in detail each of these statements.

2.1 Estimation errors and non-Gaussian models

The discussion of estimation errors and non-Gaussian models has two aspects and we draw a clear distinction between them. The first aspect concerns application of higher-order moments in order to account for a potential skewness and excess kurtosis which are two geometric characteristics of the return distribution representing deviations from normality. The usual approach is to take advantage of the classical statistical measures of skewness and kurtosis and to use them directly in an extension of the classical mean-variance analysis for portfolio construction (see Martellini and Ziemann (2010)) or to use them for the purpose of value-at-risk (VaR) estimation through a moment-based approximation known as the Cornish-Fisher expansion. It is important to note that this approach is not related to a fat-tailed distribution and is non-parametric in nature. Loosely speaking, the goal is to describe deviations from normality in geometric terms using two descriptive statistics in addition to mean and volatility.

The criticism of this approach summarized by Esch (2010) is basically centered on the fact that the classical estimators of skewness and kurtosis exhibit significant variability to the input sample and, therefore, reliable estimation is difficult. Martellini and Ziemann (2010) extend the discussion into a multivariate setting arguing that the curse of dimensionality exacerbates further the estimation problem.

The instability of the classical estimators is a well-known problem in the field of statistics and alternative robustified estimators have been developed. As noted by Martellini and Ziemann (2010), one possible approach in an asset allocation context is to introduce a structure in the covariance, co-skewness and co-kurtosis of asset returns by means of a factor model which reduces sampling error at the cost of specification error. Also, shrinkage estimators can be considered as they provide an optimal trade-off between sampling and specification error. Finally, as far as the one-dimensional problem of skewness and kurtosis estimation of a single asset is concerned, there are robust quantile-based estimators which are much more reliable than the classical estimators.

Dismissing the non-parametric approach on the basis of the properties of classical estimators for skewness and kurtosis is a hurried decision. Martellini and Ziemann (2010) conclude that there can be a significant value added in asset allocation decisions if the estimation problem is handled properly.

The extension of the mean-variance analysis with skewness and kurto-

sis is, essentially, based on a higher-order Taylor series approximation of expected utility and is well-motivated by theory. However, employing the non-parametric method for downside risk estimation leads to other difficulties. From a conceptual viewpoint, while skewness and kurtosis do describe deviations from the Gaussian distribution in geometric terms, they do not focus on the downside of the return distribution. The described deviations from normality more or less concern the central part of the distribution and there are no reasons to believe that a skewed distribution with an excess kurtosis will better describe extreme losses. For any distribution, skewness and kurtosis represent two numbers and it is impossible to describe the richness of the possible shapes of the tail behavior only in terms these two numbers.

As a result, for the purposes of risk estimation it is better to rely on other techniques especially when quantiles deep in the tail are involved. These techniques involve assuming a fat-tailed model that can describe the tail of the empirical distribution and bring up the second aspect of the discussion on estimation errors and non-Gaussian models. Fat-tailed models have parameters that are related to their tail behavior and skewness in addition to the mean and scale³ parameters meaning the number of parameters of a fat-tailed distribution exceeds the number of parameters of the normal distribution which results in additional complexity. The common argument summarized by Esch (2010) is that the notion of parsimony dictates simpler models are better because complexity may turn out to be redundant which is not easily detectable in-sample.

This is a very general philosophical argument which is counter-productive if used without further analysis. First, additional complexity may be needed in order to explain observed facts and this holds not only for finance but for the realm of science in general. If skewness and fat-tails are observed in the real world, it is only natural to choose a family of distributions that can take into account the empirical facts at the cost of introducing two additional parameters. Model selection techniques can be employed to assess whether the improvement of the extended model is statistically significant.

Second, it is incorrect to think that fat-tailed models are alternatives to the Gaussian model. In fact, most classes of fat-tailed distributions contain the Gaussian distribution as a special case. In this sense, they truly extend the framework and if the data is Gaussian, then the model would recognize this fact. There are statistical techniques for hypothesis testing that can

³Volatility describes the “scale” of the normal distribution.

verify if the fitted parameters of the fat-tailed model are significantly different from the parameter values identifying the normal distribution. The labels “non-Gaussian” and “fat-tailed” as opposed to “Gaussian” and “non-fat-tailed” are supposed to reflect the underlying properties and not to imply that these models cannot co-exist in a more general framework.

