



universität
wien

Magisterarbeit

Titel der Magisterarbeit

“Long-Run Value at Risk: Approaches, Models,
Parameters, and Assumptions”

Verfasser

Erich Arthur Stark, Bakk. rer. soc. oec.

angestrebter akademischer Grad

Magister der Sozial- und Wirtschaftswissenschaften
(Mag. rer. soc. oec.)

Wien, im April 2008

Studienkennzahl lt. Studienblatt:

A 066 915

Studienrichtung lt. Studienblatt:

Magisterstudium Betriebswirtschaft

Betreuer:

Ao. Univ.-Prof Mag. Dr. Christian Keber

Table of Contents

Table of Contents	1
Acronyms	3
Notation	4
Acknowledgments	5
1. Introduction	6
2. Laying the Groundwork	8
2.1. Overview.....	8
2.2. The Value at Risk Methodology.....	8
2.3. Standard Short-Run Value at Risk Models.....	15
2.3.1. Analytical Local Valuation VaR Models	16
2.3.2. Simulation-Based Full Valuation VaR Models.....	18
2.3.3. Hybrid VaR Models	21
2.4. Essential Parameters and Assumptions in Standard VaR Models.....	22
2.4.1. Base Currency	22
2.4.2. Confidence Level	22
2.4.3. VaR Horizon.....	23
2.4.4. Constant Portfolio Composition.....	23
2.4.5. Standard Distributional Assumptions and Parameters	24
3. Long-Run Value at Risk Models.....	26
3.1. Overview.....	26
3.2. Simple Extensions of Short-Run VaR to Long Time Horizons	26
3.2.1. Scaling VaR Quantiles Directly with the Square Root of Time Rule	26
3.2.2. Scaling Volatility Estimates in VaR Calculations.....	29
3.2.3. Other Scaling Approaches for VaR Quantiles and Models.....	41

3.3. Specific Long-Run Value at Risk Models	54
3.3.1. “A tale of two VaRs”	54
3.3.2. MaxVaR	56
3.3.3. Long-Term Value at Risk	58
3.3.4. A Term Structure of Risk	61
3.3.5. Wavelet VaR	65
3.3.6. RiskMetrics Long-Run VaR Methodologies	66
3.4. Reviewing the Critical Issues for an Optimal Long-Run VaR Model	78
3.4.1. Constant Portfolio Composition	78
3.4.2. Standard Distributional Assumptions and Parameters	78
4. Conclusion	80
Appendix	82
Abstract – English Version	82
Abstract – German Version	84
About the Author	86
Bibliography	87

Acronyms

AR	AutoRegressive
ARCH	AutoRegressive Conditional Heteroskedasticity
BIS	Bank for International Settlements
CAPM	Capital Asset Pricing Model
ECM	Error Correction Model
EVT	Extreme Value Theory
EWMA	Exponentially Weighted Moving Average
GARCH	Generalized AutoRegressive Conditional Heteroskedasticity
GARP	Global Association of Risk Professionals
GPD	Generalized Pareto Distribution
I-GARCH	Integrated Generalized AutoRegressive Conditional Heteroskedasticity
L-VaR	Liquidity Value at Risk
LM-ARCH	Long Memory AutoRegressive Conditional Heteroskedasticity
NYSE	New York Stock Exchange
POT	Peaks Over Threshold
SEC	Securities and Exchange Commission
TARCH	Threshold AutoRegressive Conditional Heteroskedasticity
VaR	Value at Risk
VARM	Vector AutoRegressive Model
VECM	Vector Error Correction Model

Notation

0 present time, i.e. evaluation time
T time units in the VaR horizon
t any point in time between 0 and T
c VaR confidence level
α_c standard normal variate (quantile) for chosen confidence level
μ drift term in mean log portfolio (asset) return
σ volatility of log portfolio (asset) return
$r_{t,T}$ portfolio (asset) return from t to T
\bar{r} mean log portfolio (asset) return [$\bar{r} = \mu - \frac{\sigma^2}{2}$]
ε residual, error or innovation in a portfolio (asset) return process
P_t portfolio value at time t
P_0 portfolio present value at time 0
P_T^c quantile of portfolio value distribution at time T for chosen confidence level

Unless otherwise mentioned the equations listed in this thesis will be in line with the listed variable definitions.

Acknowledgments

First and Foremost I want to thank my parents Heinz and Jacqueline Stark for their patience and support. As is the case with most positive things I have achieved in my life to date, the completion of this thesis and the associated master's degree is a testament to the upbringing, education, and support I have been lucky enough to receive from them. Their confidence in me and my abilities has been a driving force throughout the entire process.

With regard to my academic studies at the University of Vienna I would like to take this opportunity to express my gratitude towards professors Engelbert Dockner, Christian Keber, Erich Streissler and Josef Zechner for sparking my initial interest in the fields of finance and financial economics and then for enabling me to gain an extensive and solid understanding, knowledge, and academic foundation in these fields through the attendance and completion of numerous courses held by either the professors themselves or other lecturers of the respective departments and institutes.

As far as this master's thesis itself is concerned, I want to thank Michael Halling for providing me with invaluable input and feedback during the early stages of the process, including the initial selection of the chosen topic. Additionally I would like to express my deep appreciation towards Professor Christian Keber for providing supervision and support for the "project" of writing this thesis during its latter stages up to its final completion. A very special thanks goes to Dr. Susan Etlinger for contributing her time and vast expertise with regard to scientific publications and the English language in general to optimize the linguistic quality of the text.

I would also like to thank Martin Mayr and Josef Flandorfer of the treasury department of Telekom Austria Group for providing me with greatly appreciated support and flexibility, which allowed me to complete this thesis while working full time throughout the entire process. Furthermore, I want to mention my current and former colleagues both at the University of Vienna (the legendary "BWZ Finance Crew") and in the corporate finance department of Telekom Austria Group who have made my time at work and at the university extremely interesting, enjoyable, competitive and productive.

Last but not most definitely not least, I would like to thank my family and friends who played an essential role throughout the writing of this text. On the one hand, by providing much appreciated motivation and impetus to remain focused on completing the thesis, and on the other hand, by providing welcome diversions during the time not spent writing.

Thank you all very much!

1. Introduction

The concept of Value at Risk (VaR) has been around for more than fifteen years and has undergone vast advances from its conception up to present times, however the time horizon of interest in VaR models has for the most part remained very short.¹ The reasons for limiting the perspective of VaR calculations to a few days can be ascribed to the fact that, for the longest time the only institutions which implemented VaR models were banks and other financial institutions whose primary perspective in their risk management functions is for short horizons. The reasons for this emphasis on short time horizons are diverse, but some key factors are the growing complexity and volatility of financial instruments and markets, and related financial "disasters" on the one hand² and the regulatory requirements set forth by institutions such as the Securities and Exchange Commission (SEC), the Bank for International Settlements (BIS), and the Global Association of Risk Professionals (GARP), which banks must adhere to.³ In recent years however, the user group of VaR measures has expanded extensively; in the financial sector different investment and pension funds have had increased demand for such measures, but this demand has also reached the risk management functions of non-financial corporations. These new user groups have turned to Value at Risk to assess market risks in their portfolios and contrary to the initial user group have fundamentally different priorities as to which time horizons are of interest to them. Non-financial corporations usually report their financial statements quarterly and or annually and are therefore interested in potential volatility of portfolio values at these time horizons. Due to the nature of their business, pension funds tend to look even further into the future, as both their assets and their liabilities are usually defined for very long time horizons going far beyond a single year.⁴ With these new users comes a demand for VaR models that can provide reliable measures of market risk for the horizons which are relevant for these users.

¹ The first time the term "Value at Risk" was mentioned in a publication was in the Group of Thirty report of 1993. For a detailed history of the development of Value at Risk please refer to Holton (2002 & 2003), Field (2003) and Jorion (2001).

² For information on the evolution of derivatives see Chance (1995 & 1998) and for examples of such financial disasters refer to Stulz (2000) or part IX of Field (2003).

³ Further information on these regulations and guidelines can be found in Dale (1996) and Dimson & Marsh (1995) pertaining to the SEC, Wagster (1996) and Crouhy et al (2006) concerning the BIS and Middleton (1996) and Fight (2003) on GARP.

⁴ Bodie et al. (2005), pp. 874-880

This is exactly the topic which this thesis focuses on: presenting means of obtaining valid long-run Value at Risk figures. To facilitate this core topic the reader will initially be familiarized with the basic VaR methodology and standard short-run Value at risk models as presented in the definitive VaR and risk management literature. In a second step the shortcomings of these short-run models with regard to their applicability for longer time horizons will be identified and discussed. Using the VaR foundation and these highlighted factors as a starting point, the main part of this thesis will present different means of calculating long-run VaR, ranging from simple extensions of short-run models to specific and in some cases highly complex long-run Value at Risk models, as present in up-to-date finance publications. After a survey of these different methods of calculating long-run Value at risk, the final part of the thesis will again turn its attention to the initially identified key factors for long-run VaR and attempt to summarize feasible and, where possible, optimal solutions, sketching an ideal long-run Value at risk methodology encompassing necessary models, parameters, and assumptions.

2. Laying the Groundwork

2.1. Overview

This part will provide the reader with a solid basis for the ensuing treatment of long-run Value at Risk, the core topic of this thesis. This will be achieved by initially presenting the conceptual framework of Value at Risk, including the basic elements contained in any VaR model or calculation method. After the brief general discussion of VaR, the standard short-run Value at Risk models contained in the normative literature on financial risk management will be described to outline the foundation upon which any extensions or modifications for long time horizons are built upon.

Drawing from these standard short-run VaR models, the final part of this part of the thesis will highlight those parameters and assumptions present in these models which are most relevant and oftentimes problematic when one is concerned with VaR models for the long run.

2.2. The Value at Risk Methodology

Before describing specific means of calculating Value at Risk, a general description of the basic concept which underlies all of the different VaR models has to be presented. Fundamentally, Value at risk is a measure of financial market risk.⁵ The primary source of such market risk in financial markets is the fluctuation of market prices and rates either in absolute terms or as compared to a predetermined benchmark.⁶ Laubsch and Ulmer (1999), however, specify so-called residual market risks which are also sources of market risk, although not directly caused by fluctuations in market prices.⁷

Value at Risk, which is conceptually rooted in modern portfolio theory including the Capital Asset Pricing Model (CAPM), going back to the landmark publications by Markowitz (1952),

⁵ Jorion (2001, p. 15) distinguishes between five categories of financial risks: market risk, credit risk, liquidity risk, operational risk, and legal risk. Value at risk is primarily concerned with market risk although the boundaries between categories are often blurred (Holton, 2003, p. 21).

⁶ Longerstaey (1996), page 24

⁷ For example: spread risk (the risk caused by changes in the difference between two asset prices), basis risk (the risk caused by changes in the correlation of different asset prices), or volatility risk (the risk of changes in volatility of asset prices and not by the level of volatility itself).

Roy (1952), Tobin (1958), Treynor (1999), and Sharpe (1964) among others, provides a methodology of quantifying such market risks for single assets and large portfolios alike.⁸

Two formal definitions of Value at Risk are:

“Value-at-Risk is a measure of the maximum potential change in value of a portfolio of financial instruments with a given probability over a pre-set horizon. VaR answers the question: how much can I lose with x% probability over a given time horizon.”⁹

Jorion (2001) defines VaR as a measure for the worst loss over a target horizon with a given level of confidence, or “More formally, VaR describes the quantile of the projected distributions of gains and losses over the target horizon.”¹⁰

From these definitions the three main cornerstones of any VaR calculation can be inferred:

- The confidence level or probability, which represents the level of statistical confidence of the calculated VaR figure.
- The forecast horizon of the VaR calculation, which is the focus of this paper.
- And the base currency in which the value, i.e. the Value at Risk is measured. This is usually the home currency of the investor or company calculating VaR.¹¹

With these three parameters defined, one can specify the components of the VaR methodology which represent the elements or steps necessary for any VaR calculation independent of the model which is chosen to calculate VaR.¹²

After having defined the confidence level, forecast horizon, and time horizon for the VaR calculation, the first step is marking to market the portfolio at hand. That is, calculating the present value of the portfolio, denominated in the chosen base currency.

In most cases this present value serves as the benchmark to which the potential portfolio value fluctuations are compared.¹³

⁸ Holton (2003), pp. 15-16

⁹ Longestaey (1996), page 6

¹⁰ Jorion (2001), page 22

¹¹ Laubsch & Ulmer (1999), pp. 10-11

¹² The following listing of these VaR elements is based on Jorion (2001), pp. 108, Longestaey (1996), pp. 105, and Laubsch & Ulmer (1999), page 105.

Either during the process of calculating the present value or afterwards, one must identify and select the risk factors which will be considered when estimating the fluctuation of portfolio value over the VaR horizon. A straightforward choice is selecting every market rate or price which was necessary or used for the initial present value calculation, for example the price or price change (i.e. return) of every single asset contained in the portfolio. Beside the fact that modelling every asset price as a distinct risk factor would make calculations for large portfolios extremely cumbersome¹⁴, in many cases it is neither necessary nor the optimal choice. This is precisely where the next essential element of a VaR calculation, mapping, comes into play.

Mapping encompasses the identification of relevant and necessary risk factors and the specification of the functional dependencies of future asset values in relation to these risk factors for all assets contained in a portfolio. The main aim of mapping is to reduce the number of risk factors to a manageable number whilst maintaining accuracy of the resulting VaR calculation. The following diagram contained in Figure 1 depicts a graphical representation of mapping to illustrate the purpose of mapping in VaR calculations.

¹³ There are two exceptions: On the one hand a relative VaR can be calculated by comparing the portfolio value fluctuation to a predetermined benchmark portfolio; on the other hand one can relate these value fluctuations to the expected portfolio value at the end of the VaR horizon. The latter issue will be covered in the next part of the thesis.

¹⁴ In extreme cases it can make a VaR calculation close to impossible because the derivation of a valid correlation matrix becomes extremely difficult for large numbers of risk factors (Jorion (2001), page 168).

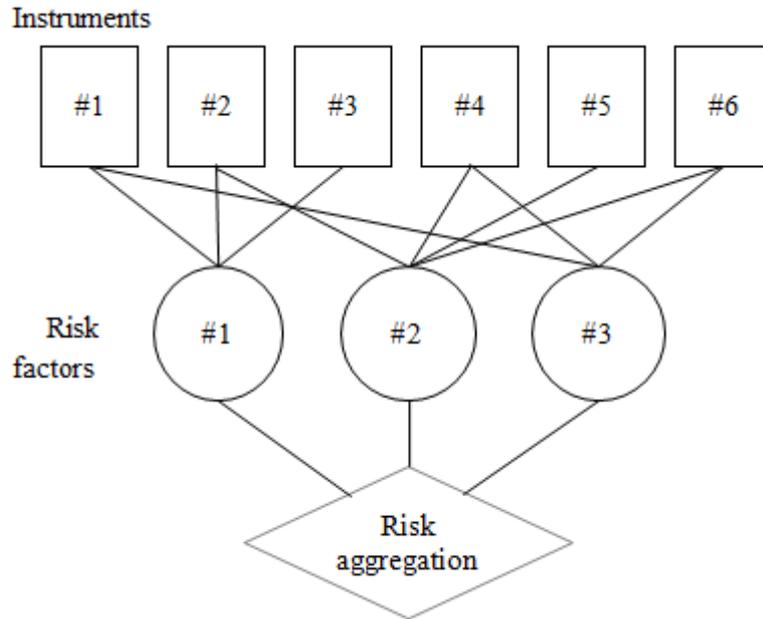


Figure 1 – Mapping - Source: Jorion (2003), page 376

Holton (2003)¹⁵ devotes two chapters to an extensive discussion of mapping in general and various techniques, but in the treatment at hand only two of the most common examples for mapping will be presented. The most prominent mapping technique presented in the RiskMetrics documents is cash flow mapping to model interest rate risk.¹⁶ In this mapping technique a select number of key interest rates in the form of zero coupon rates or discount factors for the essential maturities are selected as risk factors. Then every asset or instrument in the portfolio is dissected into the cash flows it generates and these cash flows are mapped into so-called “time-buckets”¹⁷ whereby each key rate maturity represents one of these buckets. After mapping all assets in the portfolio into these time buckets, the sensitivity of the entire portfolio with respect to each of the risk factors is identified in an aggregated form and any Value at Risk calculation is simplified by substantially reducing the number of risk factors.

A second example of a mapping technique for VaR calculations is index mapping, which is primarily applied to equities. In this mapping technique, the prices or returns on equity indices¹⁸ are selected as risk factors and the individual stocks mapped onto these indices by

¹⁵ Chapters 8 and 9 in Holton (2003), pp. 275-354

¹⁶ Longerstaeay (1996), pp. 117-121; Laubsch & Ulmer (1999), pp. 105-110; Mina & Xiao (2001), pp. 41-46

¹⁷ Longerstaeay (1996), page 39

¹⁸ These indices can be determined geographically, by sector or even a combination of the two (Laubsch & Ulmer (1999), pp. 108-109).

calculating the beta factors for each stock with respect to the relevant index in line with the CAPM. As can be imagined for portfolios containing large numbers of different stocks this can provide considerable advantages with regard to the complexity of the VaR model. There are also similar mapping approaches for other asset classes which follow the same basic principles.

After having identified both the relevant risk factors themselves and the exposures to these risk factors for the portfolio under investigation, the next step is estimating the development of these risk factors up to or over the selected VaR horizon.

This element is basically the cornerstone of the whole VaR methodology, as the accuracy and validity of these estimates ultimately determine the quality of the resulting VaR figures. Therefore, risk factor forecasting or estimation for long time horizons will be one of the focal points of the next part of this thesis.

Determining the risk factor estimates differs both in complexity and in the actual “mechanics” of deriving them for the various VaR models and will therefore be covered in the next section. The three components of this step in the VaR process, however, are the same for all models. Based on market data¹⁹, an assumption is made which distribution the risk factors, or their changes, follow. This assumption is oftentimes made implicitly by default due to the specific VaR model choice. The most common choice is the standard normal distribution and this choice will be discussed extensively throughout the thesis.

The next component is determining the specific risk factor distributions based on the market data, either by estimating the necessary distributional parameters or by estimating the entire risk factor distribution, which is the case in historical simulation models. And finally, these risk factor distributions are used to generate forecasts for the developments of the risk factors, for the category of simulation-based VaR models this is referred to as scenario generation.²⁰

The penultimate element in a VaR calculation is determining the effect of the forecast risk factor fluctuations on portfolio value at the end of the VaR horizon. There are two classes of

¹⁹ For example, directly observable current spot market prices for the risk factors, derivatives based on these assets, historical developments of these market rates, or other financial or economic data.

²⁰ Longerstaey (1996), page 151

VaR models which distinguish themselves as to the way in which this effect on portfolio value is calculated, local valuation and full valuation models.

The difference between the two is that whereas local valuation models only approximate the actual portfolio value at the VaR horizon, the full valuation models perform a full portfolio revaluation for the forecasted risk factor scenarios. Local valuation methods have advantages regarding computational intensity and simplicity, because for these methods one must only calculate or estimate sensitivities of the portfolio or the contained assets, respectively, with regard to the selected risk factors once for a given asset, which is usually done in line with the mapping procedure. The effect of the risk factor fluctuations on portfolio value can then be projected relating the sensitivities to the estimated risk factor scenarios. The main advantage of local valuation is that the sensitivities only need to be calculated once and can then be used for calculating the effect of multiple different estimated risk factor scenarios, the disadvantage is that the method only calculates approximations of the real effect of the risk factor changes on portfolio value.²¹

Full valuation methods, however, perform a full revaluation with the actual pricing functions for the portfolio assets, taking the previously identified mapping relationships into account. Such a revaluation must be performed specifically for every single different risk factor scenario so that the increase in precision comes at the price of a greater computational load.

Once the effect of the estimated risk factor fluctuations on portfolio value has been estimated one can determine the “worst case” risk factor scenario which is in line with the specified VaR confidence level and in the final step of any VaR calculation, this “worst case” portfolio value is related to the portfolio’s present market value to determine the “worst case” loss for the selected horizon and confidence level which is the quantification of market risk that a Value at Risk figure represents.

Summarizing these elements and steps in any VaR calculation:

1. VaR parameter definition (currency, horizon, confidence level)
2. Calculation of the portfolio’s present value (marking to market)
3. Mapping (assets and exposures to risk factors)

²¹ In the most common case, a linear approximation of the actual pricing functions is applied. For details please see Longerstaey (1996), page 26.

4. Risk factor estimation or forecasting (distributions, parameters and scenarios)
5. Revaluation of the portfolio (local or full valuation in line with steps 3 and 4)
6. Calculation of the VaR figure (relating “worst case” value to relevant benchmark)

These steps are necessary for every VaR calculation and therefore present in any VaR model, whereby the first two steps are exactly the same for every model, but the other four steps are where the models differ.

Before focussing on the specific VaR models, a brief and very simple numerical example will be provided to promote a basic understanding of the concepts just presented.

- An investor located in Euroland holds a portfolio containing 10 shares of a specific company’s stock whose present value is EUR 10 per share.
- The present value of the portfolio is therefore obviously EUR 100.
- It is assumed that the expected one-day return for these stocks is zero percent and the expected volatility of the one-day returns is 10 %.
- For additional simplification it is also assumed that these stock returns are drawn from a normal distribution.

By taking these assumptions into account one can state that the expected portfolio value in one day is EUR 100.

A one-day Euro VaR figure can provide an estimate of a potential loss in portfolio value which will not be breached with a specific level of confidence expressed in a probability percentage.

Again drawing from the just presented information and assumptions one can calculate a one day Euro 95% VaR by calculating that one day portfolio return which will be surpassed with a probability of 95% (i.e. the probability of the return falling below this return is 5%) and then relating the corresponding portfolio value to the portfolio’s present value of EUR 100.

Applying the simplest VaR model which will be covered in the next section, one can calculate the 95% “worst case” portfolio return by multiplying the volatility of the one-day return of 10% with the 1-95% quantile of the standard normal distribution which is equivalent to -1.65.

The 95% “worst case” one-day return is therefore -16.5% and the corresponding 95% “worst case” portfolio value in one day amounts to EUR 83.5.

This yields a one-day Euro 95% VaR of EUR 16.5 which is the difference between the portfolio’s present value and the 95% “worst case” portfolio value in one day’s time.

This VaR figure gives the investor the information that with a probability of 95% the highest loss in portfolio value which can be incurred in one day is EUR 16.5 presenting a measure for the risk potential of the investment.

After having presented this simple example as an illustration of the basic VaR concepts, the different standard short-run VaR models will be covered in the next section, whilst highlighting the simplifying assumptions contained or employed in these models.²²

2.3. Standard Short-Run Value at Risk Models

Aside from the one primary distinction or classification of VaR models into local valuation or full valuation models, there is one other general differentiation for Value at Risk models.

