

Extreme events in financial risk management

Citation for published version (APA):

Lehnert, T. (2002). *Extreme events in financial risk management*. [Doctoral Thesis, Maastricht University]. Universiteit Maastricht. <https://doi.org/10.26481/dis.20021220tl>

Document status and date:

Published: 01/01/2002

DOI:

[10.26481/dis.20021220tl](https://doi.org/10.26481/dis.20021220tl)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

EXTREME EVENTS IN FINANCIAL RISK MANAGEMENT

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de Universiteit Maastricht,
op gezag van de Rector Magnificus, Prof. Dr. A.C. Nieuwenhuijzen Kruseman,
volgens het besluit van het College van Decanen, in het openbaar te verdedigen
op vrijdag 20 december 2002 om 10.00 uur

door

Thorsten Lehnert

Promotores:

Prof. dr. M.W.M. Donders

Prof. dr. C.C.P. Wolff

Beoordelingscommissie:

Prof. dr. F.C. Palm (Voorzitter)

Prof. dr. S. Mittnik (Christian Albrechts University at Kiel)

Prof. dr. P.C. Schotman

To Natalie and Chiara

Acknowledgements

Extreme Events in Financial Risk Management is the result of four creative years of work at the Finance Department of Maastricht University. During the period of writing this book, I was lucky to have been surrounded by a team of people who all deserve a word of thank for their warm encouragement. The work described in this dissertation would never have been written in its present form without their contribution and support. I welcome this opportunity to thank all of them.

My largest debt by far is to my daily supervisor, Christian Wolff, who showed me how exciting the world of research can be. During the years he was creating the atmosphere where I could work independently, but whenever necessary he was very close and provided me with the best constructive criticism. Looking back, I realize that he always expressed his confidence in me, especially at the most difficult moments. I thank Christian for the personal interest he took in me. I have greatly benefited from his economic intuition and I have learned a lot from his personality.

A special word of thanks goes to my second supervisor, Monique Donders, for her detailed reviews of my chapters and the helpful suggestions for improvements. Despite her busy work at the bank, she still found the time to comment on every part of the thesis and this work has greatly benefited from her critique.

I would like to thank the members of my dissertation committee, Franz Palm (chairman), Stefan Mittnik and Peter Schotman for their valuable comments. I am also very much indebted to my “paranimfen”, Martin Fuchs and Dennis Bams. Martin is a very good friend and I always enjoyed the evenings when we went out for a beer. Dennis has provided considerable help, especially in the first year and I always benefited from the discussions with him. I extend these thanks for a wonderful time to all my other colleagues at Maastricht, especially the ladies from the secretariat and the members of the LIFE soccer team.

Ein besonderer Dank gilt auch meinen Eltern für ihre moralische und finanzielle Unterstützung. Ihr habt immer an mich geglaubt und mich bei all meinen Bemühungen unterstützt. Dies mag Euch nicht immer leicht gefallen sein.

Meine abschliessenden Worte des Dankes gehen an meine Frau Natalie und an meine Tochter Chiara, die in bewundernswerter Weise die schönen und weniger schönen Momente in den letzten Jahren mit mir gelebt haben. Ich danke Dir, Chiara, dass Du mir die Augen geöffnet hast für die wirklich wichtigen Dinge im Leben und für die vielen Momente der Freude, die Du uns bescherst. Ich danke Dir, Natalie, dass Du mich nach allen Kräften unterstützt hast und ich habe es Dir zu verdanken, dass ich so weit gekommen bin. Deine Energie und Begeisterung hat vieles bewegt und Du hast mir immer die Zuversicht gegeben, dass unser "Abenteuer" zu einem guten Ende kommt. Ich hoffe, dass ich einmal die Möglichkeit bekomme, dies alles auch für Dich zu tun. Ich danke Dir für Deine Liebe und Unterstützung.

I have very much enjoyed the time spent during my Ph.D. in Maastricht. It has allowed me to broaden my horizon...

Thorsten Lehnert

December 2002

Contents

CHAPTER 1

Introduction: Aims and Scope	1
1.1 Introduction	1
1.2 The Need for Risk Management	4
1.3 Measuring Market Risk with Value-at-Risk	5
1.4 Model Risk	8
1.5 The Regulatory Environment	11
1.6 Research Objectives	12

CHAPTER 2

Modeling Conditional Leptokurtosis and Skewness in Financial Time Series	17
2.1 Introduction	17
2.2 The Econometric Framework	19
2.3 Data	24
2.4 Empirical Results	25
2.5 Concluding Remarks	35