From this perspective, the Gaussian model itself represents a significant restriction that can lead to overly optimistic risk numbers in times of crashes and realistic numbers in ordinary market conditions. Thus, as far as the modeling philosophy is concerned, it is better to start with a more general model and restrict it only if the impact of the “noise” coming from the additional complexity outweighs the benefits.

2.2 Explanations of fat-tails through Gaussian-based models

The main argument in this section is related to the last argument of the previous section – in the light of a non-explained empirical fact, it is preferable to try to describe it first with the available tools before increasing the model complexity.

There are two main approaches to explain fat-tails through Gaussian based models. First, fat-tails appear because the volatility of asset returns is dynamic. Therefore, looking at the unconditional distribution as described by histogram plots or kernel density plots we can see heavy tails but they are mostly due to the time varying volatility. Second, fat-tails appear because asset returns depend in a non-linear fashion on other factors which are distributed according to a Gaussian law.

Empirical research has firmly established the fact that volatility clustering is a stylized fact of asset returns.⁴ Different econometric models have been suggested to explain the time varying volatility and the most widely used ones are the GARCH-type models. These econometric models can be regarded as filters that transform the empirical data by explaining certain phenomena and produce new data, also known as residual, which is more homogeneous. As a consequence, it is possible to run a statistical test on the residual to verify if explaining the volatility clustering effect leads to a Gaussian distribution. Empirical research has shown even though the residual appears to be less

⁴For an empirical research, see for example Akgiray (1989).




Exhibit 1: The degrees of freedom parameter of the residuals of a time-series model fitted on the returns of the constituents of S&P500.

fat-tailed than the original data, it is not Gaussian.⁵

We illustrate this behavior in an empirical calculation which includes the stocks in the S&P500 universe during the 3-year period from January 1, 2007 to November 1, 2010. We fitted a time series model⁶ to clean the clustering of volatility effect assuming the classical Student's t model on the residual. The degrees of freedom (DOF) parameter, which governs the tail behavior, is shown on Figure 1. The smaller the value is, the heavier the tails are. If DOF has a value of about 30 and above, it can be argued that the Student's t distribution is close to the normal distribution.

We calculated that 27% of all stocks have a DOF above 7 and 13% have a DOF below 4 indicating a diverse tail behavior. These results agree with a similar calculation for the same universe done for a different time period, see Rachev et al. (2010). Figure 1 demonstrates that the time variations in the volatility cannot explain the fat-tails observed in the unconditional

⁵See Rachev and Mitnik (2000) and the references therein for extensive econometric studies.

⁶The time series model is an ARMA(1,1)-GARCH(1,1) model.

distribution of stock returns. i.e. we find significant tail thickness in the residual even though the time-series model captures the clustering of the volatility effect.

Even though asset returns do not seem to be conditionally Gaussian, GARCH-type models are a very useful tool because they can describe the time structure of volatility. Volatility, or more generally the scale parameter of the distribution, is an important factor in risk estimation because it affects all quantiles of the return distribution including the very extreme ones. Therefore, any probabilistic model for risk estimation should include a GARCH-type component.

The argument that, hypothetically, the fat-tailed behavior of asset returns can be explained through non-linear factor models and, therefore, fat-tailed distributions are redundant is weak without any empirical justification that such factor models can be reliably fitted and that the factor returns are indeed normal. This argument should be considered on a case-by-case basis and not generally.⁷ Certainly, a good non-linear factor model (or linear one for that matter) is a powerful tool to identify risk exposures and is instrumental for portfolio risk management purposes. The structure provided by such a factor model is a useful piece of information in addition to describing the shape of the portfolio return distribution. However, if it turns out that the factor model is not reliable because of low explanatory power or if our goal is only risk measurement, then having a good description of the (conditional) distribution of portfolio returns is sufficient.

Non-linear factor models represent a big family of diverse models. There are three basic types models depending on how non-linearity is introduced – (i) non-linear in the factors but linear in the parameters, (ii) non-linear in both the factors and the parameters, and (iii) no particular functional form of the non-linearities, e.g. a kernel regression. From a parsimony standpoint, it is very easy to end up with a model that overfits the data which does not imply that non-linear factor models should be disregarded. This is a class of techniques that can be helpful and should be used very carefully bearing in mind the potential issues.