One distinguishes between parametric or analytical models²³ on the one hand and simulation-based approaches on the other hand. Analytical VaR models all assume the risk factors to be drawn from a simple parametric distribution²⁴ and consist of a single specific function which is formulated and can then yield a precise VaR figure for a given set of input parameters. The VaR calculations are usually fairly simple and by definition they always apply local valuation. Simulation VaR models are more labor intensive. These entail calculating an entire risk factor distribution at the end of the VaR horizon and then deriving the associated portfolio value distribution via full or local valuation. These models can basically incorporate any kind of risk factor distribution.

²² While the different standard short-run VaR models and classes of models will be surveyed, they will not be explained in great detail as the focus at hand is highlighting the inherent assumptions and limitations with respect to an application of these models for long time horizons. For detailed explanations and derivations of the short-run models, the reader should refer to any of the cited standard VaR or RiskMetrics publications.

²³ Jorion (2003), page 371

²⁴ The standard normal distribution is the most common and straightforward choice.

2.3.1. Analytical Local Valuation VaR Models

2.3.1.1. Delta Normal VaR

The simplest of all VaR models or methods is the delta normal VaR. As the name implies, this method makes use of linear mapping procedures and of the normality assumption for the relevant risk factors.²⁵ This assumption has tremendous benefits in relation to simplifying the corresponding calculations.²⁶

As far as the process of calculating a delta normal VaR goes, the first two steps are standard. In the mapping process the assets are then mapped linearly to the risk factors to determine the exposures. This yields an exposure vector of so-called “delta equivalents”²⁷.

The next element of risk factor estimation is also highly simplified. This model assumes that the fluctuations or changes of the selected risk factors are normally distributed with a mean value of zero. Therefore the only parameter necessary to define the single risk factor distributions is the volatility, which is usually calculated from daily historical market returns and scaled to the defined VaR horizon. In addition, the other necessary element relating the distinct risk factor distributions amongst each other is a correlation matrix, also derived from historical market returns.

The model therefore introduces some significant assumptions in this element: risk factor fluctuations which follow a normal distribution with a constant mean return of zero, constant volatilities and constant correlations across the entire VaR horizon.

These risk factor forecasts are used to generate a covariance matrix for all determined risk factors and the defined VaR horizon. The actual VaR calculation then is simply a matrix multiplication of the exposure vector and the covariance matrix which produces a standard deviation of the portfolio value in units of the defined currency. VaR is finally computed by multiplying this standard deviation with the appropriate quantile of the standard normal distribution representing the defined confidence level. The simplicity of this model has

²⁵ Jorion (2001), page 206

²⁶ Linear combinations of normally distributed variables are also normally distributed; therefore, if one assumes risk factors to be normally distributed and exposures are mapped to these risk factors via linear relationships, then the distribution of portfolio value is also normally distributed (the invariance property of normal variables – from Jorion (2001), page 206).

²⁷ Mina & Xiao (2001), page 21

advantages when considering computational intensity and the modification of parameters such as the confidence level or the time horizon, as these can be changed quickly and one must only perform a simple multiplication to obtain the new result. However, this high degree of simplicity restricts a valid application of the model significantly. Even for short time horizons and the associated smaller expected fluctuations of risk factors, both the linear mapping and the distributional assumptions are problematic in many cases, but for long time horizons they would have to be deemed implausible in most cases.

2.3.1.2. Delta Gamma VaR (and other “Greeks”)

These VaR models are extensions of the delta normal VaR method. Whereas the delta normal VaR applies linear approximations of the relationships between asset values and risk factors in the mapping process, for example in the form of the first derivative of the pricing functions with respect to a risk factor, the delta gamma and related VaR models expand the approximation to include second derivatives of the pricing function and other sensitivities to produce a more accurate measure of the value fluctuations with regard to risk factor changes.²⁸ Mathematically this is achieved by employing a higher order Taylor series expansion of the pricing functions.²⁹

As far as the actual VaR calculation is concerned, the only difference with regard to the delta normal VaR is the increased complexity of the mapping and the corresponding risk factor estimation elements. In the most common delta gamma models two additional parameters are required to calculate VaR for every risk factor, however, this group of models can become very complex if one includes additional sensitivities. The simplest of these models is the delta gamma delta method which makes use of an additional simplifying assumption which makes the resulting VaR equations manageable. The assumption being that the first order derivative of the pricing function with respect to the risk factor and the second order derivative are joint normally distributed.³⁰ Once the parameters and matrices have been derived, the actual VaR calculation is once again very straightforward and equivalent to the delta normal VaR.

²⁸ Examples are the following: for a fixed income asset the delta normal VaR would use duration as a mapping function and the equivalent delta gamma VaR would also include convexity. For equity options the delta normal VaR would be expanded to include gamma or any of the other so called “greeks”. For more information on these Greeks please refer to chapter 14 in Hull (2003), pp. 299-329.

²⁹ This process is explained in Holton (2003), page 63.

³⁰ Jorion (2001), page 213

However, the advantages with regard to computational intensity that this method can have over full valuation methods can quickly disappear for large portfolios with different complex assets which lead to a dramatic increase in the number of parameters which must be calculated.

2.3.2. Simulation-Based Full Valuation VaR Models

2.3.2.1. Monte Carlo Simulation VaR³¹

Monte Carlo simulation VaR models are the most flexible VaR methods. They can basically model any and every risk factor distribution conceivable and provide a complete distribution of portfolio values subject to the different risk factor scenarios via full valuation. The precision of the yielded results depends solely on the assumptions underlying the simulations and the extent of simplification caused by the mapping procedures.

Naturally this flexibility and precision must be weighed against the computational intensity in real world VaR applications or implementations, to arrive at an optimal tradeoff for each specific case.

Just as in the analytical VaR models, the first two steps in the process are also the same in the simulation-based approaches. But the remaining elements are notably different from those in the analytical models. For one, mapping is not restricted to linear or second order approximations of the actual relations between risk factors and asset values as the portfolios are completely revalued for every risk factor scenario. The main difference, however, arises in the fourth step of risk factor forecasting. As already mentioned, contrary to the analytical methods, there are no restrictions as to which statistical distributions can be assumed. The only factors that must be kept in mind are the computational intensity and accuracy related to any such choice. Each distribution requires different parameters to be calibrated from market data and the subsequent generation of forecast scenarios from these distributions. Therefore the choices which are simplest to implement, whilst reflecting the actual distributions most accurately, must be preferred.

³¹ Monte Carlo simulations go back to Stanislaw Ulam in 1946 (Eckhardt (1987) pp. 131-137) and are used in various fields of science to find solutions numerically by producing scenarios and evaluating the results obtained for every scenario. In finance the other major application is for asset valuation of derivative assets such as complex options. For detailed information on Monte Carlo Simulations in general please refer to chapter 5 in Holton (2003), pages 193 to 225.

Despite the high degree of flexibility that Monte Carlo VaR models can potentially incorporate, the standard Monte Carlo models, as presented in Mina and Xiao (2001)³², make use of the same simplifying assumptions that are incorporated in the standard analytical models. The risk factor fluctuations or returns are assumed to follow a normal distribution with a constant mean of zero, constant volatility and correlation across the entire VaR horizon.

These assumptions facilitate the application of the geometric Brownian motion as the process from which the forecast risk factor fluctuations are drawn to simulate the necessary scenarios. The actual simulation runs are calculated in line with the following basic equation:³³

$$r_{t,T} = \left(\mu - \frac{1}{2} \sigma^2 \right) (T - t) + \sigma \varepsilon \sqrt{T - t}$$

This presents the statistical process for the single risk factor fluctuations (returns), where in the standard model the mean return is assumed to be zero, which reduces the expected return to a function of the volatility estimate, scaled by the relevant time horizon and multiplied by a random number which is drawn from the standard normal distribution.³⁴

In a framework where there are multiple risk factors, i.e. a multivariate setting, these distinct risk factor processes must be correctly related to each other. This is achieved by estimating a constant correlation matrix from historical market data just as in the standard analytic VaR methods.³⁵ In the process of the actual risk factor scenario generation, this correlation matrix is initially “decomposed”³⁶ and then applied to the random variables in order to essentially impose the estimated correlation structure onto the generated risk factor scenarios. At the end of the fourth step in the Monte Carlo VaR process, one has generated the required number of

³² Mina & Xiao (2001), pp. 18-21

³³ Mina & Xiao (2001), page 14

³⁴ These random numbers are referred to as innovations, error terms, or residuals and are the source of the random component in the simulation models.

³⁵ It is important to keep in mind that a correlation matrix is not directly dependent on the VaR horizon, but a covariance matrix which combines the correlations with volatilities is specified for a specific horizon just as single volatility estimates are.

³⁶ Basically the matrix is split into a triangular matrix by means of a so called Cholesky decomposition or Single Value decomposition, which has the necessary properties required for the actual scenario generation. For details on these elements of the Monte Carlo simulation please refer to Jorion (2001) pages 303-304 or Mina and Xiao (2001) page 19.

scenarios, each encompassing a simulated value for every risk factor as at the end of the VaR horizon.

The next element is then simply calculating the portfolio value for every simulated scenario by revaluing the assets assuming the forecast risk factor values, which yields a complete distribution of portfolio values as at the end of the VaR horizon.

In the sixth and final step of the VaR calculation the sought-after VaR figure can be read out of the portfolio value distribution by selecting the quantile which corresponds to the chosen confidence level and comparing this value to the relevant portfolio value benchmark, which in the standard model is the portfolio's present value.³⁷

Although the tool of a Monte Carlo simulation provides means to create highly complex and flexible models which are fundamentally restricted only by the available computational power, the standard models subscribe to most of the simplifying assumptions which are also present in the simple analytical models.

2.3.2.2. Historical Simulation VaR

The historical simulation VaR methods are also full valuation models and can be seen as special implementations of a Monte Carlo simulation. The first three and final two steps follow the same footsteps as in any Monte Carlo simulation, whilst the difference appears in the process of estimating and forecasting the risk factor distributions.

Contrary to Monte Carlo simulations, in which statistical distributions are chosen and the distributional parameters are then calibrated based on the market data, historical simulations use the actual historical return data directly for the forecast scenarios. For each simulation run an independent non-overlapping historical sample of the same time span as the VaR horizon is necessary.³⁸ This is the greatest drawback of the method, as the sample size of relevant or even available data is restricted even at a one day horizon, and this problem increases dramatically when the VaR horizon is extended. Although Mina and Xiao (2001) present two possibilities of increasing the sample size for multi-day VaR horizons, neither provides

³⁷ This approach is called "percentile ranking method". Das (2006) pp. 95-96

³⁸ This implies that, to generate 500 runs for a one day VaR horizon one needs 500 distinct historical one day returns which, assuming 250 trading days in a year, is equivalent to two years of market data.

satisfactory relief.³⁹ This issue will be discussed further in the next part of the thesis as there are some long-run VaR approaches based on historical simulation. An additional drawback of historical simulations is that the time period from which the data is sampled must be deemed representative for possible future developments, so that if the historical data covers a period of uncharacteristically low or high volatility the derived VaR will most certainly be systematically flawed. For these reasons, in any historical simulation a compromise between the accuracy (larger number of scenarios) and the quality of representation of the expected risks (shorter sample timeframe) must be found.⁴⁰

In spite of these shortcomings, historical simulations have some significant advantages over Monte Carlo simulations. By abstaining from making any distributional assumptions, the drawbacks of any distributional misspecifications can be avoided. The actual historical distributions of market data obviously incorporate any special empirical distributional characteristics which might not be reflected accurately in a standard statistical distribution such as the normal distribution.⁴¹

The actual scenario generation and VaR calculation in a historical simulation model is performed by randomly selecting the necessary number of scenarios from the dataset of daily historical returns and performing a full valuation of the portfolio subject to each scenario, which just as in a Monte Carlo simulation provides a full distribution of portfolio values from which VaR can be determined.⁴²

2.3.3. Hybrid VaR Models

In addition to the local valuation analytical VaR models and the full valuation simulation based VaR models there are VaR models which combine features from these two model classes which are therefore referred to as “hybrid” models in this thesis.

³⁹ For details on these two options please refer to Mina and Xiao (2001), pages 26 to 27.

⁴⁰ Jorion (2001), pp. 223-225

⁴¹ Examples which are commonly discussed in finance publications are the skewness and heavy tails present in financial time series. For an overview of the history of return distributions, please see Malevergne and Sornette (2006), pages 37 to 43.

⁴² Das (2006, pp. 95-96) also suggests an alternative for calculating a VaR figure from the simulated portfolio value distribution, in which, similar to the delta normal method, a sample mean and standard deviation are calculated and VaR is computed from these parameters.

Two such VaR models are the Delta-Gamma-Monte Carlo method and a so-called Grid Monte Carlo method. The Delta-Gamma Monte Carlo method is no different from a standard Monte Carlo VaR model except for the fact that for every simulated risk factor scenario only a local valuation based on the Delta-Gamma models is performed. This procedure could also be applied when scenarios are derived via historical simulation. The grid Monte Carlo method on the other hand reduces the computational load by only performing full valuations for a small number of risk factor scenarios, so called “grid points” and then applies linear interpolation of portfolio values between these grid points.⁴³

These hybrid methods try to provide a compromise solution in the tradeoff between the accuracy of simulation-based approaches and the computational efficiency of local valuation models.

2.4. Essential Parameters and Assumptions in Standard VaR Models

In the final section of this part, the parameter choices, explicit, and implicit assumptions in the standard short-run VaR models that were just presented will be summarized and critically evaluated with regard to their validity and applicability in long-run VaR calculations.

The first three parameters are those which define any VaR figure and must be defined in the first step of any VaR calculation: the base currency, confidence level, and time horizon.

2.4.1. Base Currency

The choice of base currency can in general be deemed straightforward and unproblematic, irrespective of whether one is concerned with a short or a long time horizon.

2.4.2. Confidence Level

The choice of confidence level on the other hand can pose problems or, at the very least, must be kept in mind when extending the VaR horizon significantly. On the one hand some models such as the historical simulation can become prone to large estimation errors for very high confidence levels with this problem increasing as the sample size decreases due to longer sample intervals. On the other hand, different VaR models achieve different levels of

⁴³ Jorion (2001), pp. 210-215

precision for the different levels of confidence, so that this issue should be kept in mind when choosing the VaR model and determining the confidence level to be calculated.⁴⁴

2.4.3. VaR Horizon

The third of these parameters goes to the heart of this thesis. The VaR horizon in standard models usually lies between a single day and one month, with the most common choices being 1 day, 5 days, or 10 days - the time horizon set forth in the BIS's regulatory requirements. The time horizons this paper focuses on are of a much greater dimension than the standard short-run choices, specifically, horizons in the range between 3 months and 2 years where the focal point lies at a time horizon of 12 months.⁴⁵

The fact that the time horizons are not limited to a few days is the dominant reason why most of the standard assumptions and parameters cannot be deemed feasible in long-run VaR calculations.

2.4.4. Constant Portfolio Composition

This is one of the critical, implicitly made assumptions present in all standard short-run VaR models. The portfolio is evaluated and marked to market in the second step of any VaR calculation and all ensuing evaluations of the effect of market rate fluctuations are performed assuming the portfolio composition remains unchanged throughout the VaR horizon.⁴⁶

This assumption might seem trivial and unproblematic for the short horizons of 1 to 10 days, but for longer time horizons some critical issues arise in connection with this assumption.

Many if not most assets yield intermediate cash flows during their existence, some at predefined and others at unforeseeable points in time.⁴⁷ Other assets, on the other hand, can require additional inflows during their term, where a textbook example is a margin call in

⁴⁴ The most common choices are 95% and 99%.

⁴⁵ These time horizons are of relevance for the primary user groups of such long run VaR figures, as will be explained in the next part of the thesis.

⁴⁶ Composition remains constant in the sense that the number and/or nominal value of the assets which are contained in any portfolio stay fixed. Obviously the composition as far as market values are concerned can change in different scenarios.

⁴⁷ Standard examples include equity dividends or fixed income coupon payments.

connection with a futures contract.⁴⁸ Such cash flows could cause substantial shifts from or into a potential cash position in a portfolio or even a liquidation of assets to cover potential cash shortfalls. Although this aspect can be neglected for a one day horizon because portfolio revaluations usually take place at daily intervals, they cannot be ignored for longer time horizons, let alone horizons of a year or longer.

A closely related issue is that very many financial assets have a specific term and an either fixed or maximum maturity date.⁴⁹ For a twelve month time horizon the probability of a portfolio containing such assets is obviously relatively high. These are just a few examples of possible events that could force a change in portfolio composition, the other element that certain underlying strategies or restrictions might lead to an active change in composition.⁵⁰ Holton (2003) refers to such occurrences as “intra-horizon events”.⁵¹

2.4.5. Standard Distributional Assumptions and Parameters

As previously mentioned, all standard short-run VaR models except for the historical simulations employ the assumption that risk factor fluctuations (returns) are drawn from a multivariate normal distribution with a constant mean of zero, constant volatility, and constant correlation. The normality assumption is not even without criticism for short horizons, but its validity for long-run VaR models must and will be evaluated. As far as the distributional parameters are concerned, the assumption which is most questionable is the zero mean assumption. This simplification goes against fundamental principles of financial theory, because if a risk-bearing asset has an expected return of zero, this would violate any risk-return considerations and remove any financial incentive of holding such assets.⁵² Assuming constant volatilities and correlations over such long time horizons is also by no means unproblematic, however as will be illustrated in the next part of the thesis, any deviations from these simplifications are quite complex. A related issue is how these distributional parameters, most importantly the volatilities, are estimated. The standard models rely solely

⁴⁸ Reilly & Brown (2003), page 908

⁴⁹ Most fixed income and derivative instruments have predetermined maturities, with some derivatives such as American options being exercisable at any point in time up to maturity.

⁵⁰ If a portfolio is managed in line with a predetermined strategy, certain events could lead to a required adjustment of the portfolio composition.

⁵¹ Holton (2003), page 290

⁵² Cuthbertson & Nitzsche (2005), pp. 124 & 381-383

on historical spot market data and calculate one day volatilities either by an equally weighted historical average or an Exponentially Weighted Moving Average (EWMA).⁵³ When estimating a volatility forecast for horizons in excess of a single day, the one day estimates are simply scaled to the required horizon by the so called “square root of time rule”.⁵⁴ This simple scaling approach can lead to very poor results and as the volatility estimates are without a doubt the most critical parameter in any VaR calculation, this issue will be thoroughly discussed in the next part. Furthermore, the development of different volatility estimation models and approaches has made significant advances over the past few years so that any implementation of these simple approaches should be carefully assessed.⁵⁵

This concludes both the current section, which highlighted the critical issues which must be considered when dealing with VaR for long time horizons, and also the second part which should provide the reader with a fundamental understanding of both VaR in general and the standard short-run VaR models and their deficiencies in particular.

⁵³ Mina & Xiao (2001), pp. 14-16; Longerstaeay (1996), pp. 78-82

⁵⁴ Longerstaeay (1996), page 87

⁵⁵ For examples of such models please refer to Dunis et al. (2005), pp. 129-292

3. Long-Run Value at Risk Models

3.1. Overview

After having presented and discussed basic short-run value at risk models to provide the reader with a solid understanding of the related theories and framework in the previous part of this thesis, this part will treat the core topic of Value at Risk for long time horizons.

At first the simple approaches for calculating long-run VaR figures will be surveyed and critically evaluated. These simple approaches generally make use of a short-run VaR approach or model and merely extend the time horizon to the desired length by applying a multiplier to the short-run VaR itself or to an input parameter of the model.

Then, in the second section of this part of the thesis, specific long-run VaR models will be treated, placing special emphasis on their critical differences compared to the simple approaches.

3.2. Simple Extensions of Short-Run VaR to Long Time Horizons

3.2.1. Scaling VaR Quantiles Directly with the Square Root of Time Rule

The simplest extension of short-run VaR to a longer time horizon is using a VaR figure calculated for a short time horizon such as a single day and then scaling up this 1 day VaR figure to the desired time horizon by multiplying it with the square root of the number of days in the desired VaR horizon.

$$VaR_{T \text{ day}, 99\%} = VaR_{1 \text{ day}, 99\%} \times \sqrt{T}$$

Interestingly, although the most simplistic, this approach is listed as permissible in the BIS documents on the Basel Capital Accord⁵⁶ which are regulatory documents setting forth rules for banking risk management systems.⁵⁷ Some of the first publications on VaR refer to this approach as the appropriate means of calculating a long horizon VaR and mention horizons up

⁵⁶ BIS (2005), page 40

⁵⁷ These recommendations and guidelines were adopted as law within the European Union in the summer of 2004, further increasing their relevance. (Brummelhuis & Kaufmann, 2007, page 39)

to one month as illustrations.⁵⁸ This method is also listed in other prominent publications on the subject.⁵⁹ It should be noted, however, that even the original RiskMetrics Technical Document of 1996, although mentioning this approach as allowed according to the “Basel proposals”⁶⁰, does not favor the technique. Danielsson, Hartmann and de Vries (1998) are very critical of some of the Basel proposals including the multiplier, going so far as to state that these proposals provide incentives to underestimate VaR as much as possible.⁶¹ Contrary to the BIS the United States Federal Reserve does not support this scaling method.⁶²

Although the Basel framework chooses a VaR horizon of ten days as its benchmark in which this crude extension might not deviate too much from the results of more precise methods, it does not restrict the use of this means of horizon extension to only ten days.

Fundamentally, this VaR scaling method is based on the more prominent and less restrictive approach of scaling volatility estimates to the desired VaR horizon by the square root of time rule⁶³ which are then used as inputs in VaR calculations.

These two methods are only equivalent, however, when calculating VaR with the simplest VaR model, the analytical Delta Normal VaR approach.

If:

$$VaR_{1\ day,99\%} = \sigma_{1\ day} \times 2.65$$

And:

$$\sigma_{t\ day} = \sigma_{1\ day} \times \sqrt[2]{t}$$

Then:

$$VaR_{t\ day,99\%} = (\sigma_{1\ day} \times \sqrt[2]{t}) \times 2.65 = VaR_{1\ day} \times \sqrt[2]{t}$$

When calculating the underlying short-run VaR with any other model, these two approaches are no longer equivalent.⁶⁴ In such cases, the obtained VaR figures are even more imprecise

⁵⁸ Smithson & Minton (1996a), pp. 25-27; Smithson & Minton (1996b), pp. 38-39; Jordan & Mackay (1997), pp. 292-294

⁵⁹ For Example: Jorion (2001), pp. 112, 252; Das (2006), page 50; Duffie & Pan (1997), page 34

⁶⁰ Longerstaey (1996), page 37

⁶¹ Danielsson et al. (1998), pp. 101-103

⁶² Laubsch & Ulmer (1999), page 12

⁶³ Jorion (2003), page 66

⁶⁴ Iacono & Skeie (1996), page 8

and the additional implicit assumptions (e.g. distributional stationarity)⁶⁵ of the approach are even more problematic than those inherent when scaling volatilities by the square root of time.