CHAPTER 3

Modeling Scale-Consistent Value-at-Risk with the Truncated Lévy Flight	37
3.1 Introduction	37
3.1.1 Modeling on different Time Scales	37
3.1.2 Physics meets Finance	38
3.2 Truncated Levy Flights	39
3.2.1 The Truncated Lévy Distribution	39
3.2.2 Scaling Properties	41
3.2.2 The Gaussian Process as a Special Case	43
3.3 Model Specifications	43

3.3.1	The Generalized Augmented GARCH Process	43
3.3.2	Density Forecasting and VaR	45
3.4	Empirical Results	47
3.4.1	In-Sample Analysis	47
3.4.2	Out-of-Sample Analysis	51
3.5	Conclusions	53

CHAPTER 4

An Evaluation Framework for Alternative VaR Models		55
4.1	Introduction	55
4.2	Econometric Framework	57
4.2.1	Standard Time Series Models	57
4.2.2	Modeling Extreme Events with Lévy-Stable Distributions	58
4.2.3	VaR and Parameter Uncertainty	59
4.2.4	Backtesting	61
4.3	Data	62
4.4	Results	62
4.4.1	Properties of the alternative VaR models	63
4.4.2	Uncertainty in the reported VaR	66
4.4.3	Out-of-Sample Violations	69
4.5	Concluding Remarks	75

CHAPTER 5

Implied Volatility Forecasting		77
5.1	Introduction	77
5.1.1	The Information content of Implied Volatilities	78
5.1.2	A 'New' Measure of Local Volatility	79
5.2	Model Specifications	80
5.2.1	The GARCH Option Pricing Model	80
5.2.2	Model Calibration	83
5.3	Data & Methodology	86
5.4	Empirical Results	90
5.4.1	In-Sample Valuation Errors	90
5.4.2	'Historical' and 'Implied' Volatility Forecasts	92
5.4.3	Granger Causality Test	93
5.4.4	Out-of-Sample Accuracy of the Volatility Forecasts	94
5.5	Concluding Remarks	96

CHAPTER 6

Explaining Volatility Smiles	99
6.1 Introduction	99
6.2 An Option Valuation Model	101
6.2.1 Option Pricing under GARCH	101
6.2.2 Conditional Leptokurtosis and Skewness	101
6.2.3 Local Risk Neutralization	102
6.2.4 Alternative Option Valuation Methods	105
6.3 Data & Methodology	106
6.3.1 DAX 30 Index Options	106
6.3.2 Methodology	106
6.4 Pricing Results	108
6.4.1 In-Sample Analysis	108
6.4.2 Out-of-Sample Valuation Errors	111
6.5 Conclusions	115

CHAPTER 7

Summary and Concluding Remarks	117
7.1 Introduction	117
7.2 Summary	118
7.3 Directions for further research	121

Bibliography	123
---------------------	------------

Nederlandse samenvatting	133
---------------------------------	------------

Curriculum Vitae	139
-------------------------	------------

CHAPTER 1

'In the end, there may be no absolute truth about financial values. ... When you propose a model of the physical world, you're pretending you can guess the structure God created. But when you propose a financial model, you're pretending you can guess another person's mind. ... It's not that physics is "better"; rather than finance seems to be harder.'

Introduction: Aims and Scope

1.1 Introduction

This is the overall conclusion with which Emanuel Derman, managing director in the Firmwide Risk department of Goldman, Sachs & Co., tries to explain why techniques of physics and applied mathematics work less well in finance. Derman, a well-known finance expert coming from particle physics, speculates about the imperfect science of financial modeling. In particle physics, tools like differential calculus, Fourier series and stochastic calculus for describing continuous motion are extremely successful, because once the dynamics are known, the parameter values are universal. In contrast in finance, while analyzing the dynamics of financial prices, we have to incorporate expectations about the future value of certain parameters to describe the true data generating process. Therefore our

theory is based on these expectations and the true price process depends on realizations of these unknown parameters¹. The major difference between physics and finance is probably that nature is predictable but the behavior of people isn't (Derman (2001)). Although the achievements of international financial economists are indisputable manifold, statistical physics still offers an enormous amount of statistical tools that are extremely interesting for financial modeling. Techniques from physics are extremely valuable for the analysis of extreme complex problems. A part of the dissertation deals with a stochastic process developed in the physics literature.

In general the dynamics of financial markets constitute a major challenge for financial econometricians and physicists alike. Financial markets are strongly affected by external news, which are mostly unexpected, both in date and nature. Using the statistical approach we try to decipher this randomness, find appropriate statistical tools for its description and develop a form of predictability. Analyzing the information content of historical observations we make certain assumptions about the frequency of possible price changes. These frequencies may reflect some mechanism of the markets themselves and will remain stable over time. Of course, there is an evolution in the way markets behave and therefore the statistical description of financial fluctuations will be imperfect. But nevertheless under the assumption that the behavior of financial markets changes slowly, the statistical description allows us to describe the near future and at least it is helpful to measure the risks involved in financial trading activities.