Rachev et al. (2010) illustrate the time variations of tail thickness in the daily returns of the DJIA index from October 1997 to October 2009. The

⁷The term *non-linear factor model* is quite general and without a specification of the non-linear function defining the relationship is of no use. In fact, any random variable can be represented as a non-linear factor model of factors having arbitrary pre-specified distributions.

example indicates that tail thickness may increase in times of crashes along with volatility implying that factors unexplained by the time-series model may impact downside risk. Unless these factors can be explained completely by means of linear or non-linear factor models, it is important to adopt a fat-tailed model which can provide a reasonable statistical description of the data.

3 Full distribution modeling

Apart from the non-parametric approach based on higher-order moments, there are two approaches for modeling the tails of the return distribution. One of them is to assume fat-tailed distribution and the other one is based on EVT. The two approaches are different in nature. A fat-tailed distributional assumption represents a model for the entire distribution while EVT-based models are designed to describe the maximum loss over a given horizon meaning that they are designed to describe the extreme tail only.

As mentioned before, a GARCH-type model for the scale of the return distribution is of crucial importance. Therefore, the two approaches mentioned in this section have to be applied to the residual produced by applying the GARCH filter to the observed asset returns. In this section, we consider the full distribution modeling approach represented by two Student's t distribution and tempered stable distributions.

3.1 The Student's t distribution

First introduced in 1908, the Student's t distribution is probably the most commonly used fat-tailed distribution as a model for asset returns. Like the normal distribution, classical Student's t densities are symmetric and have a single peak. Unlike the normal distribution, Student's t densities are more peaked around the centre and have fatter tails.

This property is illustrated on Exhibit 2. In fact, the normal distribution is a special case when the DOF parameter approaches infinity. For practical purposes, however, the plot indicates a very small difference for a DOF of 30. While this property makes the Student's t distribution acceptable for asset returns modeling, the real reason behind its widespread use is its ease of application – numerical methods are easily implementable and are widely available.

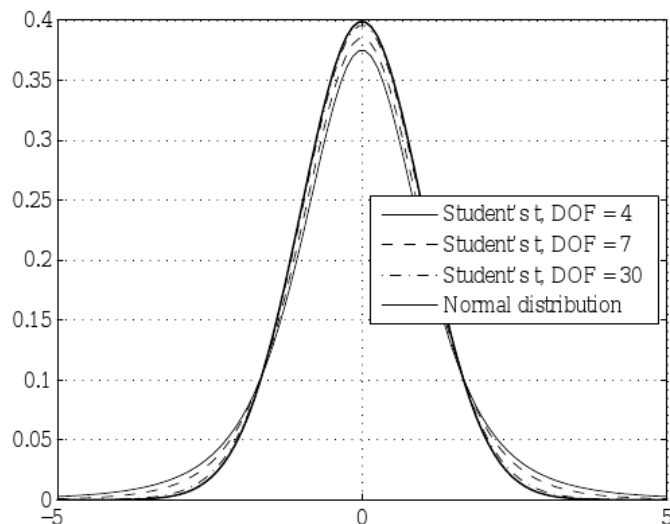


Exhibit 2: The density of Student's t distribution for different choices of the DOF parameter.

We noted that it is unreasonable and potentially dangerous to accept the Gaussian distribution simply on the grounds of parsimony without any statistical analysis. In a similar fashion, fixing the tail thickness to a given value for all assets makes little sense. A risk estimation model based on Student's t distribution with the DOF parameter fixed to 5 is suggested by Zumbach (2006). Even though the values of the DOF parameter calculated for the S&P 500 universe on Exhibit 1 imply that an average value of 5 is perhaps reasonable, fixing the DOF parameter leads to a significant limitation⁸ – it does not allow for a smooth transition between Gaussian data and fat-tailed data. In contrast to the Gaussian model which is overly optimistic in times of crashes, this model will significantly overestimate the risk for assets with returns that are close to being normally distributed since it fails to account for another stylized fact – tail thickness varies between assets and across time.

The volatility in the model suggested by Zumbach (2006) is estimated by applying weights to the observed returns with a logarithmic or exponential

⁸The motivation is given in Section 4 of Zumbach (2006).

decay based on a predefined parameter.⁹ This forces the relative importance of the observations in the past to be the same for all risk drivers and across all time periods. While this universal parameter makes these models simpler and easier to grasp, there is an important trade-off between simplicity and precision: these models are less accurate and only work “on average” in a universe of risk drivers. Even though the approach in Zumbach (2006) deviates from the traditional GARCH-type framework, an implementation with the classical Student’s t distribution for the residual without the deficiency of fixing the DOF parameter is available, for example, in the GARCH toolbox of MATLAB.