This specific fact is one of the points of criticism brought forth in the literature against this practise.⁶⁶

In their paper on time-scaling of risk, Daniélsson and Zigrand⁶⁷ discuss this approach of scaling quantiles of the portfolio value profit and loss or return distribution in great detail.

The authors argue that the underlying distributional assumptions made when scaling quantiles of a return distribution are even more restrictive than when scaling volatilities and that these assumptions “are violated in most, if not all, practical applications.”⁶⁸

Daniélsson and Zigrand go on to analyse the performance of the square-root-of-time rule (applied to quantiles) for returns following different data generating processes, such as GARCH⁶⁹, stable but non-normal or jump diffusion processes.⁷⁰ They find that applying this scaling of quantiles underestimates VaR for almost every assumed return process with the underestimation increasing with the time horizon and confidence level.

They also analyse the effect of time scaling when a positive drift is assumed and discover that even in such cases risk, i.e. VaR, is underestimated for longer horizons. However, their results show that when a positive drift is present, the square root of time rule does perform best when scaling to around ten days, due to the fact that at this time horizon the systemic underestimation caused by the approach is balanced by the overestimation of risk induced by a positive drift. This result could temper the criticism of the Basel guidelines to some extent.

One of the most interesting results of their analyses is that the shortcomings of time scaling of quantiles are not due to the potential insufficiencies caused by time scaling of volatilities. The shortcomings of this approach would arise on top of or independent of potential shortcomings due to time scaling of volatilities, which will also be discussed in this thesis.

⁶⁵ Culp et al. (1998), page 26

⁶⁶ Dowd (2005), pp. 157, 184

⁶⁷ Daniélsson & Zigrand (2006)

⁶⁸ Daniélsson & Zigrand (2006), page 2702

⁶⁹ GARCH = Generalized Autoregressive Conditional Heteroskedasticity

⁷⁰ For an introduction and detailed information on different data generating processes please refer to Schmid and Trede (2006).

Summing up their findings, the time scaling of VaR quantiles underestimates VaR in almost all analyzed cases and “configurations” and this underestimation is not caused by potential failures caused by time scaling of volatilities.

Contrary to these results Blake, Cairns and Dowd (2000) argue that the square root of time rule applied to VaR quantiles grossly overestimates the actual VaR even at short horizons such as five or ten days with the error increasing with the time horizon. Their results however are deduced from the fact that they assume a positive mean return for the portfolio, so that the two analyses are not directly comparable.⁷¹

In short, research results on direct scaling of VaR quantiles are by no means unambiguous.

3.2.2. Scaling Volatility Estimates in VaR Calculations

A more common and potentially more precise extension of short-run VaR figures or models to longer time horizons is the application of the square root of time rule in its original form. That is, scaling the volatility parameter used in a VaR model directly to the desired time horizon and then calculating VaR with this input variable.

This approach of extending the VaR horizon is definitely the most widely discussed method in the risk management and VaR literature.⁷² However, the square root of time rule extends beyond VaR and is notably present in other fields of finance.

This type of time scaling of volatility originates from the geometric Brownian motion⁷³, which is a continuous time stochastic process oftentimes chosen as the process which equity price changes are assumed to follow.⁷⁴ One of, if not the most prominent application of this assumption, is the Black-Scholes formula for the pricing of European equity options.⁷⁵ The Black-Scholes formula contains a volatility estimate which, similar to a VaR calculation, is the most critical input variable. This estimate can be derived by calculating a 1 day volatility

⁷¹ Blake et al. (2000), pp. 3, 7

⁷² For example: Jorion (2001) and Longerstae (1996)

⁷³ This process was introduced in the previous part of the thesis.

⁷⁴ Luenberger (1998), pp. 308-309

⁷⁵ Black & Scholes (1973), pp. 637-654

estimate and then scaling it by the square root of time to an annualized figure.⁷⁶ The time scaling of volatility estimates by the square root of time rule can also be found in other models of option pricing, such as Monte Carlo simulation for the pricing of complex derivatives which was mentioned in the discussions on the analogous VaR method.⁷⁷

Although slightly less restrictive than a direct scaling of VaR quantiles, applying the square root of time rule to volatility estimates does garner its share of criticism.

Diebold et al. (1998)⁷⁸ devote an entire article to a detailed analysis of volatility scaling via the square root of time rule and highlight the major weaknesses of this approach. The authors illustrate that such volatility scaling implies strict statistical assumptions which are consistently violated by asset return data and therefore lead to inappropriate long horizon volatility estimates. Furthermore they argue that this scaling of volatility does not generate volatility estimates which are generally too high, too low, or too conservative, but rather amplifies volatility fluctuations when it should in fact dampen the larger volatility fluctuations which are present in short horizons. However, Diebold and colleagues do not completely condemn the approach as totally unsuitable or biased to either produce estimates which are always too low or too high. They highlight the fact that in spite of its shortcomings, the square-root-of time rule produces results which are correct on average and due to its simplicity does have its place, but that if possible one should implement more precise scaling methods or volatility estimates for long time horizons calculated directly from long run returns.

In another publication which covers time scaling of volatility estimates, Gopikrishnan et al. (2000) analyze time scaling in financial time series in a general framework and find that for some asset classes such time scaling can be valid for horizons up to sixteen days, but that such scaling breaks down for longer time horizons.⁷⁹

An important point to consider when, or rather, before scaling volatility estimates, is which model is used to calculate the short-run volatility estimate⁸⁰ and if scaling would be a grave

⁷⁶ Reilly & Brown (2003), page 978

⁷⁷ Hull (2003), pp. 410-414

⁷⁸ Diebold et al. (1998), pp. 104-107

⁷⁹ Gopikrishnan et al. (2000), page 364

⁸⁰ i.e. which process or distribution the volatility itself or the underlying returns are assumed to follow.

violation of the applied model. Daniélsson (2002) states that applying the square root of time rule to a short-run volatility estimate generated from a conditional volatility model would imply a complete absence of proficiency in statistical concepts of risk estimation.⁸¹

Christoffersen and Diebold (2000)⁸² also take up this issue in their publication on volatility forecastability. They argue that the evaluations of volatility forecastability present in the literature are usually flawed in the sense that any results are not only dependent on the data and different time horizons but on the chosen or assumed volatility model. Therefore, one could not be certain that any results obtained by such an analysis are really due to the data, but rather could be due to the choice of model. In their own work the authors aim to assess volatility forecastability for different horizons without making model assumptions, thereby avoiding any potential distortion of their results caused by a potentially incorrect model choice.⁸³

In the first part of their analysis the authors formulate and test their model of a general assessment of volatility forecastability. Their general results are that the forecastability is very strong at short horizons and although it decreases with the horizon, the results for horizons up to twenty trading days are still reasonable. In the second part of their analysis, Christoffersen and Diebold (2000) apply their framework to financial market data, including foreign exchange, equity, and bond market return data for daily intervals. As far as equity and foreign exchange returns are concerned they find that for these markets volatility is significantly forecastable for horizons up to ten trading days with forecastability of up to three weeks in some markets. The bond markets are covered separately due to the fact that the underlying historical market data is usually in the form of annual yields as opposed to bond prices, so that exact returns cannot be calculated in a straightforward fashion. Therefore, in a first step an approximation is performed to obtain bond return data from bond yields. The calculated returns are then used to assess volatility forecastability. The results for bond returns are stronger than those for equity and foreign exchange, with forecastability present for horizons up to fifteen or twenty trading days, although the authors mention that some of this additional

⁸¹ Daniélsson (2002), page 1285

⁸² Christoffersen & Diebold (2000)

⁸³ For detailed information on the methods and analyses please refer to the original publication. (Christoffersen & Diebold, 2000)

forecastability could be due to the structural differences between bond markets and other capital markets.⁸⁴

Summarizing their findings, they state that when the relevant time horizons exceed ten or twenty trading days, depending on the markets, volatility forecasts are no longer reliable and therefore cannot be of grave significance for risk management applications due to a lack of precision. They also highlight in which respects their conclusions contradict certain research on the topic such as some of the propositions regarding the forecast performance of ARCH models contained in the RiskMetrics framework⁸⁵ or Jorion (1995)⁸⁶. On the other hand, they list different articles which seem to support their results. For example, they cite West and Cho (1995) who use a different approach in their analysis of foreign exchange market data but come to the similar conclusion that volatility in these markets is forecastable for horizons up to five trading days.⁸⁷ Finally Christoffersen and Diebold cite those studies of volatility which are based upon intraday (very high frequency) data such as five minute intervals, and at first glance would seem to support forecastability for very long time horizons.⁸⁸ They infer that when these time steps are in line with the data which is used to calculate the forecasts, i.e. five minute steps, then the resulting time horizon in days for five thousand five minute periods would be around seventeen days, which would also support their findings.

With regard to the topic at hand, time scaling of volatility estimates to long time horizons, these results are somewhat positive in the sense that they would support the calculation of a volatility estimate based upon daily market data for horizons exceeding a single day. On the other hand, this paper focuses on time horizons which are much longer than the five to twenty days which are propagated as the boundaries of such estimates, so that for long horizons such scaling would seem to fail.⁸⁹

⁸⁴ It is proposed that yield curve models (e.g. Brennan and Schwartz (1979) or Cox et al. (1985)) provide some explanations for the differences in volatility forecastability.

⁸⁵ Longerstaey (1996)

⁸⁶ Jorion (1995)

⁸⁷ West & Cho (1995)

⁸⁸ Up to 1000 or even 5000 time steps as presented in Andersen and Bollerslev (1997) for example.

⁸⁹ Christoffersen & Diebold (2000), pp. 12-22

In addition to the more general treatments of time scaling and forecastability of volatility estimates for longer time horizons which were just presented, there is also literature on time scaling of volatility estimates for various specific volatility models or processes.

Andersen and Bollerslev (1998)⁹⁰ do not treat time scaling of volatility estimates per se, but assess the quality of volatility estimates produced by ARCH models in general and GARCH(1,1) models as a specific example and argue that the forecast quality for daily volatility forecasts can be improved by using high frequency intraday data. They also liken this approach to work done by Schwert (1989, 1990)⁹¹ and Schwert and Seguin (1990)⁹² in which daily return data is used to calculate monthly volatilities. One of Andersen and Bollerslev's (1998) conclusions is that their approach could be extended to longer horizons. The authors do point out however, that when extending to longer horizons, correct modelling of long-term dependencies in volatility gains substantially in importance, and refer to research on this subject by Baillie et al. (1996).⁹³

In his articles on ARCH volatility models, Engle (2001 & 2004)⁹⁴ discusses issues to consider when extending a VaR horizon from one day to longer time horizons if the volatility estimate is calculated via an ARCH⁹⁵ model calibrated to daily return data. In his analysis he points out that when forecasting volatility many steps ahead with such a dynamic volatility model, eventually the long run volatility dominates the forecast and the day to day noise loses relevance. In Engle (2001) he describes a way to generate a long horizon volatility forecast when implementing a GARCH(1,1) model. The approach is straightforward: one generates sequential daily volatility forecasts step by step until the desired horizon is reached. Ultimately the volatility forecast will move closer and closer to the long run average volatility.⁹⁶ This is also one of the two main issues Engle highlights in his 2004 paper, as he

⁹⁰ Andersen & Bollerslev (1998), pp. 885-905

⁹¹ Schwert (1989), pp. 1115-1153; Schwert (1990), pp. 77-102

⁹² Schwert & Seguin (1990), pp. 1129-1155

⁹³ Baillie et al. (1996), pp. 3-30

⁹⁴ Engle (2001), pp.157-168; Engle (2004), pp. 405-419

⁹⁵ In his 2001 paper he focuses his investigations on GARCH(1,1) models which is one of the most common variations in financial applications, and in his 2004 paper he uses a TARCH model (=Threshold ARCH) as discussed in Zakoian (1994), which is an asymmetric volatility model.

⁹⁶ This is one of the three parameters in such a volatility model, next to the previous period's forecast and new information in the form of the most recent squared residual.

points out that when forecasting ahead for longer than a single day, the level of current volatility in relation to the long run average is relevant. If the current volatility is above a long run average it can be expected to decrease as the horizon increases and vice versa. The second important consideration, which he deems more relevant, is the effect of asymmetry in volatility estimates in a multi-period horizon.⁹⁷ He illustrates this phenomenon with a binomial tree, which is replicated in Figure 2.

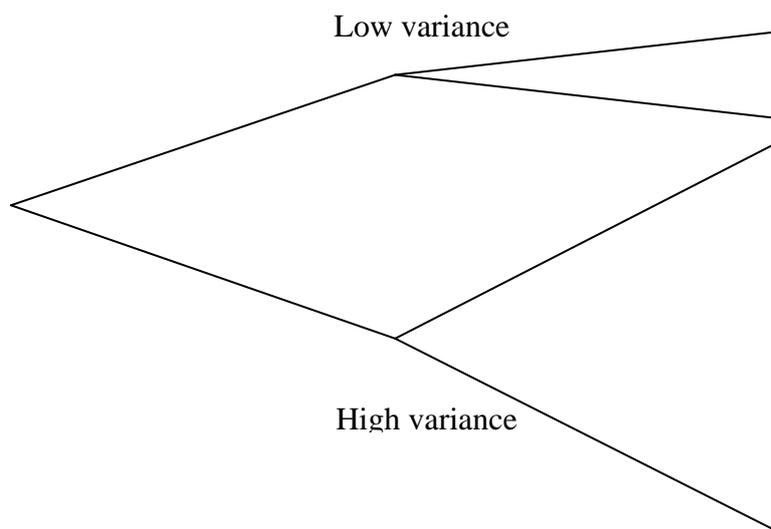


Figure 2 – Asymmetric Volatility – Source: Engle (2004), page 416

The probability of a positive return occurring is equal to that of a negative return in every period or day, i.e. the distributions are symmetric at every time step. The volatility in the next period however depends upon the direction of the previous move – if the return was negative the subsequent volatility is higher and if it was positive the volatility is lower. This relationship between return direction and volatility introduces asymmetry of volatility in a multi-period setting and would lead to a substantial increase of risk with respect to VaR, because this risk measure focuses on precisely those cases in which negative returns arise.

When implementing a long horizon or multi-period VaR that reflects this volatility asymmetry, simple time scaling is not possible. Engle proposes a simulation based approach which simulates the daily moves and corresponding changes in volatility up to the end of the VaR horizon, similar to the standard Monte Carlo simulation method.⁹⁸

⁹⁷ Engle refers to research by Nelson (1991) who finds that negative returns lead to higher subsequent volatility than equivalent positive returns

⁹⁸ Engle (2004), pp. 405-419

A different and frequently cited method of “scaling” volatility estimates calculated from GARCH models was proposed by Drost and Nijman (1993).⁹⁹ The authors analyse the relations between short horizon GARCH models based on high frequency, e.g. daily, data and longer horizon GARCH models based on lower frequency, e.g. weekly or monthly, data and whether or not one can obtain valid long horizon volatility estimates from a model calibrated to high frequency data. They argue that in applications of GARCH models it is usually assumed that a GARCH process at a short horizon is consistent with a GARCH process at a longer horizon, by making assumptions as to the independence of the innovations in the process. However, according to Drost and Nijman, these assumptions are not valid, so that simple scaling techniques as previously presented in this paper are not applicable to GARCH volatility estimates.

Their analyses are performed for GARCH models in general and are not limited to financial market data. However they do place a special emphasis on the GARCH(1,1) models and test their findings on exchange rate data so that the relevance and applicability of their findings for financial markets is illustrated. Basically Drost and Nijman develop means to derive the GARCH parameters for long horizon (e.g. monthly) volatility estimates from the parameters of a short horizon GARCH model which was calibrated with high frequency (for example daily) data.¹⁰⁰

Summarizing their results it can be stated that in general, generating weekly volatility estimates by using the proposed equations and daily data provides estimates that are very close to estimates calculated directly from weekly data and seem to be even better than the direct estimates. The results for generating monthly estimates from either daily or weekly data are similar, with the inferred volatility estimates seemingly outperforming the direct estimates based on monthly data.

These results would have to be deemed favourable for long horizon VaR models as the data availability of high frequency data is obviously greater than that of lower frequency data.

⁹⁹ Drost & Nijman (1993), pp. 909-927

¹⁰⁰ For the proofs and derivations of the authors’ model please refer to the original text.

In other research on this topic one can come across both support for¹⁰¹ and criticism of¹⁰² Drost and Nijman's model and findings. However those articles which are critical of the approach do not propose any alternatives. Therefore Drost and Nijman's model remains one of the few published means of generating long horizon GARCH volatility estimates from a short horizon GARCH model.

To conclude the discussions on time scaling or horizon extension of volatility forecasts, a very extensive comparative study on the subject will be reviewed. Embrechts and colleagues (2005)¹⁰³ evaluate various methodologies which can be applied to model the development of market risk factors for a one year horizon and test these on financial market data. Due to the fact that their article focuses on a time horizon of one year, which is exactly in the range that is the focus of this paper, and that they cover the main volatility models which have been taken into consideration in the treatment of volatility scaling that was just presented, their article serves as a very useful summary of this topic. Their analysis encompasses four different classes of processes which can be employed to model risk factors: random walks, autoregressive processes, GARCH(1,1) models and a static process originating from extreme value theory which replicates heavy-tailed distributions.

The authors treat each one of these model classes by describing or deriving an appropriate scaling rule to extend the time horizon from that of the applied data to the annual horizon and then providing the necessary equations to calculate both a VaR figure as well as an expected shortfall measure. After providing the means to obtain both scaled volatility estimates as well as VaR figures for long time horizons, the different models are tested for forecasting accuracy with regards to various classes of financial market data, including foreign exchange, equity indices, ten year government bonds and single stocks.

One of the interesting general issues that are brought forth by Embrechts et al. is that when trying to find the optimal data frequency for the calibration of any risk factor process, the tradeoff between the increase in sampling error due to smaller sample sizes¹⁰⁴ and the estimation improvement due to a decrease of data dependence with lower frequency data must be considered. Although this issue is also brought up in other publications such as Daniélsson

¹⁰¹ Diebold et al. (1998)

¹⁰² Christoffersen & Diebold (2000)

¹⁰³ Embrechts et al. (2005), pp. 61-90

¹⁰⁴ Jordan and Mackay (1997) present a bootstrapping technique which can be applied to mitigate this problem on page 268.

(2002)¹⁰⁵, the authors illuminate this issue further by pointing out that for random walk processes, lower frequency data are the better choice due to the normality assumption and that for GARCH(1,1) processes the opposite is true, as lower frequency data tend to violate necessary stationarity conditions so that these models should only be calibrated to data with horizons shorter than one month.¹⁰⁶

As far as the specific scaling methods for the random walk and GARCH processes are concerned, the authors proceed in line with the volatility scaling methods which have already been presented. For the random walk model the returns are adjusted for the mean return (i.e. “normalized”) and then the square root of time rule is applied to the short horizon volatility estimates to obtain annual volatility estimates. With regard to the GARCH(1,1) model, the Drost and Nijman approach, which was just discussed, is implemented.

The other two models have not yet been discussed in this paper, therefore the necessary steps for time scaling when applying these models will be explained briefly.

For autoregressive processes¹⁰⁷ returns are “detrended” initially (i.e. the trend is subtracted from the returns), similar to the procedure performed in the case of a random walk. The AR(p) process is then fitted or calibrated to the “detrended” return data and from this short horizon process scaled returns are estimated for the required time horizon and finally a long horizon volatility estimate is calculated from these scaled returns. The exact step by step procedures including the necessary equations can be found in the text of Embrechts and co-workers.¹⁰⁸

The fourth class of risk factor processes discussed is that of heavy-tailed distributions as derived from Extreme Value Theory (EVT).¹⁰⁹ A heavy-tailed distribution is defined by its tail index usually denoted with the greek letter alpha (α). This parameter is descriptive of the shape of the distribution’s tail. In their publication the authors use the “Hill estimator”¹¹⁰ to calculate α . In simplified terms, the “Hill estimator” describes the shape of the tail (those returns which fall below a certain threshold, for example below a 95% confidence level) by relating the average return in the tail to the return at exactly the threshold level.

¹⁰⁵ Danielsson (2002), page 1285

¹⁰⁶ Kaufmann & Patie (2003), pp. 1-61

¹⁰⁷ Kim et al. (1999), page 9

¹⁰⁸ Embrechts et al. (2005), page 65

¹⁰⁹ For detailed information on EVT refer to Beirlant et al. (1996), Embrechts (2000) and Embrechts et al. (2003).

¹¹⁰ Hill (1975), pp. 1163-1174

For financial market return data, alpha commonly takes on values between two and five.¹¹¹

The scaling approach presented for these heavy-tailed processes is applied directly to the VaR quantiles as opposed to the other models for which the volatility estimates are scaled.

The scaling factor is $k^{1/\alpha}$, with k being the number of periods to scale up to. This approach will be covered in more detail in the next section on alternative VaR quantile scaling methods. Therefore, it will not be further elaborated upon in this section of the thesis.

For the analysis itself and the evaluation regarding which of the four processes and in which form, as far as the data calibration is concerned, is optimal, Embrechts et al. (2005) use three different measures of forecasting precision. Two of these measures examine the quality of expected shortfall figures, which are not of direct interest as far as VaR is concerned. On the other hand one could confidently assume that a volatility forecast which leads to accurate expected shortfall estimates would also produce accurate VaR figures. In addition, the third criterion (the frequency of exceedances¹¹²) is related directly to VaR and is employed by the Basel committee on Banking Supervision when evaluating the quality of banks' VaR systems. The authors backtest the models by calculating these measures for numerous financial market data sets.¹¹³ As far as the calibrations go, the random walk and AP(p) process are calibrated to daily (1 day), weekly (5 days), monthly (22 days), quarterly (65 days) and annual (261 days) data, the GARCH(1,1) process to daily and weekly data, and the heavy-tailed distribution to daily, weekly, and monthly data.

¹¹¹ See Dacorogna et al. (2001) and Straetmans (1998) for empirical analyses.

¹¹² The number of times that the actual return in the backtesting period exceeded the forecast VaR figure is counted and should be as close as possible to the chosen confidence level - i.e. for a 95% VaR ideally there should be five exceedances in one hundred forecasts

¹¹³ Including time series for four exchange rates, five stock indices, government bonds issued by five countries and twenty two different single stocks which were elements of one of the analyzed indices.