Recently have academics begun to study the risk management practices of financial institutions. In the light of high-profile failures of risk management (such as Barings Bank or Long Term Capital Management), corporations increasingly recognize that managing their exposure to financial risks is naturally one of their core businesses. A practical risk management tool for financial institutions nowadays is Value-at-Risk (VaR). The introduction of the RiskMetricsTM methodology by JP Morgan (1996) has enabled companies to compute a simple measure of market risk for given portfolio of assets. It is prescribed by the Basle Committee to report to the (international) supervisor and also used as an internal management tool, among other things, to see if traders are within their limits. VaR is an estimate of the size of the potential loss to be expected over a given period for a given probability. The new

¹ E.g. the value of an option depends on the expected volatility of the underlying asset over the lifetime of the option.

methodology has also aroused the interest of academics, because the measurement of VaR is a challenging statistical problem. Crucial for an accurate VaR estimate is the precise determination of the 'extreme events', which results from the distribution function of the underlying return series in the tails.

For modeling financial returns and determining the downside risk of financial positions exact knowledge of the frequency of extreme events is crucial. In the literature two different approaches can be distinguished: unconditional versus conditional modeling. In unconditional models, based on extreme value theory (EVT), the tail behavior of the return distribution is modeled over time assuming that the tail events are independent and uncorrelated and ignoring the conditional heteroskedasticity of most financial assets (see e.g. Christoffersen and Diebold (1998b)). Usually a huge data set is necessary to derive reasonable tail estimates, but the tail characteristics seem to be stable over time. Tail estimates are often used to approximate the downside risk exposure of financial assets. Embrechts et al. (1998) and Longin (1997) use parametric estimation techniques for the tails of the distribution in the i.i.d. case. Danielsson and de Vries (1997) use a semi-parametric method based on a tail estimator for the largest risks and an empirical distribution function for smaller risks. Mc Neil (1998) uses parametric estimation techniques, but shows how to correct for the clustering effects caused by stochastic volatility. Diebold et al. (1998a) discuss the potential of extreme value theory for estimating and forecasting extreme event probabilities.

In contrast to the unconditional approach, the class of GARCH models has been very successful in modeling the significant volatility clustering and non-i.i.d. properties of financial data (Bollerslev et al. (1994)). Various studies show the potential of ARCH-type models driven by fat tailed distribution for estimating the risks involved in financial trading. Pownall and Koedijk (1999) extend the standard RiskMetricsTM (1996) approach and allow for conditional leptokurtosis using a tail estimator to determine the degrees of freedom of the Student-t distribution. McNeil and Frey (2000) combine EVT and GARCH-modeling and confirm that the model outperforms existing methods. Mittnik et al. (2000a, 2000b) demonstrate that more general GARCH structures and skewed fat tailed distributions improve the precision in-sample and out-of-sample.

In particular, to calculate VaR for a certain time horizon we have to forecast the probability distribution of portfolio returns over this horizon. It often turns out to be a problem of having a proper volatility model. Since volatility is unobservable, the accurate

measurement and forecasting of volatility is crucial for successful risk management and analysis. The available models, GARCH or stochastic volatility, based on historical returns seems to work quite well in-sample but have limited potential for out-of-sample forecasting. In recent years there is some support for the hypothesis that more relevant information is provided by implied volatilities from daily option prices and there is less relevant information in volatilities constructed using historical returns (GARCH, stochastic volatility or realized volatility), see among others Blair, et al. (2001). If option markets are efficient, option prices contain information about the future price process of the underlying asset over the lifetime of the option. In particular for risk management purposes analyzing option prices means to decipher information about the expectation of market participants about the volatility of the underlying asset over the lifetime of the option. Therefore the volatility measure backed out from option prices is a forward-looking measure and extremely interesting for risk management purposes.

Modeling volatility and asymmetric tail fatness is not only relevant for the measurement of the downside risk potential of asset portfolios, but it is also important for option pricing. In addition to improving the prediction of volatility, conditional leptokurtosis and skewness can also partly explain strike price and maturity patterns across option values such as the pronounced smile or better skew in implied volatilities in the index option market. The size of violations is too extreme to be explained by market imperfections. The choice of a specific option pricing model has important implications in terms of the implementation of risk management strategies using options or the pricing and the implementation of risk management strategies for portfolios of options. Relying on a wrong pricing models and therefore the failure to hedge properly can expose an institution to extreme swings in portfolio values from unanticipated changes in the underlying risk factors. Recent experiences with losses at "arbitrage" hedge funds like Long Term Capital Management in 1998 show the crucial importance to many, including regulators who supervise financial institutions.