Finally, the classical Students t model is symmetric. In cases where there is a significant asymmetry in the data, it will not be reflected in the risk estimate.

3.2 Tempered stable distributions

As noted in the introduction, the class of stable distributions has a special place among non-Gaussian full distribution models. It was suggested as a model for asset returns because it contains the normal distribution as a special case sharing a remarkable property – only this class of distributions can approximately describe the behavior of a stochastic system influenced by many small, regular, and independent random factors. Since price changes are driven by many random factors, it is reasonable to assume that stable distributions could represent a model for their approximate behavior (see Rachev and Mitnik (2000)). This was the reason the stable Paretian hypothesis was suggested in the 1960s.

Like the Students t distribution, stable Paretian distributions have a parameter responsible for the tail behavior, which is called the tail index or index of stability. In contrast to the DOF parameter, the index of stability is between zero and two. The closer it is to two, the more Gaussian like the distribution is; smaller values of the index of stability imply a fatter tail. Unlike the Students t distribution, stable distributions allowed for skewed representatives, see Exhibit 3.

The two empirical inconsistencies with the stable Paretian hypothesis is that it implies infinite variance for asset returns and a tail behavior which

⁹An example is a decay parameter of the equally weighted moving average method equal to 0.94.




Exhibit 3: The density of the stable distribution for different choices of tail index α and with the skewness parameter fixed to -0.7 .

does not change with the return frequency. Therefore, the good results reported in empirical studies seem paradoxical (see Samorodnitsky and Grabchak (2010)). Efforts to obtain classes of distributions which allow for aggregation across frequencies, have a finite volatility and do not deviate much from the shape of stable distributions lead to the development of tempered stable distributions. In a nutshell, tempered stable distributions are derived from the class of stable distributions through a process called tail tempering. Tail tempering is achieved by modifying only the tails of stable distributions so that they remain thicker than the Gaussian tails but do not lead to an infinite volatility. The technique is described in Kim et al. (2008), Kim et al. (2010), and Bianchi et al. (2010).

Even though modifying the natural tail behavior of stable distributions may seem artificial, the idea behind the tail tempering technique is based on solid scientific grounds and is related to how temporal aggregation works. Generally, lower frequency logarithmic returns can be represented as sums of relatively more irregular higher frequency logarithmic returns. If we think of the Gaussian behavior of the lower frequency return as some sort of a limit

behavior, then an explanation for the change in the tail behavior can be the convergence rate to that limit. For example, monthly returns and weekly returns can be represented as sums of daily returns, the only difference is in the number of summands. Intuitively, the convergence rate would be faster if there are more summands (higher vs lower frequencies) which are also more regular (normal vs extreme market conditions).

The tail tempering technique arises from results in probability theory dealing with the problem of estimating the rates of convergence in limit theorems indicating that the shape of the distribution of the sum looks like a stable distribution at the center but does not have as heavy tails, see Samorodnitsky and Grabchak (2010).

4 Extreme value theory

EVT has been applied for a long time when modeling the frequency of extreme events, including extreme temperatures, floods, winds and other natural phenomena. From a general perspective, extreme value distributions represent distributional limits for properly normalized maxima of independent random quantities with equal distributions, and therefore can be applied in finance as well, see Embrechts et al. (1997). In contrast to the other distribution families mentioned in the introduction, EVT represents a model for the tail of the distribution only. Therefore, in practice, one needs to combine EVT with a model for the remaining part of the distribution if needed.

The method behind EVT can be intuitively described in the following way. Consider a sequence of returns at a given frequency. The maximum loss can be approximately described through a limit distribution known as the generalized extreme value distribution (GED). Another way to model extreme losses is to consider the exceedances over a high threshold. EVT indicates that asymptotically, as the high threshold increases, the exceedances can be described by the generalized Parreto distribution (GPD).

There are many empirical studies applying EVT directly to the return time series ignoring the clustering of volatility effect (see, for example, Marinelli et al. (2007) and Sheikh and Qiao (2009)). While at times this may be a reasonable simplification for shorter time intervals, in practice the stylized facts mentioned in the introduction can represent significant deviations from hypothesis that returns are independent and identically distributed (iid). As a consequence, the stylized facts have to be modeled separately and EVT




Exhibit 4: The daily losses of S&P500 from November 3, 2000 to November 5, 2010 and the maxima of losses in blocks of 6 months.

should be applied to the residual of the time series model for which the iid assumption is more adequate.