The results across data and model, including the optimal calibration and over- or underestimation tendencies, if present, are summarized in the Table 1:¹¹⁴

	Exchange Rates	Bonds	Stock Indices	Single Stocks
Random walk	++++ 22 & 65 days minor overest.	++++ 65 days or less minor overest.	++++ 65 days or less minor overest.	+++ 261 days underest.
AR(p)	+++ 22 & 65 days underest.	+++ 22 days or less Underest.	++ 1 day underest.	+ 65 days underest.
GARCH(1,1)	+++ 5 days minor overest.	++ 1 day overest.	++++ 1 day minor underest.	++++ 5 days minor underest.
Heavy-tailed	++++ 22 days minor overest.	+ 22 days major underest.	+++ 22 days minor overest.	++ 22 days minor underest.

Table 1 – Source: Embrechts et al. (2005), pp. 80

The interpretation of this table should be fairly straightforward. The number of plusses (+) designates the quality of the estimates with four plusses being the best, the number of days listed are the optimal calibration of each model for the specific class of data and the under- and or overestimation tendencies are listed differentiated from minor to major.

As can be deduced from this table, the random walk model with a constant trend and normal innovations, calibrated on monthly or quarterly data is the best performing model for three of the four classes of data. The random walk also performs very well for the fourth class of data although in this case the optimal calibration is to annual data directly. The GARCH(1,1) model also performs very well on average and is the optimum model for the single stock and stock index data sets. Interestingly the heavy-tailed distribution model always yields the best results when calibrated to monthly data and produces very accurate forecasts for foreign exchange and stock index data. The AR(p) model generates above average estimates for exchange rate and bond data when scaled to monthly returns but shows a distinct tendency to underestimate risk across all analyzed asset classes.

¹¹⁴ The detailed results of all measurements for all models and calibrations can be found in the original text in tables 1 to 4 on pages 80 and 81.

In addition to presenting these results the authors perform further tests on the four models, including the testing of a simulated random walk¹¹⁵, replacing the normally distributed innovations by Student-t distributed innovations¹¹⁶, and variance analysis which focuses directly on the evaluation of the variance estimates and respective confidence intervals. The results of these additional tests and criteria increase the specificity of and reinforce the initial results. Surprisingly, rather than improving the results, Student-t distributed innovations in the random walk model decrease the quality of forecasts, so that the choice of normal innovations is affirmed. The variance analysis shows that volatility estimates generated by GARCH(1,1) models calibrated to daily or weekly data can fluctuate dramatically, leading to very wide confidence intervals and that those produced by AR(p) models are also unsatisfactory, as the confidence intervals are very asymmetric and in some cases do not even contain the “real” standard deviation of the underlying distribution. The estimates and confidence intervals for the random walk model, however, are superior and moderately stable. The simulation and variance tests also establish that the optimal calibration data for forecasts of a one year horizon are weekly or monthly returns.

These results might seem surprising as the most basic process which makes use of the simplest scaling rule seems to provide the best results when looking to estimate market risks for a horizon of one year. Two interesting issues are pointed out, however, on the one hand the fact that calibration to weekly or monthly data yields the best results and, on the other, that a good estimate of the trend of the returns must be found to achieve valid results.

This subsection presented research on the topic of time scaling of volatility estimates generated by different models or processes and the restrictions and problems associated with the various approaches. Summing up the main conclusions to be drawn from the surveyed literature, it can be noted that the key issues for achieving accurate long-run volatility estimates are choosing an appropriate scaling method for the chosen volatility or return model and calibrating the chosen model to return data at an intermediate horizon. Much of the criticism levied against scaling of volatilities by means of the square root of time rule originates from studies that scale daily volatility estimates to longer horizons and point to the

¹¹⁵ This aims to eliminate any distortions in the results that might stem from a miscalibration or misspecification of the random walk process by using a simulated process - i.e. it is correctly specified ex-definition.

¹¹⁶ Student-t distributions exhibit fatter tails than normal distributions.

large fluctuations of the estimates and the exaggeration of long-run volatility produced by this approach.¹¹⁷ As was suggested by Diebold et al. (1998) and then proven by Embrechts et al. (2005), this issue is mitigated by using data at weekly or monthly horizons as a basis for ensuing scaling procedures.

3.2.3. Other Scaling Approaches for VaR Quantiles and Models

After having devoted a significant amount of attention to the scaling of volatility estimates, as this parameter is the most important in any VaR model, the final subsection in this section on extensions of short horizon VaR models to long time horizons will present alternative VaR scaling approaches. One such approach was already touched upon in the discussion of Embrechts and colleagues' 2005 paper. These scaling methods can be applied to VaR figures or quantiles directly but are somewhat more complex or less restrictive than scaling quantiles by the basic square root of time rule.

A minor modification of the simple square root of time rule approach for the scaling of VaR quantiles was suggested by Al Janabi¹¹⁸. He proposes the following scaling relationship:

$$VaR_{T \text{ day}} = VaR_{1 \text{ day}} \times \sqrt{\frac{T+1}{2}}$$

This scaling factor is derived from the consideration that the minimum VaR horizon should be the period that it would take to unwind a position in an asset¹¹⁹ and its primary aim is to reflect the additional risk posed by illiquid market conditions. His argument supporting this factor is that in such conditions a position would be unwound by selling off equal parts linearly over the course of t-days until the entire position has been unwound.¹²⁰ Although his argument and scaling factor is tailored towards such liquidity risk, he extends the applicability of his scaling factor to “calculate the VAR for any time horizon.”¹²¹

¹¹⁷ Diebold et al. (1998)

¹¹⁸ Al Janabi (2007), pp. 41-42

¹¹⁹ Danielsson (2002), page 1285

¹²⁰ In their article on Optimal Execution of Portfolio Transactions, Almgren and Chriss (2001) take liquidity risk and its potential effect on VaR during the unwinding of a position into account and provide a modified VaR model taking such liquidity risk into account (L-VaR).

¹²¹ Al Janabi (2007), page 42

This proposed scaling factor would generate substantially lower VaR figures than those produced by the standard square root of time rule. However, as no additional arguments in support of this approach are given, an implementation should be considered with care.

Another alternative scaling method for VaR quantiles was already mentioned briefly in the previous subsection when the VaR scaling approach for heavy-tailed distributions as presented in Embrechts et al. (2005) was discussed.

The so called “alpha-root rule” originates from Extreme Value Theory, where alpha reflects the tail index of a distribution. This scaling technique is explained in detail in a publication by Daniélsson and de Vries (2000).

In their article, they present a VaR model which combines historical simulation with a modelling of the distributional tails in line with EVT. The authors argue that independence conditions which are present when returns are assumed to follow a normal distribution are in fact present in tail distributions. They point to empirical evidence which shows that the occurrence of one “tail event”, i.e. an extreme negative return, is not indicative of a second sequential tail event, so that the independence assumption holds when focussing on the tails of return distribution. However, the probability of an extreme return occurring is significantly higher than in normal distributions. As far as the time scaling of VaR quantiles for their model is concerned, the authors point to work on heavy-tailed distributions by Feller (1971). The scaling factor derived from Feller’s analysis is the following:¹²²

$$VaR_{T \text{ day}} = VaR_{1 \text{ day}} \times T^{\frac{1}{\alpha}} = VaR_{1 \text{ day}} \times \sqrt[\alpha]{T}$$

For alpha equal to two this scaling factor is exactly the standard square root of time rule which stems from normal distributions of returns. For values of alpha below two the volatility is infinite.¹²³ However, for financial time series, alpha is generally above two, in line with the assumption of a heavy-tailed distribution. Daniélsson and de Vries state that for the distributional tails of these heavy-tailed distributions, the self-additivity property, which is the basis for applying the square root of time rule on normal distributions, is also present. This is their basis for the alpha-root rule.

¹²² Derived from Daniélsson and de Vries (2000).

¹²³ Fama & Miller (1972), page 270

The scaling factor is related to the shape of the distributional tail, represented by the tail index. Due to the fact that this approach is applied for values of alpha which are greater than two, the scaling factor is smaller than the square root of time rule.

This fact seems contradictory to statements that normal distributions underestimate the VaR in relation to heavy-tailed distributions. However, as the authors explain, the probability of tail events is higher in heavy-tailed distributions leading to higher one period VaR figures. This effect balances the effect of a smaller multiplier when VaR figures are scaled to longer time horizons. For time scaling from a one day VaR to a ten day VaR this issue can be illustrated by the following Table 2:

VaR confidence level		95%	99%	99.5%	99.9%	99.95%	99.995%
heavy-tail VaR	1 day	\$ 0.9	\$ 1.5	\$ 1.7	\$ 2.5	\$3.0	\$ 5.1
	10 day (alpha root)	\$ 1.6	\$ 2.5	\$ 3.0	\$ 4.3	\$ 5.1	\$ 8.9
“normal” VaR	1 day	\$ 1.0	\$ 1.4	\$ 1.6	\$ 1.9	\$ 2.0	\$ 2.3
	10 day (square root)	\$ 3.2	\$ 4.5	\$ 4.9	\$ 5.9	\$ 6.3	\$ 7.5

Table 2 – Source: Daniélsson & de Vries (2000)

This table shows VaR figures (in millions of U.S. Dollars) for an analysis that Daniélsson and de Vries performed on a portfolio with one hundred million U.S. Dollar portfolio value.

The results are in line with what could be expected. For lower confidence levels (95% and 99%) “normal” VaR is greater than the heavy-tail VaR at both one and ten day horizons. With increasing confidence levels the normal VaR increasingly underestimates tail risk and the VaR figures fall below heavy-tail VaR, first at the one day horizon (99.5%) and eventually also for the ten day horizon (99.995%).

In their tests of the approach on empirical equity data the authors calculated an average alpha of 4.6 which leads to a scaling factor of 1.7 for ten days as opposed to the square root of ten which is 3.7.

This approach is an appropriate alternative to the standard square root of time rule when assuming a heavy-tailed distribution. Although Daniélsson and de Vries only analyse scaling of daily values to horizons of ten days, the alpha-root rule can produce good results for long time horizons when calibrated to monthly data, as was previously mentioned.¹²⁴

¹²⁴ Embrechts et al. (2005)

A somewhat similar and yet significantly different VaR quantile scaling methodology is presented in two papers by McNeil (1999 & 2000), one of which is co-authored by Frey (2000).¹²⁵ They conceive a short horizon VaR model which is also based on Extreme Value Theory to reflect a heavy-tailed return distribution. However, contrary to Danielsson and de Vries (2000) who assume a historical distribution for their volatility estimate and apply a Hill estimator to model the distributional tails, McNeil and Frey (2000) calibrate the historical return data to a GARCH(1,1) model and then go on to apply a different technique to estimate the distributional tails. The reasons given for choosing a GARCH model to generate the volatility estimate, which is then used to estimate the tail distribution, are in line with standard arguments in support of GARCH models - i.e. they more accurately reflect the tendencies of volatility clustering seen in financial market return data. For the tail distribution itself they select a Generalized Pareto Distribution (GPD)¹²⁶ as the distribution function, which is the alternative model aside from the Hill estimator within the Peaks-Over-Threshold (POT) group of models for implementations of heavy-tailed distributions in line with Extreme Value Theory. The authors argue that a GPD approach for heavy-tail modelling has advantages over applying the Hill estimator, such as the fact that it produces results which are more stable in the face of different data sets and that it is more flexible, as it can be applied to heavy-tailed distributions as well as to distributions which do not exhibit heavy tails, whereas the Hill estimator is explicitly tailored towards heavy-tailed distributions. They support these arguments with backtesting results for VaR quantiles at a confidence level of 99% or above. Based on this short horizon VaR model they go on to propose a Monte Carlo simulation approach to produce scaled VaR forecasts for longer time horizons. In this simulation they use their short horizon model to generate a distribution of residuals, go on to draw from this distribution to compute sequential return estimates for each day up to the selected horizon and then use the same technique, which was originally applied to the underlying daily return data¹²⁷, to obtain the long horizon VaR quantiles.

¹²⁵ McNeil (1999), pp. 93-113; McNeil & Frey (2000), pp. 271-300; It should be noted that although McNeil (1999) was published before McNeil and Frey (2000), the work presented in the 1999 publication draws from a preprint version of the 2000 publication.

¹²⁶ For details on this distribution function and the authors' estimation procedure, please refer to the original text and for details on GPD in general, see Embrechts et al. (1998), pp. 96-100.

¹²⁷ I.e. calibrate a GARCH model to the generated distribution and estimate the tail distribution via GPD. See McNeil & Frey (2000), pp. 296-297 for details.

To evaluate the performance of their scaling methodology, McNeil and Frey (2000) compare the accuracy of the approach with that of applying the square root of time rule to the initial one day VaR figures produced by their model. Backtesting is performed for five different empirical data sets, including returns for two equity indices, one single stock, a foreign exchange rate, and a commodity price. The forecast quality with regard to these five time series is assessed for scaling horizons of five and ten days and confidence levels of 95% and 99%. The results of their comparisons are that the simulation approach outperforms the square root of time method for the majority of the tested time series. However, when taking a closer look at the detailed backtesting results, one cannot completely agree with the dominance which McNeil and Frey (2000) attest to their model. For one of the equity indices the square root of time rule underestimates VaR drastically, whereas the simulated values only produce minor underestimations. As far as the other data sets are concerned, the results are not so one-sided. For the foreign exchange and commodity data the forecast quality is similar with relative advantages amongst the two approaches depending upon the chosen confidence level and scaling horizon. For the single stock and the second equity index, the simulation method does outperform the square root of time rule noticeably, as this scaling technique significantly underestimates VaR for this asset class.

Based on the simulated, scaled VaR quantiles, the authors derive so-called “implied scaling factors” (λ - lambda) for different volatility regimes¹²⁸ and confidence levels, which can be applied in a VaR scaling relationship for time horizons up to fifty days:¹²⁹

$$VaR_{T \text{ day}} \approx VaR_{1 \text{ day}} \times T^\lambda$$

The implied scaling factors are presented in the Table 3 below:

		Confidence level	
		95%	99%
Volatility regime	Low volatility	0.65	0.65
	Average volatility	0.60	0.59
	High volatility	0.48	0.47

Table 3 – Source: McNeil (1999), page 104

¹²⁸ In support of a differentiation of VaR scaling dependent on current volatility levels, McNeil (1999) cites a publication by Hull and White (1998).

¹²⁹ Based upon McNeil (1999), page 104.

Relating this lambda scaling relationship to the square root of time rule and the previously discussed alpha-root rule leads to the following conversions

for the square root of time rule:

$$\alpha = 2 \rightarrow \lambda = 0.5$$

for the alpha root rule in general:

$$\lambda = \frac{1}{\alpha}$$

and per definition:

$$\alpha > 2 \rightarrow \lambda < 0.5 .$$

Interpreting these results the authors propose that if one were to apply a simple time scaling factor, the exponent should be greater than one half, which is the exponent used in the square root of time rule, except for cases in which volatility is high. They go on to criticize the alpha-root rule¹³⁰, which typically produces even smaller scaling factors.

Putting these results and the proposed VaR scaling model in general into perspective with respect to long time horizons, it has to be noted that extension of the methodology to time horizons which are significantly longer than the proposed ten or fifty days could be problematic. It would seem that some of the elements of the approach, such as use of GARCH models, are strongly dependent on calibration to daily data, and applying the model to very long time horizons (i.e. well beyond fifty days) with data calibration to weekly or monthly returns, could turn out to be quite laborious if not superfluous, due to the fact that for such time horizons, a GARCH volatility forecast is dominated by the long-run average as opposed to current volatility levels. Additionally, the quality of results would need to be validated separately for such scenarios.

The last publication which will be covered in this subsection is also the most recent. Brummelhuis and Kaufmann (2007)¹³¹ analyse and compare different approaches for the time scaling of VaR quantiles when the return distribution is assumed to follow either a random walk, GARCH(1,1), or AR(1)-GARCH(1,1) process. They perform their analysis for three different error or innovation distributions to differentiate between the standard normal distribution and heavier-tailed distributions, represented by two different Student-t distributions. For these nine different combinations of return processes and error distributions

¹³⁰ Daníelsson & de Vries (2000)

¹³¹ Brummehuis & Kaufmann (2007), pp. 39-94

they initially compare the performance of seven different scaling approaches for one day, 99% VaR quantiles to a horizon of ten days when the data set available for calibration of the selected model is limited to 250 daily log returns. The choices of the 99% confidence level and the ten day horizons are motivated by the regulations in the Basle Capital Accord.¹³²

The seven included scaling methods provide an overview of some of the approaches which were previously discussed and others which are interesting alternatives. The benchmark is the standard square root of time rule. Two of the approaches do not really apply scaling techniques, but are simple procedures which highlight the problems associated with a limited data basis in long horizon VaR. One is simply calculating non-overlapping ten day returns from the 250 data points, which yields a sample of only 25 ten-day returns on which to base the VaR calculations. This approach is obviously an extreme example of limited underlying return data. The second of these approaches is calculating 241 overlapping ten-day returns from the 250 days of data, which in the literature is also discarded as inherently flawed as it introduces serious problems related to serial correlation in the derived return data. The next three approaches are variants of the bootstrapping technique previously mentioned¹³³ which provides a means of mitigating the issues related to limited datasets. This technique, similar to a historical simulation approach, generates a distribution of ten-day returns by sequentially selecting ten daily returns from the original data set and simply summing up these ten log returns to compute one data point in a ten-day return distribution. This procedure is repeated until 10000 ten-day returns are generated, which are then used as the distributional basis for VaR calculations. The three variants of this approach are created by variations in the selection procedure for the ten drawings for each of the ten-day returns. The first procedure of “Random resampling”¹³⁴ is based on combination with repetition.¹³⁵ The second selection rule is defined to ensure independence amongst the ten daily returns chosen to generate a single ten-day return. This is achieved by ensuring that the time span between each of the drawn daily returns is at least ten days. The third and final variation is set up in exactly the opposite manner. It aims to ensure dependence amongst the ten selected daily returns by selecting the ten returns out of 231 successive subsets of twenty consecutive daily values. This procedure is repeated 44 times to produce 10000 data points for dependent ten-day returns.

¹³² BIS (2005), page 40

¹³³ Jordan & Mackay (1997), page 268

¹³⁴ Brummelhuis & Kaufmann (2007), page 48

¹³⁵ i.e. Randomly and with repetition - every one of the 250 daily returns can be selected in every drawing.

The final scaling method is very similar to the EVT approaches implemented by McNeil and Frey (2000)¹³⁶ and Danielsson and de Vries (2000)¹³⁷ in that the tails (the top ten and bottom ten percent) of the assumed daily return distribution are modelled to a GPD process and the middle part of the distribution is derived from the 250 daily returns. From this distribution of daily returns 10000 simulation runs of ten daily returns per run are calculated to generate a ten-day return distribution.

For their performance evaluations of the different scaling techniques Brummelhuis and Kaufmann (2007) first calculate a “true” ten-day Value at Risk for each of the nine process and error term combinations. They go on to calculate six different performance measures for each possible scaling procedure applied to every one of these nine combinations. Their prime performance criterion is a comparative measure relating the mean deviation of the scaled VaR figures from the true ten-day VaR to the mean deviation resulting from the ten-day VaR calculated via the square root of time rule, which serves as the benchmark. In total for each process and error term combination, six relative performance assessments are completed.¹³⁸

The results of this, in effect empirical analysis, are as follows. The two approaches which apply a calculation of actual ten-day returns (overlapping and non-overlapping) from the underlying dataset distinctly underperform the square root of time rule for all process and error term combinations. It can be observed, however, that the relative performance of these methods improves for the heavier tailed error term distributions. These results are in line with the problems associated with these two approaches – serial correlation and insufficient data set, respectively. The third method, which performs poorly in relation to the square root of time rule, although it does outperform the first two approaches, is the bootstrapping based method with dependent sampling. Unfortunately the authors do not derive or provide an explanation for this underperformance, although it would seem counterintuitive.

The results of the performance analyses for the remaining three scaling approaches are summarized in the following Figure 3¹³⁹ which shows the mean relative deviations from the true VaR figures.

¹³⁶ McNeil & Frey (2000), pp. 271-300

¹³⁷ Danielsson & de Vries (2000), pp. 236-269

¹³⁸ i.e. everyone of the six alternative scaling procedures is benchmarked against the square-root-of-time rule.

¹³⁹ Brummelhuis & Kaufmann (2007), page 62

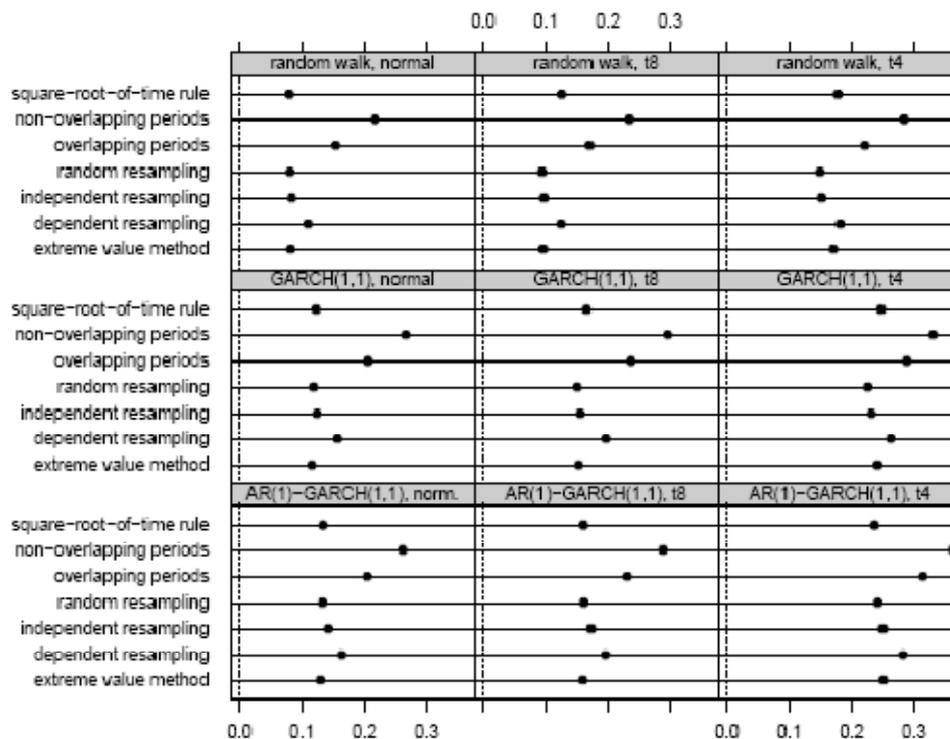


Figure 3 – Source: Brummelhuis & Kaufmann (2007), page 62

Basically the three approaches, random resampling, independent resampling, and the EVT method achieve very similar results across processes and innovation distributions. The only significant differences are present for Student- t_4 distributed innovations, whereby the EVT method underperforms the other two scaling techniques. The best results are achieved for the combinations of the random walk process with the Student- t distributed errors.

As far as the relative performance of these approaches in relation to the benchmark, the square root of time rule, is concerned, the performance is equal to the benchmark for all processes when implementing normally distributed errors, and in the case of the AR(1)-GARCH(1,1) process also for the Student- t distributed errors. On the other hand, the three scaling approaches significantly outperform the square root of time rule for Student- t distributed errors in combination with a random walk process, and also show measurable performance advantages for the GARCH(1,1) process. These advantages stem from the fact that the square root of time rule overestimates ten day VaR for heavy-tailed distributions.

In their presentation of the performance results, Brummelhuis and Kaufmann (2007)¹⁴⁰ go on to interpret the additional information contained in the other performance measures which they applied. These show that the alternative scaling approaches display a tendency to

¹⁴⁰ Brummelhuis & Kaufmann (2007), pp. 61-65

underestimate VaR, whereas the square root of time rule tends to overestimate VaR. As previously mentioned, this tendency, inherent in the square root of time rule, is amplified for the random walk process with Student-t innovations. For the AR(1)-GARCH(1,1) process all scaling rules show a tendency of underestimation independent of the chosen error distribution, with the square root of time rule exhibiting the smallest bias.

Fundamentally, the other performance measures support the initial assessment of relative performance, i.e. the square root of time rule performs very well in relation to all other scaling approaches, providing results equal in accuracy to those achieved by the random resampling, independent resampling, and EVT approaches in most cases, only being outperformed in some instances with heavy-tailed innovations. As far as these three alternative scaling approaches are concerned, the authors argue that the random resampling approach has advantages over the other two, because it attains the same level of precision whilst exhibiting significantly simpler implementation properties.

When evaluating these results in general and comparing them to those presented by McNeil and Frey (2000)¹⁴¹, one should keep in mind that they are related specifically to a ten day horizon and a 99% confidence level, and that Brummelhuis and Kaufmann (2007) analyse unconditional VaR estimates as opposed to McNeil and Frey who deal with conditional VaR estimates.¹⁴²

In addition to their research on time scaling of 99% VaR quantiles to ten day horizons, which is the main focus of their paper, the authors provide some insights into a few general issues which are of great relevance for long-run VaR calculations. As their main analysis assumes a zero trend in the underlying return data, one of the issues they investigate is how a non-zero trend could be properly integrated into the VaR calculation and scaling procedure. The suggested procedure is very similar to the one proposed by Embrechts, Kaufmann and Patie (2005)¹⁴³.

First, a one day trend is derived by calculating the mean daily return from the dataset of daily returns. Then the original daily return data is “de-trended” (normalized or centered) by subtracting the one day trend from the return data. This modified dataset serves as the basis for a ten day VaR calculation and in a final step one simply subtracts the ten day trend¹⁴⁴ from

¹⁴¹ McNeil & Frey (2000), pp. 271-300

¹⁴² The authors also point to additional research on conditional VaR estimates by Brummelhuis and Guégan (2005).

¹⁴³ Embrechts et al. (2005), pp. 61-90

¹⁴⁴ The ten day trend is simply the one day trend multiplied by ten.

the VaR figure to produce an adjusted ten day VaR figure which accurately reflects a non-zero trend in the underlying return data. The second issue discussed is the potential presence of auto-regression in the underlying return data, which is relevant in AR(1)-GARCH(1,1) processes.

They argue that for a stationary AR(1) daily log-return process (X_t) in the form:¹⁴⁵

$$X_t = \lambda \times X_{t-1} + \epsilon_t \quad \text{with } |\lambda| < 1 \text{ and } \epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{1}) \quad \text{i. i. d.}$$

one can calculate a ten day VaR with the following formula:

$$VaR_{10 \text{ day}, 99\%} = \frac{1}{1-\lambda} \times \sqrt{\mathbf{10} - 2\lambda \frac{1-\lambda^{10}}{1-\lambda^2}} \times q_{99\%}^{\mathcal{N}}.$$

The ten day square-root-of-time VaR estimate is given by:

$$\widehat{VaR}_{10 \text{ day}, 99\%} = \sqrt{\frac{\mathbf{10}}{1-\lambda^2}} \times q_{99\%}^{\mathcal{N}}$$

where $q_{99\%}^{\mathcal{N}}$ represents the 99% quantile of the standard normal distribution.

By relating the two latter equations for calculating a ten day VaR to one another, the authors derive a quotient for the relative underestimation of true ten day VaR by the scaled VaR given by:

$$\frac{VaR_{10 \text{ day}, 99\%} - \widehat{VaR}_{10 \text{ day}, 99\%}}{VaR_{10 \text{ day}, 99\%}} > 0.68 \times \lambda.$$

For large values of lambda (λ) this could lead to significant underestimations of VaR if the square root of time rule is applied. However, the authors cite empirical results for returns of different financial asset classes which on average generated lambda values around 0.04 so that any underestimation tendencies would seem to be rather limited.

The final generalisation presented by Brummelhuis and Kaufmann (2007)¹⁴⁶ is also the most relevant for this paper. They derive a fairly simple formula to scale one day VaR quantiles to T-day VaR quantiles with a fixed confidence level, for AR(1)-GARCH(1,1) processes, which they argue outperforms the square root of time rule.

Building on the AR(1) process (X_t) which was just presented, the AR(1)-GARCH(1,1) process is defined as follows.

$$\sigma_t^2 = \mathbf{a}_0 + \mathbf{a} \times (X_{t-1} - \lambda X_{t-2})^2 + \mathbf{b} \times \sigma_{t-1}^2$$

¹⁴⁵ All relevant equations are drawn from page 71 in Brummelhuis and Kaufmann (2007) and slightly modified.

¹⁴⁶ Brummelhuis & Kaufmann (2007), pp. 74-78

This process reduces to a standard GARCH(1,1) process for lambda equal to zero, so that the scaling formula can easily be applied to such a process.

The authors' derivations yield the following result.

$$\frac{VaR_{\alpha\%,T \text{ day}}}{\sigma_{(T)}} \rightarrow q_{\alpha\%}^{\mathcal{N}} (T \rightarrow \infty)$$

So that for large values of T, i.e. very long time horizons, the scaled VaR quantile can be defined as

$$VaR_{\alpha\%,T \text{ day}} \approx \sigma_{(T)} \times q_{\alpha\%}^{\mathcal{N}}$$

with

$$\sigma_{(T)}^2 = \frac{\sigma_{\infty}^2}{(1-\lambda)^2} \times \left(T - 2\lambda \times \frac{1-\lambda^T}{1-\lambda^2} \right) \quad \text{and} \quad \sigma_{\infty}^2 = \frac{a_0}{1-a-b}.$$

This provides a closed form approach to scale unconditional VaR quantiles to long time horizons, which uses quantiles of the standard normal distribution. An analogous formula for conditional VaR quantiles is also presented in the original text.

Summing up the most important findings contained in Brummelhuis and Kaufmann's (2007) extensive publication brings forth evidence that the square root of time rule for scaling VaR quantiles performs very well for various combinations of underlying return processes and assumed innovation distributions. For heavy-tailed distributions and GARCH(1,1) or AR(1)-GARCH(1,1) processes some alternative scaling models tend to outperform the square root of time rule, but when it is outperformed, this simple scaling approach usually errs on the "safe" side in the sense that it tends to overestimate VaR.

Finally the authors derived general scaling formulas for AR(1)-GARCH(1,1) processes with which both unconditional and conditional VaR quantiles can be scaled to very long time horizons, although unfortunately the authors did not provide backtesting results for horizons longer than ten days.

This concludes the section on simple methods of extending a short horizon VaR model to long time horizons in the form of various scaling approaches applied either directly to VaR quantiles or to volatility estimates which serve as the pivotal input parameter in any VaR calculation.

The simplest scaling approach in the form of the square root of time rule, applied to volatility estimates as well as quantiles, was compared to multiple alternative scaling techniques applicable to quantiles or volatility estimates for different possible return processes, such as

random walks, GARCH processes (including or excluding an autoregressive component), models based on EVT, and others. Research on the effects of implementing an assortment of possible innovation or error distributions within these different processes was also presented. Somewhat surprisingly, when taking all of the presented research into account, the square root of time rule must be deemed an acceptable and fairly accurate scaling approach for the long horizons in the focus of this thesis. The approach is outperformed in some cases by alternative approaches depending on the horizon, confidence level, and foremost the underlying return process; however, across all possible combinations of processes and different asset classes, the approach performs very well for very long time horizons. Where possible it seems that scaling volatility estimates rather than VaR quantiles directly should be preferred as this procedure is less restrictive. And finally, the limited research on very long time horizons, as opposed to that on horizons up to a few weeks, points out that when possible, the return process should be calibrated to intermediate returns such as weekly or monthly if the available time series history is sufficiently long, and that trends in the return data should be accounted for.

This section on simple extensions of short-run models to long time horizons was fairly elaborate. The reasons being that, on the one hand, if such a simple extension can provide acceptable results it has vast advantages in any implementation due to the fact that it can simply be applied on top of any short-run VaR systems already in place, and does not require the implementation of a distinct and entirely new model and system for calculation of long-run VaR figures. Additionally, the amount of research performed and published on scaling approaches of VaR and volatility estimates is much more extensive than the publications on specific long-run VaR models which are the subject of the next section. Lastly, the subsection on the time scaling of volatilities is a very important element in the discussions on long-run VaR in general and will be very useful in the ensuing treatment of explicit VaR models, as every long-run VaR model needs a forecast for long-run volatility.

3.3. Specific Long-Run Value at Risk Models

Although the literature on different aspects of Value at Risk has grown to become quite extensive throughout the years, there are few publications specifically dealing with Value at Risk for long time horizons. After an extensive literature search there were a few publications to be found which either present long-run VaR models explicitly or as part of the discussion of a related issue.

In a first part of the treatment of long-run Value at Risk models, the approaches and models presented in these publications will be surveyed and critically evaluated, keeping in mind the critical issues and assumptions highlighted in the previous part of the thesis.

3.3.1. “A tale of two VaRs”¹⁴⁷ - Which VaR should be considered?

One of these critical issues, that of a zero mean return for assets or risk factors, is of distinct importance in computing a long-run VaR, because there are two different methodical approaches to calculating VaR. These are equivalent in a short-run or zero mean return environment, but very different in a long-run context. The two approaches, which are discussed in Kupiec (1999)¹⁴⁸ differ with respect to which asset or portfolio value the VaR-quantile of the portfolio value at the end of the VaR horizon is compared to when calculating the VaR. One can either compare this “worst case” portfolio value to the present value of the portfolio at the time of evaluation, or to the expected mean value of the portfolio at the end of the VaR horizon.

Expressed in equations the two different measures are:¹⁴⁹

$$\begin{aligned} VaR(c) &= | \min [P_0(\bar{r} + \alpha_c \sigma), 0] | \\ VaR^{\bar{r}}(c) &= | P_0(\alpha_c \sigma) | \end{aligned}$$

In short-run horizons and when a mean return of zero is assumed, these two different Value at Risk metrics are identical. When extending to long time horizons and assuming a positive mean return, the two different approaches lead to very different results, with the difference increasing with both the time horizon and the assumed mean return or drift.

¹⁴⁷ Kupiec (1999), page 42

¹⁴⁸ Kupiec (1999), pp. 41-52

¹⁴⁹ Kupiec (1999), page 42 with modifications.

Obviously the expected mean value of a portfolio increases with time when a positive mean return is assumed, whereas the present value is set at the time of the evaluation. The distribution of portfolio values at the end of the VaR horizon depends on various distributional assumptions and parameters, but in the simple log-normal case for given positive mean return or drift, it depends upon the volatility of the returns.

For this case the detailed equations presented are as follows:

$$VaR(c) = \left| \min \left(P_0 \left[e^{\left(\mu - \frac{\sigma^2}{2} \right) T + \alpha_c \sigma \sqrt{T}} - 1 \right], 0 \right) \right|$$

$$VaR^{\bar{}}(c) = \left| \min \left(P_0 \left[e^{\left(\mu - \frac{\sigma^2}{2} \right) T + \alpha_c \sigma \sqrt{T}} - e^{\left(\mu - \frac{\sigma^2}{2} \right) T} \right], 0 \right) \right|$$

Summarizing arguments brought forth by Kupiec¹⁵⁰, both approaches can produce problematic results. When calculating VaR in relation to the initial portfolio value, the positive mean return can, with increasing time horizon, dominate the risk element and lead to a negative VaR figure, whereby the VaR “worst case loss” is in fact a profit.¹⁵¹

On the other hand, when VaR is calculated vis-à-vis the mean expected portfolio value at the end of the horizon, the VaR increases with time just as the mean terminal value does and the resulting VaR figure can exceed the initial portfolio value. Such a result is not appealing as it would imply that the loss can exceed the amount which was invested, a scenario which is “counter-intuitive (if not absurd)”¹⁵² as Kupiec puts it and not possible in a limited liability context.¹⁵³

The difference between these two methodical approaches to VaR must be kept in mind when looking into long-run VaR models.

¹⁵⁰ Kupiec (1999), pp. 42 - 44

¹⁵¹ i.e. the VaR quantile of the terminal portfolio Value distribution is greater than the present value of the portfolio.

¹⁵² Kupiec (1999), page 43

¹⁵³ There are of course portfolio elements such as short positions in certain assets which could in fact lead to such unbounded losses, but not for “real assets” held in long positions, e.g. holding bonds or equity.

3.3.2. MaxVaR¹⁵⁴

The first of the specific long-run VaR models covered is titled “MaxVaR” and focuses on the effects of interim risk in a mark to market environment. This source of risk was already discussed in the comments on the assumption of constant portfolio composition in the previous part of the thesis. The framework Boudoukh et al. (2004) develop aims to quantify this interim risk with so called MaxVaR figures which represent “worst case” loss estimates on or before the end of the VaR horizon.

MaxVaR therefore “considers the probability of seeing a given low cumulative return on or before the terminal date”¹⁵⁵ as opposed to VaR, which only considers the probability of seeing a given low (cumulative) return on the terminal date. Interesting by-products of their results are adjustment factors for regular VaR calculations, subject to different specific calculation parameters, to come up with an approximation for MaxVaR based upon a certain VaR figure.

The model assumes that the asset or portfolio value follow a log-normal process, just as the standard VaR models do. It does however consider the possibility and effects of a non-zero positive drift parameter. The fundamental difference between MaxVaR and VaR is that, with P_t being the relevant asset or portfolio value at time t ¹⁵⁶, VaR only takes the distribution of P_T , specifically, a certain quantile of this distribution, into account, whereas MaxVaR focuses on the worst case path P_t could follow between 0 and T. Basically, instead of only looking at the “lowest level” of S at the end of the VaR horizon, MaxVaR aims to identify the lowest value S reaches at any point throughout the entire VaR horizon. To formulate their model in terms of an equation or theorem, Boudoukh et al. define a minimum cumulative continuously compounded return over the VaR horizon as the return between the minimum value S_t takes over the entire VaR horizon and S_0 . Their theorem in turn defines the probability of this minimum cumulative return lying below a certain threshold level, which leads to the result of identifying a MaxVaR figure for a given threshold and confidence level. For details on the derivation and proofs of their model and theorem please see the authors’ original article.

The authors then calculate MaxVaR and associated VaR figures in terms of multiples of time and the standard deviation of asset or portfolio value, for different confidence levels and

¹⁵⁴ Boudoukh et al. (2004), pp. 14-19

¹⁵⁵ Boudoukh et al. (2004), page 14

¹⁵⁶ t ranging from 0, the time of evaluation, to T, the end of the VaR horizon.

different levels of expected mean return or drift. To summarize their findings they present a quotient of MaxVaR to VaR for every considered parameter combination and show that for a zero expected return MaxVaR can exceed VaR by between approximately eleven and nineteen percent depending on the confidence level.¹⁵⁷

Interestingly, for zero expected return these results are independent of both the horizon and volatility. For the case of zero expected return the authors also present MaxVaR figures for discrete interim sampling¹⁵⁸ and show that this reduces MaxVaR inversely related to the number of samples taken, “regular” VaR being at one end of the spectrum with a single sample taken and their general MaxVaR taking infinite samples at the other end of the spectrum. Their explanation for this relationship is that the discrete sampling can lead to “missing” the lowest asset value.

The results for positive expected returns are the following. Although both VaR and MaxVaR decrease in absolute terms due to the positive drift, the relationship between VaR and MaxVaR changes substantially. Fundamentally, the relationship is no longer independent of the horizon and volatility, due to the fact that the positive drift has a greater effect on the distribution of P_T than on any P_t before the end of the VaR horizon (T). In general terms: the effects of interim risk, as defined by the authors, increase with drift and time but decrease with volatility. For the parameters chosen they show that depending on the chosen confidence level, analogous to the zero expected return case, MaxVaR can exceed VaR by between nineteen and seventy-five percent.¹⁵⁹ These differences are without a doubt substantial and for certain assets and long time horizons, the authors look at one year, they could exceed even these high levels.

Summarizing this approach in terms of the focal points of this thesis, it can be stated that the proposed model deviates from the standard VaR models in terms of two elements. On the one hand it includes the possibility of integrating a non-zero mean expected return for a VaR calculation, which is a very important point for long time horizons. The other element, which is the essence of the MaxVaR concept, is that the authors quantify the probability of reaching certain threshold value levels at any time during the VaR horizon and thereby quantify “interim risk” or, as previously mentioned, certain elements of intra-horizon events or risk.

¹⁵⁷ The factor decreases as the confidence level increases, 19% for 95% confidence and 11% for 99% confidence.

¹⁵⁸ As opposed to the continuous approach generally followed in their framework.

¹⁵⁹ Boudoukh et al. (2004), page 18

This extension of standard VaR models is valuable in long-run scenarios and can be easily applied to improve such standard models; however, due to the simplicity of the approach it seems to be quite limited in its application. It could, for example, give the user an indication of the probability that a certain asset or risk factor value would be breached within the considered VaR horizon, so that one could get an idea of how probable the risk of a margin call is, but the model does not provide means to integrate potential effects of such an event directly into the model. Such an event would, among other things, cause portfolio composition to change and this effect could not be directly reflected as MaxVaR conforms to the standard assumption of constant portfolio composition.

The authors do discuss some of the standard assumptions¹⁶⁰ they apply, deeming them to be markedly problematic. However these discussions merely cite some of the common arguments against these assumptions but leave potential solutions up to future research. As far as the distinction between the two different VaR methodologies is concerned, this approach calculates VaR or MaxVaR by comparing the potential losses to the present value of the portfolio.

3.3.3. Long-Term Value at Risk¹⁶¹

The next approach was published in 2004 and the basic analysis underlying the presented VaR model has distinct similarities to the findings contained in Kupiec (1999)¹⁶² for the VaR methodology relating to the portfolio's present value.

¹⁶⁰ For example, the assumptions of independent log-normal return distributions with constant parameters (correlation, standard deviation, drift) or constant portfolio composition.

¹⁶¹ Dowd et al. (2004), pp. 52-57

¹⁶² Kupiec (1999), pp. 41-52

Dowd and collaborators define VaR as:¹⁶³

$$VaR(h) = P - P_{cl} = P - e^{[\mu h + \alpha_{cl} \sigma \sqrt{h} + \ln P]}$$

with:

- h* timeunits in the VaR horizon
- P* current portfolio value
- cl* VaR confidence level
- μ random mean log portfolio return (>0)
- σ volatility
- P_{cl} (1-cl) percentile of the terminal portfolio value
- α_{cl} standard normal variate for chosen confidence level

In terms of the standard variable definitions for this text, this formula is adjusted to the following:

$$VaR(T) = P_0 - P_T^c = P_0 - P_0 \left(e^{\bar{r}T - \alpha_c \sigma \sqrt{T}} \right)$$

This is basically the standard random-walk VaR model with normally distributed returns. The authors do not restrict their findings to this assumption, however, mentioning that the distribution of returns could be substituted to accommodate empirically proven characteristics of financial return data, such as heavy tails, without altering their main conclusions.¹⁶⁴

They do not perform any analyses as to potential effects of such modifications, however, leaving these issues up to further research.

Obviously the essence of their analysis is not embodied in the equation presented for their long-run VaR model. However, with this basic model as a foundation, Dowd et al. (2004) compare the effect of different volatility and mean return levels on the development of VaR figures across time horizons and confidence levels and find that with a positive mean return, for increasing time horizons VaR initially increases then tapers off and eventually decreases indefinitely.¹⁶⁵ This is due to the effect of the mean return and its compounding, which eventually dominates the increase of volatility.

For the case of zero mean return, VaR increases with time and approaches its upper bound of 100% of the initial portfolio value.

¹⁶³ Equation 2 in Dowd et al. (2004), page 53

¹⁶⁴ They mention Student-t distributed returns, or other alternatives to incorporate phenomena such as mean reversion or volatility clustering (GARCH). Dowd et al. (2004), pp. 53, 57

¹⁶⁵ i.e. the VaR (“worst case”) loss is in fact a profit.

This is also the main point of criticism the authors raise with respect to scaling VaR quantiles by the square root of time rule. They argue that this approach gravely overestimates VaR and that, when applying this rule, VaR has no such upper bound and increases to infinity with time.¹⁶⁶

The derivation of the input parameters, mean return and volatility, for their VaR model is where the authors differentiate themselves from other related treatments. They argue against the estimation of these parameters for smaller sub-periods, such as single days, and then applying these estimates to sequential sub-periods for the entire VaR horizon. Their reasoning revolves around possible trends in the return data, as the scaling of such a trend with increasing time horizons would eventually lead to an unrealistically high or low value and erroneous VaR figures. If, on the other hand, the data are best described by a zero trend, forecasting the variables for multiple sub-periods would be unnecessary as with increasing time horizon the value would eventually level out at some long-run average, yielding the sub-period forecasts unnecessary.¹⁶⁷ For these reasons the authors suggest simply applying an estimated mean return and volatility for the entire VaR horizon. Unfortunately, they do not provide suggestions as to how one could derive such long run average estimates.

These arguments are somewhat in line with the results presented by Embrechts et al. (2005) who find that the best results for calculating volatility and return estimates for a one year horizon are achieved with data at intermediate horizons.

In addition to the issues of parameter estimation and return distribution assumptions, Dowd et al. (2004) also examine the assumption of constant portfolio composition throughout the VaR horizon. Although they also subscribe to this simplifying assumption¹⁶⁸, they do not do so without briefly discussing it and providing a suggestion for deviating from this restriction, which is in contrast to most arguments in the relevant publications. The authors argue that an additional instrument could be added to the portfolio, thus providing a means of accurately representing a predetermined asset allocation strategy. As an alternative they suggest a

¹⁶⁶ The approach presented by Brummelhuis and Kaufmann (2007) to integrate a non-zero mean return can provide relief regarding this issue, however, when the mean return is or is assumed to be zero, this approach cannot be applied.