The two limit distributions, GEV and GPD, give rise to two approaches of EVT-based modeling — the block of maxima method and the peaks-over-threshold method.

4.1 The block of maxima method

The idea behind the block of maxima approach is essentially to take advantage of the limit behavior described by GEV. We divide the data into consecutive blocks of equal size and then focus on the series of the maxima of the returns in these blocks. As a next step, we fit the GED to the series of maxima, which can be performed via the method of maximum likelihood.

The block of maxima method is illustrated in Exhibit 4 with the daily losses of S&P500 in the period between November 3, 2000 and November 5, 2010 with the size of the block equal to 6 months. This example also illustrates that applying naively EVT directly to the time series of losses may lead to difficulties. The maxima of losses between 2003 and 2007 occur in a low volatility environment while the ones in 2008 occur in a very high

volatility period and, therefore, they are not generated by one and the same distribution.

There are two practical problems with the block of maxima method: (i) the choice of the size of the blocks and (ii) the length of the time series. While for daily financial time series, the size of the block is 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. Even these recommendations are rather arbitrary because the size of the resulting sample of maxima can vary up to four times which naturally can exercise a big impact on the estimates.

Concerning the second problem, 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.

4.2 The peaks-over-threshold method

The peaks-over-threshold (POT) method arises from the limit result leading to GPD. Like the block of maxima method, the parameters of GPD can be fitted using only information from the respective tail. The process is straightforward – choose a value for the high threshold and fit GPD to the part of the sample which exceeds the threshold.

The approach is illustrated in Exhibit 5 which also shows why applying POT directly on the returns series is unreasonable when there is a significant clustering of volatility – the larger exceedances are from a high volatility environment and cannot be compared to the smaller exceedances in normal markets. Further on, they appear in clusters which is in conflict with the iid assumption.

There are two big challenges stemming from this restriction: (i) a very large sample is needed to get a sufficient number of observations from the tail, which is also present in the block of maxima method, and (ii) we need to know where the body of the distribution ends and where the tail begins in order to set a good value for the threshold.

Concerning the sample size, academic publications indicate that minimum requirements are in the range of 5,000 to 10,000 observations. However, it should also be noted that using a sample this large minimizes any current fat-tailed market behavior and would only be suitable for long-term projections.




Exhibit 5: The daily losses of S&P500 from November 3, 2000 to November 5, 2010 and the exceedances over a threshold of 3%.

For financial time series, where the standard time window for risk estimation is two years of daily data, a sample size of just 500 observations is far too short to ensure an accurate GPD fit. Goldberg et al. (2008) suggest 1,000 days, with approaches to generate synthetic data where enough observations are not available.

The second issue may seem easy to resolve by resorting to statistical methods that would indicate where the tail begins. However, no reliable methods of this type exist. Typically, this high threshold is chosen subjectively by looking at certain plots, such as the Hill plot or the mean excess plot, which are standard in EVT. As a consequence, identifying the threshold between the body and the tail is a matter of subjective choice based on visual inspection, which cannot be achieved on a large scale. Kuipers test has been suggested in Goldberg et al. (2008) as a numerical method for determining the optimal threshold selection. However, this test is very difficult to automate for large universes, because the resulting optimization problem does not have good optimality properties, with the global minimum being hard to find. The fact that the choice of this threshold has a great impact on the parameter estimates of GPD and, therefore, on the final risk estimates in smaller samples is acknowledged in Goldberg et al. (2008) and Castillo and

Daoudi (2008).

In using GPD, one is always faced with the classical tail-estimation trade-off problem. For the GPD estimates to be unbiased, they must be fit with the largest possible threshold. Unbiased estimators are obtained when the threshold is infinity. This implies that one should use a very small number of extreme observations from the original sample. On the other hand, using such a small number of observations drastically increases the variance of the estimators. GPD estimation fits become more of an art than a science in balancing this trade-off between being unbiased and small variance of the estimators. It should be noted that fitting the parameters of GEV has the same deficiencies, see Hull and Welsh (1985).