¹⁶⁷ A good example in support of this argument is the approach of producing sequential volatility estimates for a GARCH process, whereby eventually the volatility estimates converge to the long-run average component.

¹⁶⁸ This is achieved by implicitly reinvesting intermediate cash flows into the portfolio in a way which maintains the relative weights among the contained assets.

simulation approach which estimates VaR by calculating sequential sub-periods; however, this procedure is a contradiction to their general line of argumentation.

Fundamentally this article does not present any groundbreaking results, but does cover some of the essential elements and assumptions necessary for an accurate long-run VaR model and provides some valuable insights on long-run VaR.

3.3.4. A Term Structure of Risk

In contrast to the three previously presented articles which focused on a few specific issues when calculating a long-run VaR, Guidolin and Timmermann (2006a) present research and a comparative study for different possible long-run VaR models, some of which provide possibilities to handle or incorporate many of the identified critical issues present in a long-run VaR calculation, which the aforementioned models cannot.¹⁶⁹

For their analyses they use monthly return data from markets in the United States of America for three asset classes. Specifically they focus on equity data in the form of a value weighted portfolio for the New York Stock Exchange (NYSE), T-bills¹⁷⁰, and ten year U.S. Treasury bonds.

The authors focus on the equity and bond data, however, and base all their models and calculations on the excess returns of equity or bonds above the respective monthly T-bill rates, by subtracting these T-bill rates from the gross monthly returns of these two assets.¹⁷¹ Applying these data the authors calibrate five different VaR models and compare their performance for monthly horizons ranging from one month to two years, and two confidence levels, 95% and 99%.

The VaR models differ in their flexibility with regard to five different elements: Incorporation of time varying mean asset returns, time varying volatility, time varying correlations, flexible portfolio composition, and accurate reflection of the fat tails seen in empirical return data.

Regarding the flexibility of portfolio composition, the differentiation among the models is a fundamental one. If portfolio composition is assumed to be constant, the models are tailored

¹⁶⁹ Guidolin & Timmermann (2006a)

¹⁷⁰ T-bills are short term government debt securities issued by the U.S. Treasury. The authors make use of T-bills with one month, i.e. thirty day maturities.

¹⁷¹ For their asset return calculations they include possible dividend or coupon payments during a month.

directly to portfolio returns and are therefore univariate. When keeping portfolio composition flexible, the models describe single asset returns and must be multivariate.

Guidolin and Timmermann (2006a) estimate three different multivariate models. A fairly basic multivariate model is defined as the benchmark. It implements constant mean returns, correlations, and volatilities with standard normal innovations and provides a means to incorporate serial correlation in asset returns by including an autoregressive component. This model is moderately restrictive in comparison to the presented alternatives but, due to its multivariate nature, can allow for changing portfolio composition from month to month. In contrast the most flexible of all models under investigation is a multivariate Markov regime switching model.¹⁷² This model can integrate flexibility regarding all five of the just mentioned issues. Although mean return, volatility, and correlations are fixed for each state, the necessary flexibility is created by modeling the different states accordingly. For example, one could model a more probable “regular” state in which volatilities and mean returns are average and correlations reflect average market situations and a second state of market turmoil in which returns are negative, volatilities high and standard correlations break down, which occurs with a small probability. The challenge of producing a precise and efficient model lies in its calibration and the identification of significant different states in line with the underlying data.¹⁷³

The third multivariate model is a GARCH(1,1)-M model, based on Engle and Kroner (1995), with Student-t distributed innovations to replicate heavy tails. This model is similar to the univariate models which have been presented in this paper and is only restricted with regard to correlations, which are constant.

Aside from these three multivariate models which can model all asset returns separately, the authors include a univariate two component GARCH(1,1) model drawn from Engle and Lee

¹⁷² Regime switching models assume a discrete number of different regimes, i.e. states which have a certain probability of occurring. For each state a specific multivariate model is defined with distinct mean returns, volatilities, and correlations. In the model at hand Guidolin and Timmermann (2006a) define the four different regimes as a “realization of a first-order Markov chain with constant transition probabilities.” (Guidolin & Timmermann (2006a), page 287) The transition probabilities, which are summarized in a transition matrix, state the probability of moving from one state to another.

¹⁷³ For more information on multivariate regime switching models the authors refer to publications by Ang and Bekaert (2002) and Guidolin and Timmermann (2006b).

(1999) This model differs from basic GARCH models in two respects. On the one hand, the two component feature allows both the short-run element of volatility as well as the long-run average volatility to fluctuate, which provides additional advantages in replicating the heavy tails present in the relevant empirical distributions. On the other hand, this GARCH model includes the asymmetry feature that was previously mentioned.

The final method the authors include in their survey is a nonparametric block bootstrap, which draws a large number of samples equivalent in length to the VaR horizon and derives risk measures by calculating averages across the samples.¹⁷⁴

Guidolin and Timmermann (2006a) perform their evaluation of VaR measures based on the different models for five different portfolios to assess differences in performance and accuracy related to the three asset classes. The five portfolios are 50% stocks and 50% bonds, 50% bonds and 50% T-bills, 100% bonds, 50% stocks and 50% T-bills, and finally 100% stocks. In the first step of their analysis the authors use Monte Carlo simulations to calculate so-called “term structures of risk” for all models and portfolios at both the 95% and 99% confidence levels.¹⁷⁵ A general observation made with regard to the benchmark model is that VaR rises at first for the shorter time horizons until it reaches its peak and falls slowly after this peak. This is in line with the conclusions presented by Dowd et al. (2004). The initial analysis of unconditional VaR shows that the fluctuations across models and time horizons can be very large, with the greatest variations occurring at the longest time horizons. Depending on the portfolio, the 24 month VaR predicted by the five different models can range from approximately 0% to more than 30% for the 50% bonds and 50% T-bills portfolio or take values between approximately 20% and 43% for the 100% equity portfolio. As far as general tendencies for specific models are concerned, one can observe that the benchmark model has a propensity to generate the smallest VaR estimates for all time horizons and portfolios, except the 100% bond portfolio for which the regime switching model produces even smaller VaR figures. The multivariate GARCH(1,1)-M model produces the highest VaR estimates at long time horizons for four of the five portfolios and in general does not exhibit the universal tendency of reaching a peak and declining subsequently. It shows steady increases with time for all portfolios with distinct differences in the slope of the increases

¹⁷⁴ This method can be classified in the historical simulation category. For more information on block bootstrap methods in general please refer to Lahiri (1999) or Lahiri (2003).

¹⁷⁵ i.e. they chart the VaR figures over the entire time horizon up to the 24 month limit for every model and portfolio to highlight the different evolution of VaR over the full horizon for the various models.

amongst the portfolios. The univariate GARCH model tends to lead to the highest VaR figures at short horizons, but declines slowly for longer horizons and the bootstrap method produces average results across time horizons and portfolios. Based on these results one would have to take the asset classes and specific time horizon into account when choosing a model to implement.

As these results were unconditional VaR figures independent of the exact situation regarding the relative volatilities and current state at the time of evaluation, the authors evaluate potential effects of varying the initial states or volatilities at the time of evaluation to compare differences between unconditional and conditional VaR across time horizons.

These deviations can only occur in the regime switching and GARCH models, of course. The analysis shows that the effect is greatest for variations in the initial state of the regime switching model and smaller for the GARCH models, with noticeable differences amongst the portfolios. The results also contradict the arguments commonly brought forth regarding the irrelevance of initial states when looking at large time horizons, as even at the longest time horizon of two years the differences caused by varying starting points remain very significant.¹⁷⁶

Aside from these general and descriptive analyses of the distinct characteristics of the different VaR models for different time horizons, starting points, and portfolio compositions, the authors evaluate the predictive performance and precision by comparing out of sample forecasts with the actual empirical returns for horizons up to twelve months and confidence levels of 95% and 99%.

To assess the performance of the different models Guidolin and Timmermann (2006a) compare the actual number of exceedances to the expected number for the chosen confidence level and an additional measure as proposed by Christoffersen and Diebold (2000).

The results of the performance evaluations do not identify a single model as unambiguously optimal. The relative performance depends upon the asset class, horizon, and confidence level. The benchmark model tends to underestimate VaR especially at horizons longer than six months. The multivariate GARCH model overestimates VaR distinctly for all portfolios containing bonds across all time horizons but performs well for the 100% equity portfolio at

¹⁷⁶ It is argued that these differences are related to the fact that the cumulative return is of importance so that initial differences would not be evened out as time tends to infinity.

the 99% confidence level. The univariate GARCH model achieves better results at shorter horizons, but underestimates VaR at longer horizons for the 99% confidence level and across time horizons for 95% confidence. The bootstrap model outperforms the other models across portfolios and time horizons for the 95% confidence level, for which the other approaches exhibit a tendency of underestimating VaR. At the 99% confidence level the overall top performing model is the multivariate regime switching model.

In summary the bootstrap model is the top performer for a confidence level of 95%, whereas at 99% the regime switching model shows the highest level of precision.

This result is encouraging because when looking at the higher confidence levels, above 95%, the model that provides the most flexibility and can best deal with the critical issues for long-run VaR as highlighted in this paper is also the most precise.

3.3.5. Wavelet VaR¹⁷⁷

In two articles on applications of wavelet analysis¹⁷⁸ with regard to the calculation of CAPM¹⁷⁹ parameters and estimation of returns via the proposed CAPM implementations, Fernandez (2005 & 2006) presents VaR calculations for different time horizons. These two publications are focused on CAPM specific issues and only deal with VaR as a byproduct, so they will only be covered in a few words.

Basically, after estimating CAPM parameters from daily return data for the different time-scales, which are horizon intervals derived from the wavelet analysis¹⁸⁰, the author calculates VaR estimates for the underlying data for horizons up to one year. Due to the fact that these VaR calculations are based on CAPM models estimated to very specific emerging markets, the general applicability of the VaR element of the performed research is very limited. The only assets covered are stocks and exchange rates and no other asset class could even be included in the model as the CAPM is only applied to foreign and domestic equity. Additionally, when implementing such a VaR model one would automatically be restricted to

¹⁷⁷ Fernandez (2005) and Fernandez (2006)

¹⁷⁸ For detailed information on wavelet methods in general and in relation to financial markets Fernandez points to other publications such as Ramsey (1999 & 2002) or Ramsey and Zhang (1997).

¹⁷⁹ Capital Asset Pricing Model originally proposed by Sharpe (1964), Lintner (1965) and Mossin (1966).

¹⁸⁰ Seven different time scales are presented: 2 to 4 days, 4 to 8 days, 8 to 16 days, 16 to 32 days, 32 to 64 days, 64 to 128 days, and finally 128 to 256 days which is equivalent to one year.

assuming the CAPM as the mapping function for the stocks included in the portfolio VaR calculation at hand. This would restrict the approach even further and force all restrictions and potential flaws of the CAPM onto the VaR model. In spite of these cautionary comments, the results of the VaR analyses will be summarized briefly. Fernandez calculated 95% VaR figures measured in U.S. Dollars portfolio value per day for time scales 1 to 6, i.e. for horizons up to 128 days, and found that these values decreased steadily with increasing time horizons. Although the horizons are aggregated into time intervals and the analysis is limited to the 95% confidence level and a horizon of 128 days, these results are in line with other empirical evidence which suggests that VaR increases with time at a rate less than linear.

3.3.6. RiskMetrics Long-Run VaR Methodologies

3.3.6.1. LongRun VaR

The first of the two long-run VaR approaches conceived by the RiskMetrics Group can be derived from some of the Technical Documents published up to the year 2001, specifically the original Technical Document of 1996¹⁸¹, the Return to RiskMetrics document (2001)¹⁸², the CorporateMetrics Technical Document (1999)¹⁸³ and most importantly, the LongRun Technical Document (1999)¹⁸⁴ which is the key to calculating any risk measure for long time horizons in line with the RiskMetrics methodology.

Basically the model is a Full Valuation Monte Carlo Simulation method, as presented in the previous part of the thesis, with a time horizon of up to two years.

The elements of the calculation are the same as in the short run. The value of the assets is mapped to risk factors, these risk factors are simulated for the VaR horizon, and the distribution of portfolio values at the end of the VaR horizon depending on the generated scenarios is calculated from which VaR can then be estimated. Fundamentally, the building blocks are identical except for the risk factor scenario generation, which is more complex and detailed in the LongRun Technical Document (1999).¹⁸⁵

¹⁸¹ Longerstaeay (1996)

¹⁸² Mina & Xiao (2001)

¹⁸³ Lee et al. (1999)

¹⁸⁴ Kim et al. (1999)

¹⁸⁵ Kim et al. (1999)

The scenarios include daily price paths from the time of evaluation up to the end of the VaR horizon; the process is split up into two separate steps.

The first part is so-called Level I simulation. This entails estimating exact distributions¹⁸⁶ for all risk factors at all Level I dates or horizons. The standard choice for these Level I dates or horizons are monthly intervals, so that for a VaR horizon of one year there would be twelve Level I dates from one up to and including twelve months.

Aside from the distinct distributions for all risk factors at all Level I horizons, the second element necessary for completing Level I simulations is a covariance matrix.

This covariance matrix, however, is massively more complex than in the simple case of estimating an equivalent matrix for a single time horizon.

In fact the necessary covariance matrix can be interpreted as a large matrix consisting of many smaller matrices, which can also be referred to as a tensor, in the following form:¹⁸⁷

$$\Sigma_X = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \cdots & \Sigma_{1n} \\ \Sigma_{12} & \Sigma_{22} & & \\ & & & \\ \Sigma_{1n} & & & \Sigma_{nn} \end{bmatrix}$$

This matrix contains ‘n’ times ‘n’ sub matrices with ‘R’ times ‘R’ elements¹⁸⁸, respectively, representing all correlations amongst all risk factors across all Level I horizons.

In two extreme cases this matrix simplifies to two common covariance matrices. If ‘R’ is equal to one, i.e. there is only one Level I date, the matrix reduces to the simple covariance matrix for any standard VaR calculation, as every one of the sub matrices only contains a single element. In the other extreme case of ‘n’ equalling one, i.e. only one risk factor is considered, the matrix only contains autocovariances for a single risk factor across different Level I horizons. In the general cases the full covariance matrix is present, containing not only the covariances across risk factors for a single date or across all dates for a single risk factor but also all combinations of different risk factors at different dates with other risk factors at other dates.

However, as the estimation of ‘n’ squared times ‘R’ squared covariance values would be extremely laborious if not impossible in many cases, RiskMetrics presents some

¹⁸⁶ The document does provide flexibility, but the basic assumption is a multivariate normal distribution for asset, i.e. risk factor returns, so that the distributions are defined by estimating means and variances (Kim et al. (1999), page 137).

¹⁸⁷ Equation 5.2 in Kim et al. (1999), page 138

¹⁸⁸ ‘n’ being the number of risk factors and ‘R’ the number of Level I horizons

simplifications that reduce the complexity of estimating all elements of the matrix substantially while preserving the consistency of the estimation approach in general.¹⁸⁹

With this covariance matrix and the expected means and volatilities for the risk factors at all Level I dates, one can generate simulation paths containing risk factor scenarios for all Level I dates that are in line with the correlation structure and the respective risk factor distributions.

The second element of the scenario generating procedure is the Level II simulation which fills in the “gaps” between the Level I horizons so that every scenario contains a complete path for every risk factor for every single day from the evaluation date up to the end of the VaR horizon.

It is assumed that the daily risk factor values follow a random walk between two Level I dates but take the values generated from the Level I simulations at the Level I dates. The intermediate random walk is therefore in line with the estimated volatility. As the risk factors are assumed to follow a random walk with normally distributed innovations, i.e. a Geometric Brownian Motion, this procedure is called a “Brownian bridge”¹⁹⁰.

The scenarios for these Brownian bridges are simulated between Level I dates by letting the distributional parameters, means, variances, and covariances gradually move from those of one Level I horizon to those of the subsequent Level I horizon in line with an algorithm which is presented in the LongRun Technical Document (1999)¹⁹¹.

After performing both Level I and Level II simulations, the end result is a complete distribution of risk factors for every single day in the VaR horizon. From these risk factor scenarios and distributions, associated portfolio values can be calculated for each date and scenario and thereby the paths of portfolio value across the entire VaR horizon.

This property of the proposed VaR model is unique amongst the presented long-run VaR approaches. It provides possibilities of implementing flexibility regarding some of the restrictive assumptions that are not possible for other methods, computational resources permitting. As full valuation can be performed for any point in time in the VaR horizon, the

¹⁸⁹ The details regarding the covariance matrix estimation are presented on pages 138 to 141 and procedures to ensure the positive definiteness of the matrix are described on pages 151 to 152 of the LongRun Technical Document (Kim et al. (1999)).

¹⁹⁰ Kim et al. (1999), page 143

¹⁹¹ Kim et al. (1999), page 145

restriction of constant portfolio composition can be relaxed and daily portfolio changes are possible without any additional modifications of the VaR model itself. Time varying means, volatilities, and correlations can also be implemented.

Aside from presenting this forecasting approach for calculating a long-run Monte Carlo simulation VaR, the LongRun Technical Document¹⁹² also outlines different methods for estimating the necessary distributional parameters of volatility and mean for various risk factors associated with the main asset classes of equity, foreign exchange, interest rates, and commodities.

The authors distinguish two main approaches:

- The first is classified as “Forecasts based on current market prices”¹⁹³ and derives the distributional parameters directly from market data on the specific risk factors or assets themselves, such as spot prices and derivative instruments including forwards, futures, and options. This approach makes use of fundamental concepts of financial market and asset pricing theory¹⁹⁴ to derive market expectations for future distributions of asset prices from the current market prices of these different instruments. For example, forward and futures prices are used to estimate mean returns of risk factors and option prices to derive estimates of future volatility for various time horizons. The authors propose three different combinations of deriving the distributional parameters for a long-run VaR model from these current market prices. The simplest approach is the same as in standard short-run models. A zero mean return is assumed and the volatility estimate generated from historical return data. The other two approaches both assume the expected return to equal the forward premium, one generating the volatility estimate from historical data and the purely forward looking approach deriving it from implied volatilities.
- The second main approach of generating the distributional parameters for forecast generation is “based on economic structure”¹⁹⁵. This approach is more complex and elaborate, as risk factor distributions are not forecast from market data directly related

¹⁹² Kim et al. (1999)

¹⁹³ Kim et al. (1999), pp. 15-80

¹⁹⁴ e.g.: The efficient markets theory, risk neutral valuation, cost-of-carry concepts, and the expectation hypothesis.

¹⁹⁵ Kim et al. (1999), pp. 81-122

to the risk factors themselves, but an entire system of historical macroeconomic and financial data is assembled and modelled econometrically. The amount of utilized time series data is much more extensive, and in a way, the risk factors' distributions are mapped to the econometric model in similar fashion as the assets are mapped to the risk factors.¹⁹⁶ Such “mapping”, which basically defines the relationships amongst the time series or risk factors, is necessary to generate forecasts for the risk factor distributions. The economic structure created by relying on fundamental economic concepts¹⁹⁷ is an essential component of such econometric forecasting models. Kim et al. (1999) also provide a survey of many publications which analyzed such relationships for various asset classes in the derivation of their own model.¹⁹⁸ The specific econometric model upon which the authors focus their attention is a Vector Error Correction Model (VECM), a combination of a Vector AutoRegressive Model (VARM) and an Error Correction Model (ECM). This econometric model can incorporate both autoregressive features of time series and the cointegrating relationships present among the different time series.¹⁹⁹ An additional feature of the selected model is that it can generate forecasts for many time series simultaneously, which is a very important feature in the multivariate case. In their paper on risk assessment for pension funds, Bosch-Príncipe et al. (2002) present an implementation of such a VECM VaR model and calibrate it to data from the Spanish market.

¹⁹⁶ For example the proposed model only integrates equity index data directly so that forecasts for specific stocks must be generated by relating them to the equity index data and other relevant time series in the econometric model.

¹⁹⁷ e.g. purchasing power parity (PPP) or monetary theory for foreign exchange rates (Kim et al. (1999), page 104)

¹⁹⁸ On pages 105 to 108 Kim et al. (1999) refer to MacDonald & Taylor (1994), MacDonald & Marsh (1997) and Mark & Choi (1997) among others regarding foreign exchange; to Edwards & Khan (1985), Hall et al. (1992) and Caporale & Pittis (1997) among others regarding interest rates; Lee (1996) in relation to equity, and finally Pindyck (1993) and Pindyck & Rotemberg (1990) with regard to commodities.

¹⁹⁹ Such a cointegrating relationship is present when two or more different time series appear to follow a random walk when observed individually but the differences between the different time series are somewhat stable. For further details on the precise derivation and selection of the chosen econometric model, please refer to the original text.

In addition to the flexibility provided in estimating the parameters for the risk factor distribution forecasts, the model presented in the LongRun Technical Document²⁰⁰ can also be expanded to incorporate economic regime switches or the possibility of structural breaks in markets and, although the normal distribution is assumed in the model, the authors provide information on possible deviations from this assumption, further increasing the flexibility of their model.

Concluding the remarks on this model, it can be noted that the framework which is presented by Kim et al. (1999) provides vast flexibility with regard to all of the critical issues identified in the previous part of the thesis, but naturally, the more flexible such a model is designed to be, the more complex it becomes in conception, implementation, and also maintenance.

3.3.6.2. RiskMetrics 2006

The final and at the same time most current long-run VaR model presented is also proposed by the RiskMetrics Group in their publications on the RiskMetrics 2006 methodology.²⁰¹

This methodology, published in the first half of 2007, is the first drastic methodological innovation brought forth by the RiskMetrics Group since the original document of 1996²⁰². In contrast to the LongRun model which was just presented, the RiskMetrics 2006 methodologies main aim is not to provide long horizon VaR figures but to provide a significant increase in accuracy for VaR calculations across all time horizons up to horizons of twelve months. Given that the majority of the literature surveyed in this thesis seems to focus on horizons up to twelve months as well, with a few exceptions covering even longer time horizons, this approach fits perfectly into the current discussions.

The benchmark Zumbach uses as a point of reference in the comparisons of his model is the basic RiskMetrics VaR approach including the refinements presented in Mina and Xiao (2001).

This benchmark model has the following features: it is derived from a random walk, the mean return is assumed to be zero, volatility estimates are generated by an Exponentially Weighted

²⁰⁰ Kim et al. (1999)

²⁰¹ Zumbach (2007a), Zumbach (2007b) and Zumbach (2007c)

²⁰² Longestaey (1996)

Moving Average (EWMA) or equally weighted process for daily volatilities, and innovations are drawn from a standard normal distribution.²⁰³

The RiskMetrics 2006 approach also relies on the basic random walk equation to model the risk factor return process, but differentiates itself as to how the mean returns, volatilities, and innovations are derived. The main aim is to fully incorporate the empirical insights on the specific features of financial time series which have been gained in recent years, fat tails, and volatility clustering in particular. An additional motivation mentioned by Zumbach²⁰⁴ is to achieve the improvements in accuracy without introducing too many additional parameters in contrast to approaches based on GARCH(1,1) processes which contain a long-run mean volatility as one of its parameters which clearly differs across different risk factors.

The VaR approach is based on an ARCH model, originally introduced by Zumbach (2004 & 2006), which extends an I-GARCH²⁰⁵ process to a so-called Long-Memory ARCH (LM-ARCH) process. The main feature is that the weights assigned to the squared returns are not related linearly to the lag as in the EWMA approach, but decay logarithmically with increasing lags. The weights are also dependent upon the horizon for which the volatility is to be estimated. This element of the model reflects what Zumbach (2007b)²⁰⁶ calls an intuitive assumption, that is, that short horizon forecasts depend more on recent data and volatility forecasts for longer time horizons place more emphasis on data which lie further in the past. Moreover, the Long-Memory ARCH process, as proposed in the Riskmetrics 2006 methodology, takes risk factor data up to six years back into account²⁰⁷ as opposed to the equally weighted volatility which uses one year of historical data and the EMWA approach which, depending on the decay factor, reflects information going back up to approximately one year, for a decay factor of 0.97. Figure 4 presented below illustrates these points nicely.

²⁰³ The model is set up as a multivariate model for short horizons of less than three months, so that the zero mean return assumption is deemed unproblematic with reference to evidence presented in Kim, Malz and Mina (1999).

²⁰⁴ Zumbach (2007b), page 2

²⁰⁵ With an Exponentially Weighted Moving Average decay factor of 0.94 or 0.97.

²⁰⁶ Zumbach (2007b), page 7

²⁰⁷ Zumbach (2007a), page 25

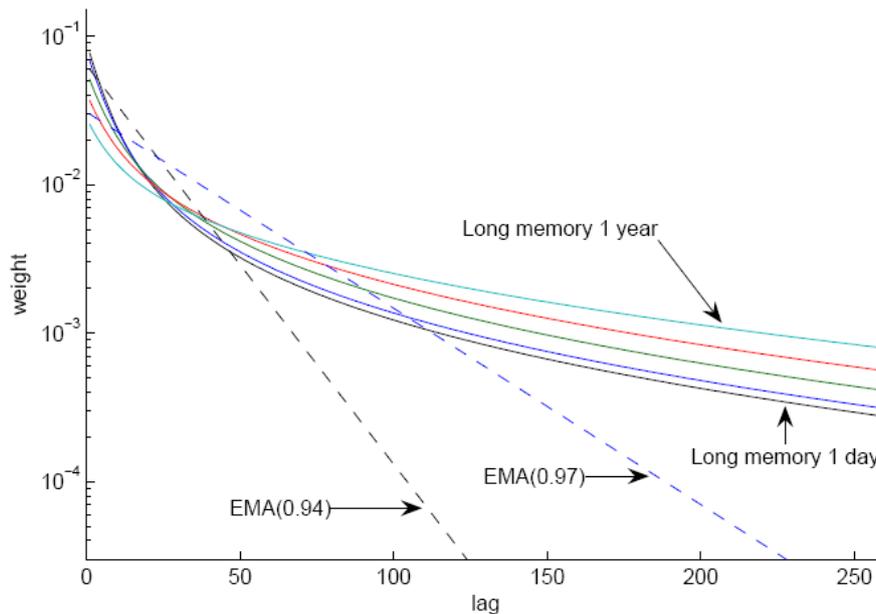


Figure 4 - Source: Zumbach (2007a), page 11

Zumbach²⁰⁸ also compares his approach to those of Embrecht et al. (2005) and McNeil and Frey (2000) which were presented in the previous section. He criticizes the facts that their approaches implement a GARCH(1,1) process which has the drawbacks that an exponential decay factor is implemented, that the time aggregation rule which is used²⁰⁹ cannot be applied to a different volatility process, and that the number of required parameters grows drastically with the number of risk factors when calibrating a VaR model.

The complete derivations and equations necessary for any implementation of Zumbach's proposed VaR model can be found in Zumbach (2007a)²¹⁰ and the results of the extensive backtesting which was performed to calibrate and verify the model are detailed in Zumbach (2007c). In this presentation of his model the focus will lie upon the key features he conceives that lead to the substantial improvements in accuracy of this VaR model.

One of these elements is the LM-ARCH process which produces volatility estimates that differ with respect to the time horizon but can be calibrated in its main parameters once, retaining validity across different risk factors. As already mentioned, this sensitivity with respect to the forecast horizon is achieved by adjusting the weights to place more emphasis on more recent or on the increasingly lagged data. Aside from this flexibility with respect to the

²⁰⁸ Zumbach (2007a), page 6

²⁰⁹ The Drost and Nijman approach published in 1993.

²¹⁰ Specifically in the appendices A and B (pp. 51 – 57).

time horizon, the main innovative feature of this approach is the fact that the volatility process has three main parameters all independent of the specific risk factors and only one additional parameter which must be adjusted to reflect the specificities of each risk factor. The volatility process is proposed and calibrated in a way that incorporates the critical empirical features common to most if not all financial time series such as heavy tails and volatility clustering, while also integrating the long memory of these data sets but without over-fitting to single assets which would diminish the valuable property of generality, and necessitate deriving multiple specific process parameters for the different risk factors. This is a significant improvement in comparison to other methodologies such as the RiskMetrics LongRun approach in which a mean return and volatility estimate for each Level I date has to be estimated separately.

The second element is incorporating the autoregressive tendencies that sequential returns for financial time series exhibit and non-zero expected mean returns. Although Zumbach (2007a)²¹¹ argues that any autoregressive trends should be zero for liquid and free floating assets, because if this were not the case market participants would exploit any such correlations until they were eliminated by market forces, one cannot deny their presence in the empirical data. As an example, short-term interest rates are mentioned which are fixed by central banks and the timing and extent of any changes is usually fairly predictable for time horizons in the area of a month and trends for longer periods can usually be foreseen in line with economic policy in the respective regions.²¹² As far as the zero mean return assumption is concerned, the other contradictory example which is listed are returns on stocks and equity indices which require an expected mean return above zero in line with their risk profile. Zumbach (2007a) performs an empirical analysis on this issue and finds that for most asset classes no significant autocorrelation can be identified, but that for interest rate data two interesting phenomena present themselves. On the one hand, there seems to be significant positive autocorrelation for time horizons between ten days and three months in line with the central bank argument, but on the other hand, for longer horizons above a year, the time series exhibit negative autocorrelation which can be explained by mean reversion tendencies. The incorporation of an non-zero expected mean return and autoregressive features of returns into

²¹¹ Zumbach (2007a), page 29

²¹² This argument was somewhat contradicted empirically by the events in financial markets in the fall of 2007 when particularly short-term interest rates displayed extreme and unpredictable fluctuations caused by liquidity problems in the financial markets.

the model, which significantly increases the complexity of the model, is presented in Appendix B of Zumbach (2007a)²¹³.

The third essential feature of the proposed methodology is the extensive treatment of the innovations or errors contained in the random walk process. Zumbach (2007a)²¹⁴ performs a detailed analysis of these errors. The starting point for his analysis is calculating empirical error distributions from the numerous financial time series he employs throughout his papers by relating the returns forecast by his model to the actual ex-post returns for the relevant periods. Using these empirical error distributions as the database, he performs different analyses on these distributions with the aim of identifying any distinctive and/or detrimental features which would need to be corrected, and also of determining which of the commonly assumed statistical distributions most accurately replicates them. First he tests for autocorrelation within these error distributions, to verify the assumption that any autocorrelation in these time series would be modelled in the volatility and mean return forecasts, so that the errors should not display any such behaviour. The results of the analysis for the different asset classes confirm the assumption that no meaningful autocorrelation is present in the error distributions, so that no modification or correction in this regard is necessary. The next issues are the mean and variance of the error distributions, which are required to be equal to 0 and 1, respectively.²¹⁵

In his analysis Zumbach (2007a) confirms that the mean of the empirical error distributions across asset classes does not measurably deviate from zero, validating the mean return forecasts produced by the model. As for the variance, on the other hand, there are large deviations from 1 with increasing time horizons for some asset classes. To reduce this systematic overestimation of risk caused by very high variances in the empirical error distributions, an additional component is included in the model. This element is a scaling factor²¹⁶ dependent on the forecast horizon which corrects the residuals and significantly reduces the problem of variances exceeding the ideal value of 1.

After analysing the properties of the empirical error distributions, these distributions are compared to some standard distributions often implemented in VaR models, such as the standard normal distribution or Student-t distributions with different degrees of freedom.

²¹³ Zumbach (2007a), pp. 53-56

²¹⁴ Zumbach (2007a) pp. 32-45

²¹⁵ If this were not the case, it would introduce a systematic bias into any VaR calculations.

²¹⁶ The scaling factor is presented in equation 17 on page 13 of Zumbach (2007a).

Obviously, for applications in VaR models, the consistency of the distributional tails is of primary relevance.

The results of these comparisons are that for very short time horizons the standard normal distribution significantly underestimates risk in the tails and is outperformed notably in replicating the empirical distributions by a Student-t distribution with 5 degrees of freedom²¹⁷. For increasing time horizons, particularly in excess of three months, the results are no longer so clear and one cannot attest advantages of correct replication to either a standard normal or a Student-t distribution. In an additional analysis on this issue Zumbach finds that choosing a heavy-tailed distribution improves the results, but that introducing a time dependence feature into the number of degrees of freedom does not provide any considerable advantages and that the sensitivity of results for Student-t distributions with relation to the chosen degrees of freedom is very low in the interval between 4.5 and 8.

In short, when choosing amongst standard distributions, one should select a Student-t distribution with degrees of freedom between five and eight.

In a final benchmarking of the three key elements of the described RiskMetrics 2006 methodology against the standard RiskMetrics VaR approach, Zumbach (2007a)²¹⁸ highlights the contribution of each of the three elements to the improvement in accuracy of the VaR model and the cumulative improvement achieved by including all three elements.

Concluding the presentation of this VaR model three points must be made:

- Firstly, comparing the accuracy of this approach to the standard VaR models as presented in the previous RiskMetrics documents²¹⁹ across time horizons, the backtesting results which are presented in Zumbach (2007c) show that the improvements are dramatic for all presented measures of accuracy. For some, the improvement is of such a magnitude that the new methodology is more accurate at time horizons up to three months than the standard model is at a one day horizon.

The following graph in Figure 5, depicting one of the accuracy measures²²⁰, illustrates this point nicely.

²¹⁷ Three degrees overestimate tail risk, while eight underestimate it.

²¹⁸ Zumbach (2007a), pp. 48-50

²¹⁹ Longerstaeey (1996) and Mina & Xiao (2001)

²²⁰ An aggregated error measure derived in Zumbach (2007a) on page 45, for which values as close as possible to zero are most desirable.

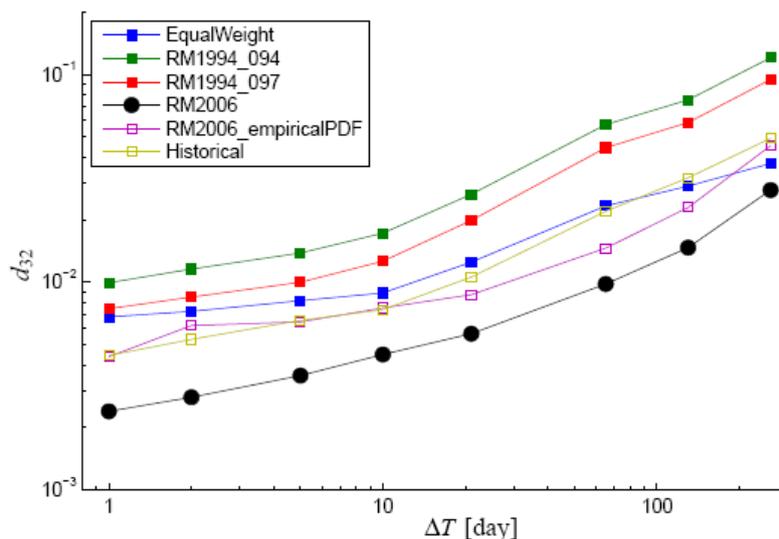


Figure 5 – Aggregated Error Measure – Source: Zumbach (2007c), page 28

- Second, the model is univariate and Zumbach does not present a means of extending the framework to a multivariate case. This diminishes the practical use significantly until such an extension is provided. One could apply the standard extension to a multivariate framework by applying a correlation structure to the errors of the single risk factors but this would require additional backtesting. Alternatively, the approach could be applied to aggregated portfolio returns. This would, however, introduce the restrictions regarding correlation and portfolio compositions inherent with these approaches into any such VaR calculations.
- Finally, the RiskMetrics 2006 framework is based on daily data, as opposed to the one presented by Dowd et al. (2004) and Zumbach (2007b) argues that basing such a model directly on returns for longer time horizons would equate to “essentially throwing away most of the information.”²²¹ Taking the results which are presented into account, one would have to agree with Zumbach as his model seems to incorporate the advantages of complex VaR models based on ARCH approaches, without being exposed to their drawbacks which usually present themselves at longer time horizons.

²²¹ Zumbach (2007b), page 12

3.4. Reviewing the Critical Issues for an Optimal Long-Run VaR Model

To conclude the elaborate presentation of approaches for calculating long-run VaR, ranging from very simple scaling approaches to very complex models, the final section of this part of the thesis will revisit the main restrictive issues, parameters and assumptions inherent in short-run VaR models which were pointed out in the previous part and summarize the possible improvements which were presented in the previous section. This could provide something of an outline as to which features an ideal long-run VaR model could or should incorporate.

3.4.1. Constant Portfolio Composition

Although this is a critical issue the models and approaches which provide flexibility with regard to this restriction are scarce. The MaxVaR approach brought forth by Boudoukh et al. (2004), provides a simple approach of quantifying some of this additional intra-horizon risk, but as far as an explicit modelling of flexible or adaptive portfolio composition is concerned, only those models that model the complete paths of risk factor development across the entire VaR horizon can really incorporate such flexibility. Examples would be the LongRun methodology²²² or some of the models presented in Guidolin and Timmermann (2006a). All multivariate models and scaling approaches cannot deviate from the constant portfolio composition assumption and implicitly do not accurately model the actual portfolio composition for long time horizons due to the reasons which were identified in the previous part of the thesis.

3.4.2. Standard Distributional Assumptions and Parameters

Quite a few models provide means of incorporating error distributions that are not so restrictive as the standard normal distribution and the analyses which were presented showed that in some cases choosing another distribution with heavier tails can be advantageous; however, surprisingly in some of the approaches the standard assumption performed best so that one cannot discard it as generally inferior.

As far as constant volatilities and mean expected returns are concerned, there are both univariate and multivariate models that can introduce flexibility with regard to these parameters and these were shown to perform very well. Regime switching models for

²²² Kim et al. (1999)

example or the two RiskMetrics approaches can incorporate varying volatilities and mean expected returns. Deviating from the constant correlation, on the other hand, can only be achieved with multivariate models, but especially the regime switching models can achieve this in a way that can accurately replicate actual empirical behavior of financial time series.

The scaling approaches, conversely, cannot incorporate flexibility with regard to time varying distributional parameters or portfolio composition, but as was shown, the relative performance of some of these approaches can be fairly good. Of course one should try to start with a precise and realistic short-run VaR model and attempt to incorporate as much precision and consistency within these approaches to get the most out of these simpler approaches.

4. Conclusion

The demand for accurate and efficient long-run Value at Risk models has grown over the last few years leading to an increase in publications on the topic and also in proposed models able to handle the special challenges that appear when extending a VaR horizon to horizons of a year or more. In a first step, this thesis identified the critical issues and shortcomings which are present and in some cases inherent in standard short-run VaR models and approaches. Building on this foundation, the simple approaches of extending the horizon of a VaR calculation while implementing a short-run model, by means of scaling either the VaR quantiles directly or the volatility estimates used to calculate VaR, were surveyed and somewhat surprisingly, comparative studies show that these approaches can perform well despite their restrictive assumptions. One assumption that must be amended, however, is that of a mean return of zero, because this would contradict fundamental concepts of financial theory when covering time spans of multiple months or even years. Of course, ideally, an implementation of a long-run VaR model should strive to be as realistic as possible and an assumption such as constant portfolio composition over a horizon of a year cannot be deemed realistic for reasons mentioned at the outset. This issue as well as some of the other critical issues such as constant correlation volatility or mean return can be accurately accommodated by some of the specific long-run VaR models. This increase in realism and flexibility comes at a cost of increased complexity and computational demands, however, so that as with any VaR model the tradeoff along these two criteria must always be considered.

Concluding, it can be stated that although there are some models and approaches which are shown to outperform others in the literature, there is no single model or approach that can be deemed superior for all possible combinations of portfolios, horizons, or even confidence levels. Various models perform better or worse than others depending on the assets which are contained in the respective portfolio, perform better or worse depending on which level of the distributional tails one is most interested in and, very importantly of course, how long the actual VaR horizon of interest is. The optimal VaR model for a specific implementation must be found by taking the special criteria of each individual situation into account but maintaining some degree of flexibility to adapt to changes in the demands or external factors.

As Zumbach puts it “To some extent, building a risk methodology is an engineering problem: one has to pick the important parts and neglect smaller or particular effects. This approach

contrasts with an academic study, where often a rigorous answer to a well posed question is sought after.”²²³

²²³ Zumbach (2007a), page 7

Appendix

Abstract – English Version

The Value at Risk (VaR) methodology which was originally developed by banks and other financial institutions with their own risk measurement and management needs in mind was and remains the foundation for numerous VaR calculation methods, approaches, and models that were conceived over the past fifteen years. Naturally these approaches and models catered to the needs of the user group and as these institutions all had a risk management focus aligned towards very short time horizons of a single day up to a week, or a month at the very most, the models aimed to produce reliable results for precisely these short time horizons. Therefore the models themselves, the underlying assumptions, and the necessary input parameters were determined in a way that achieved optimal results for the needs at hand whilst keeping the approaches as simple as possible to facilitate efficient VaR calculations.

Over time, risk management functions in non-financial corporations and other financial institutions expanded and evolved leading to increased demand for risk measurement tools that could serve their needs. The time horizons relevant to these market participants, namely, periods of three months up to two years, exceeded by far those which the existent models and approaches were geared towards.

This led to an increasing demand for adequate modifications to existing models and innovative approaches and models that could produce accurate results for these longer VaR horizons.

This thesis is focused upon possibilities of calculating long-run VaR figures and the approaches and models that can be implemented to achieve accurate risk-measurements for long time horizons.

Initially, the general elements and aspects of the VaR methodology are presented and the standard short-run VaR models are surveyed and analyzed with regard to their applicability or even expandability towards the longer VaR horizons. During this evaluation the simplifications present in these short-run models are highlighted and critically evaluated as to their validity in long-run scenarios.

After providing for a fundamental understanding of the basic VaR concepts, standard VaR, and critical issues with regard to long-run VaR, the main part of this thesis presents an extensive survey of possible approaches and models for calculating long-run VaR for horizons up to and in excess of a twelve month horizon. These approaches range from straightforward extensions of short-run VaR models by means of applying a simple scaling factor to short-run

VaR figures, over specific but fairly intuitive long-run VaR models, up to the most current state of the art long-run VaR approaches and models which are rather complex both conceptually and also in implementation.

Throughout these discussions various comparative studies evaluating the accuracy of different approaches and their results are presented in an attempt to determine an optimal or fundamentally best-suited approach or model for calculating long-run VaR.

Somewhat surprisingly, under certain conditions, even the most simple approaches can produce very accurate results and thereby outperform more complex models with regard to efficiency.

Finally, although certain critical issues which are highly relevant in a long-run setting are identified and should be taken into account when calculating long-run VaR, no approach or model can be deemed dominant or generally superior, due to the fact that the relative performance of various approaches depends greatly upon the specific situation and needs of the respective institution looking to implement such a long-run VaR model.

Abstract – German Version

Die Value at Risk (VaR)-Methodologie, welche ursprünglich von Banken und anderen Finanzinstituten unter Berücksichtigung ihrer eigenen Anforderungen an Risikomessung und Risikomanagement entwickelt wurde, war damals und ist nach wie vor die Grundlage für unzählige VaR-Berechnungsmethoden und –Modelle. Diese wurden im Laufe der vergangenen fünfzehn Jahre entwickelt. Grundsätzlich sind diese Methoden und Modelle alle auf die Bedürfnisse der ursprünglichen Benutzergruppe ausgerichtet. Da der Risikomanagementfokus dieser Institutionen auf sehr kurzen Zeithorizonten, d.h. von einem einzigen Tag, einer Woche bis hin zu maximal einem Monat lag, sollten die entwickelten Modelle zuverlässige und präzise Ergebnisse für exakt diese kurzen Zeiträume generieren. Aus diesem Grund wurden die VaR-Modelle selbst, die Basisannahmen und die notwendigen Input-Parameter derart angelegt und bestimmt, dass sie in Hinblick auf die Anforderungen der Anwender optimale Ergebnisse lieferten. Gleichzeitig wurden die Methoden so unkompliziert wie möglich gehalten um die Berechnungen optimal effizient zu gestalten.

Im Laufe der Zeit wurde dem Risikomanagement jedoch auch in anderen Gruppen von Finanzinstituten und Unternehmen außerhalb des Finanzsektors größere Bedeutung beigemessen, sodass diese Unternehmen einen größeren Bedarf an solchen adäquaten Risikobewertungsinstrumenten hatten, die ihren Bedürfnissen bzw. Anforderungen gerecht wurden. Die Zeithorizonte, welche für diese neuen Anwendergruppen relevant waren und nach wie vor sind, und zwar drei Monate bis hin zu zwei Jahren, gingen weit über jene hinaus auf denen die bestehenden Modelle und Methoden ausgerichtet waren. Dies führte zu einer verstärkten Nachfrage nach geeigneten Modifikationen und Erweiterungen für bestehende Modelle bzw. nach innovativen Methoden und Modellen, die präzise Ergebnisse für diese längeren VaR Zeithorizonte liefern könnten.

Der Fokus dieser Magisterarbeit richtet sich auf Möglichkeiten VaR für lange Zeithorizonte zu berechnen und auf den Methoden und Modellen die implementiert werden können, um präzise Risikomessungen und -kennzahlen für lange Zeiträume zu ermitteln.