5 Model selection

Having discussed different classes of fat-tailed models which can account for the observed fat-tailed behavior of asset returns, the pragmatic mind would be interested in criteria for model selection. A few general rules can be considered which can eliminate inappropriate choices. First, models which allow for fat-tailed representatives and also for light-tailed representatives are more appealing compared to models with a fixed tail thickness, e.g. the one suggested in Zumbach (2006). In this way, if the empirical tail is close to normal the model would be able to recognize it. Rejecting a richer model only on the grounds of parsimony as argued in Esch (2010) makes sense only after an extensive statistical analysis.

Second, choosing a model may depend on the particular problem we are trying to solve. From a theoretical viewpoint, some models arise from limit theories in probability theory which make them suitable for some types of problems. Operational risk, for example, requires loss estimation deep in the tail, e.g. VaR at the 99.9% level, which makes EVT-based models good candidates. They have to be applied carefully bearing in mind the caveats mentioned in Section 4. Full distribution models are more appropriate in market risk estimation when the left tail of the portfolio return distribution is used in the calculation of multiple quantile levels some of which are not in the extreme tail, e.g. VaR at 95%, 99%, and 99.5%. Tempered stable distributions represent an appealing choice for the residual process in a time-series model but other distributions can also be used.

Third, back-testing studies should be considered when choosing between

different models. They can identify certain models as overly optimistic or overly pessimistic. Rachev et al. (2010) consider a back-testing of VaR at 99% confidence level of DJIA in which Gaussian, Student's t , stable and EVT models are compared after cleaning the clustering of volatility effect through a GARCH model. The back-testing study indicates that EVT and Student's t with the DOF parameter fixed to 5 as suggested in Zumbach (2006) lead to overly conservative VaR models. Not surprisingly, the Gaussian model, on the other hand, turns out to be overly optimistic. Additional analysis and a comparison in a more simplified setting is available in Marinelli et al. (2007).

Having selected a distributional model, one has to calculate confidence intervals of the risk statistics through non-parametric bootstrapping which provides a measure of sensitivity of the output of the model to the input sample. The same exercise can be repeated with the parameters of the fat-tailed distribution which will provide insight into the variability of the estimator. If the Monte Carlo method is involved, then the Monte Carlo variability of the risk statistics has to be explored as well in order to tune the number of generated scenarios. Studies of this type help identify factors that may cause instabilities and situations in which the model may fail.

6 Conclusion

After the financial meltdown of 2008, it seems the existence of fat tails in asset returns and the importance of this phenomenon for downside portfolio risk estimation is getting more widely recognized by practitioners. Academic literature is abundant in models that can take into account fat tails but this stylized fact should not be considered in isolation. A realistic model for asset returns should take into account all statistically relevant stylized facts. The best approach is to choose an extended model that includes the normal distribution as a special case and then test its performance on the real data. The more theoretical question of whether the chosen model is the best one among all other families of fat-tailed models is not so essential as there may be different families of fat-tailed distributions which are statistically equivalent as far as the particular modeling problem is concerned.

References

- Akgiray, A. (1989), ‘Conditional heteroskedasticity in time series of stock returns: evidence and forecast’, *Journal of Business* **62**, 55–80.
- Bhansali, V. (2008), ‘Tail risk management’, *Journal of Portfolio Management* **34(4)**, 68–75.
- 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* **55(1)**, 59–86.
- Blattberg, R. C. and N. J. Gonedes (1974), ‘A comparison of the stable and student distributions as statistical models for stock prices’, *Journal of Business* **47**, 244–280.
- Campbell, R., C. Forbes, K. Koedijk and P. Kofman (2008), ‘Increasing correlations or just fat tails?’, *Journal of Empirical Finance* **15(2)**, 287–309.
- Castillo, J. Del and J. Daoudi (2008), ‘Estimation of the generalized pareto distribution’, *Statistics & Probability Letters* **79(5)**, 684688.
- Dubikovsky, V., M. Hayes, L. Goldberg and M. Liu (2010), ‘How well can the risk of financial extremes be forecast?’, *Whitepaper, MSCI Barra, Research Insights* .
- Eberlein, E. and U. Keller (1995), ‘Hyperbolic distributions in finance’, *Bernoulli* **1(3)**, 281–299.
- Embrechts, P., C. Klüppelberg and T. Mikosch (1997), *Modeling extremal events for insurance and finance*, Springer.
- Esch, D. (2010), ‘Non-normality facts and fallacies’, *Journal of Investment Management* **8(1)**, 49–61.
- Fama, E. (1963), ‘Mandelbrot and the stable paretian hypothesis’, *Journal of Business* **36**, 420–429.
- Fama, E. (1965), ‘The behavior of stock market prices’, *Journal of Business* **38**, 34–105.

- Goldberg, L., G. Miller and J. Weinstein (2008), ‘Beyond value at risk: Forecasting portfolio loss at multiple horizons’, *Journal of Investment Management* **6(2)**, 73–98.
- Hull, P. and A. Welsh (1985), ‘Adaptive estimates of parameters of regular variation’, *Annals of Statistics* **12(1)**, 331–341.
- 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. T. Rachev, M. L. Bianchi and F. J. Fabozzi (2010), ‘Tempered stable and tempered infinitely divisible GARCH models’, *Journal of Banking and Finance* **34**, 2096–2109.
- Kon, S. (1984), ‘Models of stock returns - a comparison’, *Journal of Finance* **39**, 147–165.
- Mandelbrot, B. (1963), ‘The variation of certain speculative prices’, *Journal of Business* **26**, 394–419.
- Marinelli, C., S. D’Addona and S. Rachev (2007), ‘A comparison of some univariate models for value-at-risk and expected shortfall’, *International Journal of Theoretical and Applied Finance* **10(6)**, 1043–1075.
- Martellini, L. and V. Ziemann (2010), ‘Improved estimates of higher-order comoments and implications for portfolio selection’, *Review of Financial Studies* **23(4)**, 1467–1502.
- McNeil, A., R. Frey and P. Embrechts (2005), *Quantitative risk management: Concepts, techniques and tools*, Princeton University Press.
- Mittnik, S. and S. Rachev (1993), ‘Modelling asset returns with alternative stable distributions’, *Econometric Reviews* **12(3)**, 261–230.
- Rachev, S., B. Racheva-Iotova and S. Stoyanov (2010), ‘Capturing fat tails’, *Risk Magazine* **May 2010**, 72–77.
- Rachev, S.T. and S. Mittnik (2000), *Stable Paretian Models in Finance*, John Wiley & Sons, Series in Financial Economics.

- Samorodnitsky, G. and M. Grabchak (2010), ‘Do financial returns have finite or infinite variance? a paradox and an explanation’, *Quantitative Finance* .
- Sheikh, A. Z. and H. Qiao (2009), ‘Non-normality of market returns: A framework for asset allocation decision making’, *Whitepaper, J.P. Morgan Asset Management, JPMorgan Chase & Co* .
- Zumbach, G. (2006), ‘A gentle introduction to the RM2006 methodology’, *RiskMetrics Technology Paper* .

Working Paper Series in Economics

recent issues

- No. 30** *Stoyan V. Stoyanov, Svetlozar T. Rachev, Boryana Racheva-Iotova, Frank J. Fabozzi: Fat-tailed models for risk estimation, May 2011*
- No. 29** *Stoyan V. Stoyanov, Svetlozar T. Rachev, Frank J. Fabozzi: CVaR sensitivity with respect to tail thickness, May 2011*
- No. 28** *Young Shin Kim, Svetlozar T. Rachev, Michele Leonardo Bianchi, Frank J. Fabozzi: Tempered stable and tempered infinitely divisible GARCH models, May 2011*
- No. 27** *Takashi Kanamura, Svetlozar T. Rachev, Frank J. Fabozzi: A profit model for spread trading with an application to energy futures, May 2011*
- No. 26** *Michele Leonardo Bianchi, Svetlozar T. Rachev, Young Shin Kim, Frank J. Fabozzi: Tempered infinitely divisible distributions and processes, May 2011*
- No. 25** *Sebastian Kube, Michel André Maréchal and Clemens Puppe: The currency of reciprocity - gift-exchange in the workplace, April 2011*
- No. 24** *Clemens Puppe and Attila Tasnádi: Axiomatic districting, April 2011*
- No. 23** *Dinko Dimitrov and Clemens Puppe: Non-bossy social classification, April 2011*
- No. 22** *Kim Kaivanto and Eike B. Kroll: Negative recency, randomization device choice, and reduction of compound lotteries, April 2011*
- No. 21** *Antje Schimke and Thomas Brenner: Long-run factors of firm growth - a study of German firms, April 2011*
- No. 20** *Aaron B. Scholz: Spatial network configurations of cargo airlines, April 2011*
- No. 19** *Arne Beck: Public bus transport in Germany - a proposal to improve the current awarding system, April 2011*