In einem ersten Schritt werden die allgemeinen Elemente und Aspekte der VaR-Methodologie dargelegt und die Standard VaR-Modelle für kurze Zeithorizonte untersucht und in Hinblick auf ihre Anwendbarkeit und Ausbaufähigkeit für längere VaR Zeithorizonte analysiert. Im Rahmen dieser Evaluierung werden die Vereinfachungen, welche in den VaR-Modellen für kurze Zeithorizonte angenommen werden, hervorgehoben und in Bezug auf deren Gültigkeit bei längeren Zeithorizonten kritisch bewertet.

Nach der Vorstellung der grundlegenden VaR-Konzepte, der Standard VaR-Modelle und der kritischen Aspekte in Bezug auf „long-run“- VaR, wird im Hauptteil der Magisterarbeit ein umfassender Überblick über mögliche Methoden und Modelle zur Berechnung von VaR für lange Zeithorizonte von bis zu zwölf Monaten und darüber hinaus gegeben. Diese Methoden reichen von einfachen Erweiterungen der Standard VaR- Modelle für kurze Zeithorizonte, welche simple Multiplikatoren einsetzen um „long-run“ VaR zu berechnen, über spezifische und dennoch einfache „long-run“ VaR Modelle bis hin zu speziellen „state of the art“ Modellen die sowohl konzeptionell als auch in der Umsetzung hoch komplex sind. Im Rahmen dieser Darstellung werden verschiedene Vergleichsstudien welche sich mit der Präzision von verschiedenen Methoden beschäftigen vorgestellt, um zu eruieren, ob es Modelle oder Ansätze gibt, die bei der Berechnung von „long-run“ VaR fundamental überlegen sind.

Überraschenderweise können auch die einfachsten Methoden in gewissen Rahmenbedingungen sehr präzise Ergebnisse liefern und dadurch in Bezug auf Effizienz komplexere Modelle übertreffen. Letztlich kann man festhalten, dass kein Ansatz oder Modell als allgemein überlegen oder dominant bezeichnet werden kann, da die relative Performance verschiedener Ansätze sehr stark von der speziellen Situation und den Bedürfnissen der Institution abhängen, die ein „long-run“ VaR-Modell implementieren möchte. Man sollte jedoch gewisse kritische Bereiche, welche in der Betrachtung von langen Zeithorizonten hochrelevant sind, auf jeden Fall bei der Berechnung von „long-run“ VAR berücksichtigen.

About the Author

Erich Arthur Stark, a citizen of Austria and the United States of America, was born, reared, and schooled in Vienna, Austria. He attended the Volksschule Boerse-gasse and then the Bundesgymnasium Wien IX Wasagasse. After graduation with the Matura in 1996 he enrolled at the University of Vienna where during the first few semesters he pursued interests in the fields of International Business Administration, Law, and Business Informatics by attending corresponding courses.

He then focused on Business Administration with an emphasis on corporate finance, financial management, and operations management and attained his bachelor's (Bakk. rer. soc. oec.) degree in 2004.

He continued his studies at the University of Vienna for his Master's degree in Business Administration, specializing in banking, risk management, and financial economics - graduation pending.

In 2006 he co-authored two publications in the series "*Wirtschaft und Management*":

Thomas Happ & Erich Stark: "Treasury als Instrument zur Risikominimierung in Corporates", published in: *Risikomanagement in Unternehmen*, Volume 3, Number 4 of the series "*Wirtschaft und Management*", Mai 2006, pp. 101-114.

Robert Schwarz, Thomas Happ & Erich Stark: "Cashflow-Modellierung: Am Beispiel der Telekom Austria AG", published in *Risikomanagement*, Volume 3, Number 5 of the series "*Wirtschaft und Management*", November 2006, pp. 21-39.

Since 2003 he has been working full time as a senior financial analyst in the Corporate Finance team of the Finance and Treasury department of Telekom Austria Group. His main responsibilities in that position include analysis and valuation of financial instruments and markets, monitoring and optimization of the company's funding portfolio and hedging activities, working on various financial transactions, performing different cost of capital assessments, and interacting with rating agencies and investment banks.

Bibliography

- Al Janabi, M. A. (2007). Risk Analysis, reporting and control of equity trading exposure: Viable applications to the Mexican financial markets. *Journal of Derivatives & Hedge Funds*, 13 (1), 33-58.
- Almgren, R., & Chriss, N. (2001). Optimal Execution of Portfolio Transactions. *The Journal of Risk*, 3 (Number 2 (Winter 2000/2001)), 5-39.
- Andersen, T. G., & Bollerslev, T. (1998). Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts. *International Economic Review*, 4 (Symposium on Forecasting and Empirical Methods in Macroeconomics and Finance), 885-905.
- Andersen, T. G., & Bollerslev, T. (1997). Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns. *Journal of Finance*, 52, 975-1005.
- Ang, A., & Bekaert, G. (2002). Regime Switches in Interest Rates. *Journal of Business and Economic Statistics*, 20 (2), 163-182.
- Baillie, R. T., Bollerslev, T., & Mikkelsen, H. O. (1996). Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 74 (1), 3-30.
- Basel Committee on Banking Supervision. (2005). *Amendment to the Capital Accord to incorporate market risks*. Bank for International Settlements (BIS), Basel Committee on Banking Supervision. Basel: Bank for International Settlements (BIS).
- Beirlant, J., Teugels, J. L., & Vynckier, P. (1996). *Practical analysis of extreme values*. Leuven: Leuven University Press.
- Black, F., & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81 (2), 637-654.
- Blake, D., Cairns, A., & Dowd, K. (2000). Extrapolating VaR by the Square-Root Rule. *Financial Engineering News* (17), 3&7.
- Bodie, Z., Kane, A., & Marcus, A. J. (2005). *Investments* (5th Edition ed.). New York: McGraw-Hill.
- Bosch-Príncipe, M., Devolder, P., & Domínguez-Fabián, I. (2002). Risk analysis in asset-liability management for pension fund. *Belgian Actuarial Bulletin*, 2 (1), 80-91.
- Boudoukh, J., Richardson, M., Stanton, R., & Whitelaw, R. F. (2004). MaxVaR: Long Horizon Value at Risk in a Mark-to-Market Environment. *Journal of Investment Management*, 2 (3), 14-19.
- Brennan, M., & Schwartz, E. (1979). A Continuous-Time Approach to the Pricing of Bonds". *Journal of Banking and Finance*, 3, 133-153.

- Brummelhuis, R., & Guégan, D. (2005). Multiperiod conditional distribution functions for conditionally normal GARCH(1,1) models. *Journal of Applied Probability*, 42 (2), 426-445.
- Brummelhuis, R., & Kaufmann, R. (2007). Time-scaling of value-at-risk in GARCH(1,1) and AR(1)-GARCH(1,1) processes. *Journal of Risk*, 9 (4), 39-94.
- Caporale, G. M., & Pittis, N. (1997). Domestic and external factors in interest rate determination. *Applied Economics Letters*, 7, 465-471.
- Chance, D. M. (1998). A Brief History of Derivatives. In D. M. Chance, *Essays in Derivatives* (1st Edition ed., pp. 16-20 of 333). New Hope, Pennsylvania: Wiley.
- Chance, D. M. (1995). A Chronology of Derivatives. *Derivatives Quarterly*, 2 (Winter 1995), 53-60.
- Christoffersen, P. F., & Diebold, F. X. (2000). How Relevant is Volatility Forecasting for Financial Risk Management? *The Review of Economics and Statistics*, 82 (No. 1), 12-22.
- Cox, J., Ingersoll, J., & Ross, S. (1985). A Theory of the Term Structure of Interest Rates. *Econometrica*, 53, 385-408.
- Crouhy, M., Galai, D., & Mark, R. (2006). *The Essentials of Risk Management* (1st Edition ed.). New York: McGraw-Hill.
- Culp, C. L., Mensink, R., & Neves, A. M. (1998). Value at Risk for Asset Managers. *Derivatives Quarterly*, 5 (2), 21-33.
- Cuthbertson, K., & Nitzsche, D. (2005). *Quantitative Financial Economics: Stocks, Bonds & Foreign Exchange* (Second Edition ed.). Chichester, West Sussex, England: John Wiley & Sons Ltd.
- Dacorogna, M. M., Mueller, U. A., Pictet, O. V., & de Vries, C. G. (2001). Extremal Forex Returns in Extremely Large Data Sets. *Extremes*, 4 (2), 105-127.
- Dale, R. (1996). *Risk and Regulation in Global Securities Markets*. Chichester, U.K.: John Wiley & Sons.
- Daniélsson, J. (2002). The Emperor has no Clothes: Limits to Risk Modelling. *Journal of Banking & Finance*, 26, 1273-1296.
- Daniélsson, J., & de Vries, C. G. (2000). Value-at-Risk and Extreme Returns. *Annales d'Economie et de Statistique*, 60 (Special Issue (2000)), 236-269.
- Daniélsson, J., & Zigrand, J.-P. (2006). On time-scaling of risk and the square-root-of-time rule. *Journal of Banking & Finance*, 30, 2701-2713.
- Daniélsson, J., Hartmann, P., & de Vries, C. G. (1998, January). The Cost of Conservatism: Extreme Returns, Value-at-Risk, and the Basle Multiplication Factor. *Risk*, 101-103.

- Das, S. (2006). *Risk Management (Third Edition Revised)*. Singapore: John Wiley & Sons (Asia) Pre Ltd.
- Diebold, F., Hickman, A., Inoue, A., & Schuermann, T. (1998). Scale Models. *Risk*, 11, 104-107.
- Dimson, E., & Marsh, P. (1995). Capital Requirements for Securities Firms. *Journal of Finance*, 50 (3), 821-851.
- Dowd, K. (2005). *Measuring Market Risk (Second Edition)*. West Sussex, England: John Wiley & Sons Ltd.
- Dowd, K., Blake, D., & Cairns, A. (2004). Long-Term Value at Risk. *Journal of Risk Finance*, 2 (5), 52-57.
- Drost, F. C., & Nijman, T. E. (1993). Temporal Aggregation of GARCH Processes. *Econometrica*, 61 (4), 909-927.
- Duffie, D., & Pan, J. (1997). An Overview of Value at Risk. *The Journal of Derivatives*, 4 (3), 7-49.
- Dunis, C. L., Laws, J., & Naim, P. (2005). *Applied Quantitative Methods for Trading and Investment*. (C. L. Dunis, J. Laws, & P. Naim, Eds.) Chichester, West Sussex: John Wiley & Sons Ltd.
- Eckhardt, R. (1987). Stan Ulam, John von Neumann, and the Monte Carlo Simulation. *Los Alamos Science (Special Issue (15))*, 131-137.
- Edwards, S., & Khan, M. S. (1985). Interest rate determination in developing countries: a conceptual framework. *IMF Staff Papers*, 32, 377-403.
- Embrechts, P. (2000). *Extremes and Integrated Risk Management*. (P. Embrechts, Ed.) London, U.K.: Risk Books.
- Embrechts, P., Kaufmann, R., & Patie, P. (2005). Strategic Long-Term Financial Risks: Single Risk Factors. *Computational Optimization and Applications*, 32, 61-90.
- Embrechts, P., Klueppelberg, C., & Mikosch, T. (2003). *Modelling extremal events for insurance and finance* (corr. 4th printing edition ed.). Berlin [u.a.]: Springer.
- Embrechts, P., Resnick, S., & Samorodnitsky, G. (1998). Living on the Edge. *Risk*, 11 (1), 96-100.
- Engle, R. F., & Kroner, K. F. (1995). Multivariate simultaneous Generalized ARCH. *Economic Theory*, 11 (1), 122-150.
- Engle, R. F., & Lee, G. G. (1999). A long-run and short-run component model of stock return volatility. In H. White, R. F. Engle, H. White, & R. F. Engle (Eds.), *Cointegration, Causality*,

- and Forecasting: a Festschrift in Honour of Clive W.J. Granger* (pp. 475-497). Oxford, U.K.: Oxford University Press.
- Engle, R. (2001). GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics. *Journal of Economic Perspectives*, 15 (4), 157-168.
- Engle, R. (2004). Risk and Volatility: Econometric Models and Financial Practice. *The American Economic Review*, 94 (3), 405-419.
- Fama, E. F., & Miller, M. H. (1972). *The Theory of Finance*. Hinsdale, Illinois: Dryden Press.
- Feller, W. (1971). *An Introduction to Probability Theory and Its Applications* (2nd Edition ed., Vol. II). New York, New York: John Wiley & Sons.
- Fernandez, V. P. (2006). The CAPM and value at risk at different time-scales. *International Review of Financial Analysis*, 15, 203-219.
- Fernandez, V. P. (2005). The International CAPM and a Wavelet-Based Decomposition of Value at Risk. *Studies in Nonlinear Dynamics & Econometrics*, 9 (4), 1-27.
- Field, P. (2003). *Modern Risk Management: A History*. (P. Field, Ed.) London: Risk Books.
- Fight, A. (2004). *Understanding International Bank Risk* (1st Edition ed.). Chichester: John Wiley & Sons.
- Global Derivatives Study Group. (1993). *Derivatives: Practices and Principles*. Group of Thirty, Global Derivatives Study Group. Washington D.C.: Group of Thirty.
- Gopikrishnan, P., Plerou, V., Liu, Y., Aramal, L., Gabaix, X., & Stanley, H. (2000). Scaling and correlation in financial time series. *Physica A*, 287, 362-373.
- Guidolin, M., & Timmermann, A. (2006a). Term structure of risk under alternative econometric specifications. *Journal of Econometrics*, 131, 285-308.
- Guidolin, M., & Timmermann, A. (2006b). An econometric model of nonlinear dynamics in the joint distribution of stock and bond returns. *Journal of Applied Econometrics*, 21, 1-22.
- Hall, A. D., Anderson, H. M., & Granger, C. W. (1992). A cointegration analysis of treasury bill yields. *The Review of Economics and Statistics*, 116-125.
- Hill, B. M. (1975). A Simple General Approach to Inference About the Tail of a Distribution. *The Annals of Statistics*, 3 (5), 1163-1174.
- Holton, G. A. (2002). History of Value-at-Risk: 1922-1998. *Working Paper*, 0207001, 27.
- Holton, G. A. (2003). *Value-at-Risk: Theory and Practice* (1st ed.). San Diego: Academic Press.
- Hull, J. C. (2003). *Options, Futures, and Other Derivatives (Fifth Edition)* (5th Edition ed.). Upper Saddle River, New Jersey: Prentice Hall - Pearson Education Inc.

- Hull, J., & White, A. (1998). Incorporating volatility updating into the historical simulation method for value-at-risk. *Journal of Risk*, 1 (1), 5-19.
- Iacono, F., & Skeie, D. (1996, October 14). Learning Curve (R) Translating VaR Using $\{\text{square root}\}$ T. *Derivatives Week*, 8.
- Jordan, J. V., & Mackay, R. J. (1997). Assessing Value at Risk for Equity Portfolios: Implementing Alternative Techniques. In R. J. Schwartz, C. W. Smith, R. J. Schwartz, & C. W. Smith (Eds.), *Derivatives Handbook: Risk Management and Control* (pp. 265-309 (of 650)). New York, New York: John Wiley & Sons.
- Jorion, P. (2003). *Financial Risk Manager Handbook (Second Edition)*. Hoboken, New Jersey, United States of America: John Wiley & Sons, Inc.
- Jorion, P. (1995). Predicting Volatility in the Foreign Exchange Market. *Journal of Finance*, 50, 507-528.
- Jorion, P. (2001). *Value at Risk: The New Benchmark for Managing Financial Risk (Second Edition)*. New York, New York, United States of America: McGraw-Hill.
- Kaufmann, R., & Patie, P. (2003). *Strategic long-term financial risks - The one dimensional case*. ETH Zurich, RiskLab. Zurich: ETH Zurich.
- Kim, J., Malz, A. M., & Mina, J. (1999). *LongRun Technical Document* (First Edition ed.). New York, New York: RiskMetrics Group.
- Kupiec, P. H. (1999). Risk Capital and VaR. *Journal of Derivatives*, 7 (2), 41-52.
- Lahiri, S. N. (2003). *Resampling Methods for Dependent Data* (1. Edition ed.). New York, New York, U.S.A.: Springer.
- Lahiri, S. N. (1999). Theoretical Comparisons of Block Bootstrap Methods. *The Annals of Statistics*, 27 (1), 386-404.
- Laubsch, A. J., & Ulmer, A. (1999). *Risk Management - A Practical Guide*. RiskMetrics Group.
- Lee, A. Y., Kim, J., Malz, A. M., & Mina, J. (1999). *CorporateMetrics Technical Document* (First Edition ed.). New York, New York: RiskMetrics Group.
- Lee, B.-S. (1996). Comovements of earnings, dividends, and stock prices. *Journal of Empirical Finance*, 3, 327-346.
- Lintner, J. (1965). The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. (H. U. Government, Ed.) *Review of Economics and Statistics*, 47 (1), 13-37.
- Longerstae, J. (1996). *RiskMetrics - Technical Document* (Fourth Edition ed.). New York, New York: J.P.Morgan / Reuters (Morgan Guaranty Trust Company of New York).

- Luenberger, D. G. (1998). *Investment Science*. New York, New York: Oxford University Press Inc.
- MacDonald, R., & Marsh, I. W. (1997). On fundamentals and exchange rates: a Casselian perspective. *The Review of Economics and Statistics*, 655-664.
- MacDonald, R., & Taylor, M. P. (1994). The monetary model of the exchange rate: long-run relationships, short-run dynamics and how to beat a random walk. *Journal of International Money and Finance*, 13, 276-290.
- Malevergne, Y., & Sornette, D. (2006). *Extreme Financial Risks: From Dependence to Risk Management*. Berlin: Springer.
- Mark, N. C., & Choi, D.-Y. (1997). Real exchange-rate prediction over long horizons. *Journal of International Economics*, 43, 29-60.
- Markowitz, H. (1952). Portfolio Selection. *Journal of Finance*, 7 (1), 77-91.
- McNeil, A. J. (1999). Extreme Value Theory for Risk Managers. In P. i. Association, *Internal Modelling and CAD II: Qualifying and Quantifying Risk within a Financial Institution* (1st ed., pp. 93-113 (of 140)). London, UK: Risk Books.
- McNeil, A. J., & Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7, 271-300.
- Middleton, P. (1996). *GARP: Generally Accepted Risk Principles*. Coopers & Lybrand Global Financial Risk Management Practice. Leicester: Septre Litho.
- Mina, J., & Xiao, J. Y. (2001). *Return to RiskMetrics: The Evolution of a Standard* (First Edition ed.). New York: RiskMetrics Group.
- Mossin, J. (1966). Equilibrium in a Capital Asset Market. *Econometrica*, 34 (4), 768-783.
- Nelson, D. B. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica*, 59 (2), 347-370.
- Pindyck, R. S. (1993). The present value model of rational commodity pricing. *The Economic Journal*, 103, 511-530.
- Pindyck, R. S., & Rotemberg, J. J. (1990). The Present value model of rational commodity pricing. *The Economic Journal*, 100, 1173-1189.
- Ramsey, J. B. (1999). Regression over Timescale Decompositions: A Sampling Analysis of Distributional Properties. *Economic Systems Research*, 11 (2), 163-183.
- Ramsey, J. B. (2002). Wavelets in Economics and Finance: Past and Future. *Studies in Nonlinear Dynamics & Econometrics*, 6 (3), 1-29.
- Ramsey, J. B., & Zhang, Z. (1997). The analysis of foreign exchange data using waveform dictionaries. *Journal of Empirical Finance*, 4 (4), 341-372.

- Reilly, F. K., & Brown, K. C. (2003). *Investment Analysis & Portfolio Management (Seventh Edition)* (7th ed.). Mason, Ohio: Thompson - South-Western.
- Roy, A. D. (1952). Safety First and the Holding of Assets. *Econometrica*, 20 (3), 431-449.
- Schmid, F., & Trede, M. (2006). *Finanzmarktstatistik*. Heidelberg, Germany: Springer.
- Schwert, G. W. (1990). Stock Volatility and the Crash of '87. *Review of Financial Studies*, 3, 77-102.
- Schwert, G. W. (1989). Why Does Stock Market Volatility Change Over Time? *Journal of Finance*, 44 (5), 1115-1153.
- Schwert, G. W., & Seguin, P. J. (1990). Heteroskedasticity in Stock Returns. *Journal of Finance*, 45 (4), 1129-1155.
- Sharpe, W. F. (1964). Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. *Journal of Finance*, 19 (3), 425-442.
- Smithson, C., & Minton, L. (1996a). Value-at-Risk. *Risk*, 9 (1), 25-27.
- Smithson, C., & Minton, L. (1996b). Value-at-Risk (2). *Risk*, 9 (2), 38-39.
- Straetmans, S. T. (1998). *Extreme Financial Returns and their Comovements*. Tinbergen Institute Research Series, Ph.D. thesis no. 181, Erasmus University, Tinbergen Institute, Rotterdam.
- Stulz, R. M. (2000, June 27). Why risk management is not rocket science. *Financial Times*, p. 10.
- Tobin, J. (1958). Liquidity Preference as Behaviour Towards Risk. *Review of Economic Studies*, 67, 65-86.
- Treynor, J. L. (1999). Toward a Theory of Market Value of Risky Assets (1962). In R. A. Korajczyk, & R. A. Korajczyk (Ed.), *Asset Pricing and Portfolio Performance* (1st Edition ed., p. 320). London: Risk Books.
- Wagster, J. D. (1996). Impact of the 1988 Basle Accord on International Banks. *Journal of Finance*, 51 (4), 1321-1346.
- West, K., & Cho, D. (1995). The Predictive Ability of Several Models of Exchange Rate Volatility. *Journal of Econometrics*, 69, 367-391.
- Zakoian, J. M. (1994). Threshold Heteroskedastic Models. *Journal of Economic Dynamics and Control*, 18 (5), 931-955.
- Zumbach, G. (2007a). *The RiskMetrics 2006 methodology*. RiskMetrics Group. New York: RiskMetrics Group.
- Zumbach, G. (2007b). *A gentle introduction to the RM 2006 methodology*. RiskMetrics Group. New York: RiskMetrics Group.

Zumbach, G. (2007c). *Back testing risk methodologies from one day to one year*. RiskMetrics Group. New York: RiskMetrics Group.

Zumbach, G. (2006). Relevance of volatility forecasting in financial risk management. In M. K. Ong, & M. K. Ong (Ed.), *Risk Management: A Modern Perspective* (Vol. 1, pp. 583-606). Burlington, Massachusetts: Academic Press (Elsevier).

Zumbach, G. (2004). Volatility processes and volatility forecast with long memory. *Quantitative Finance*, 4 (1), 70-86.