1.2 The Need for Risk Management

The stock market crash in 1987, the turbulences in emerging markets, starting in Mexico in 1995, continuing in Asia in 1997, and spreading to Russia and Latin America in 1998 and the extreme market movements more recently, starting with the burst of the 'dot-com'-bubble

March 2000 has extended the interest in risk management. Additionally, technological and regulatory changes, more globally operating corporations, more complex financial products and the increased volatility of financial markets nowadays provoked an escalating debate over financial risk management practice. The growing demand for more accurate risk measurement techniques evoked a huge body of academic research published on risk management techniques and derivative valuation.

In the past two decades, the types of risks affecting the activities of financial institutions have increased dramatically. Usually, financial risks are classified into categories of market risk, credit risk, liquidity risk, operational risk and legal and regulatory risk. Market risk is the risk that a change in the market prices or rates will affect the value of positions in assets and derivatives. Since the measurement of market risk is of crucial importance throughout the following chapters, various techniques to measure market risk are summarized in the following sections. A financial institution is exposed to credit risk, because the counterparty to e.g. a derivative contract may be unwilling or unable to fulfill their contractual obligations. A change in the credit quality of the counterparty also affects the value of a position. Liquidity risk arises from the inability to make necessary payments or to refinance obligations. Institutions are also exposed to various risks related to internal human or technical failures called operational risks. The appropriateness of valuation techniques for risk measures or derivative contracts also creates operational risk, the so-called model risk. Legal risk describes the type of risk of loss when the institution is unable to enforce a certain transaction in law. A change in the regulatory requirement is also associated with legal risk. In one of the following sections the regulatory framework will in particular be discussed.

The measurement and management of these different risks is crucial for financial institutions and various procedures have been recommended (Risk Standard Working Group (1996)). While there are certainly many sources of financial risk, we concentrate here narrowly on market risk associated with changes in the price or rates of assets and/or underlying traded instruments over a certain time horizon.

1.3 Measuring Market Risk with Value-at-Risk

Financial risk management relies heavily on quantitative methods to describe the dynamics of financial markets. In this thesis we are concerned with a quantitative analysis for optimal risk management concerning equity price risk and foreign exchange risk. Risk associated with

equities can be separated in general market risk, the sensitivity of a portfolio value to changes in the level of a broadly diversified stock market index, and specific or idiosyncratic risk, which refers to a part of the price volatility of a particular stock that is determined by firm-specific characteristics. Portfolio theory teaches us that the idiosyncratic risk can be diversified away; therefore we are mainly concerned with general market risk. Risk associated with foreign currencies arises, because of imperfect correlated movements of currency prices and fluctuations in interest rates. Mainly globally operating corporations are strongly exposed to foreign exchange risks, which result in some uncertainty about the return of investments. Risks arise naturally from an open (unhedged) position, but also from a hedged position, because of imperfect correlation with the instrument, which was intended to offset the position.

Market risk can be classified into two categories: directional and nondirectional risks. Linear approximations such as beta for the exposure of a single stock to market movements or delta for the exposure of options to the underlying asset are commonly used for directional risks. Nondirectional risks consist of nonlinear exposures and exposures to volatility changes. The quadratic exposure of options to the underlying asset is measured by gamma. Market risk also arises from movements in the level or in particular volatility of market prices, known as volatility risk. Naturally, minimizing market risk presupposes a quantification of market risk. Recent proposals for the disclosure of financial risk call for firm-wide measures of risk. A standard benchmark is the Value-at-Risk methodology. VaR is being used as a basis for capital charge for market risk. Given a confidence level of α percent and a time horizon of t days, the VaR of a portfolio is the loss that is expected to be exceeded with a probability of only $1-\alpha$ percent over the next t days. Intuitively, one expects, under normal market conditions, that a 95-percent one-day value-at-risk will be exceeded roughly once every month. Formally VaR describes the quantile of the projected distribution of gains and losses over a target horizon.

However, one should not treat one's measure of value-at-risk, even if accurate, as the level of capital necessary to sustain the firm's risk. VaR allows the comparison of different positions and is merely a benchmark for relative judgments, such as the risk of one desk relative to another, the risk of one portfolio relative to another, or the risk of one volatility environment relative to another, and so on. VaR actually gives no information about the maximum loss that can occur or the distribution of extreme losses within the lower quantile.

This is especially interesting in the case of extreme events where a single huge catastrophic loss might push the firm into bankruptcy. To account for the so-called non-subadditivity of VaR, a modification is often introduced. The expected shortfall or conditional VaR determines how big expected losses are in the $1-\alpha$ tail region that is not captured in the standard VaR measure.

Initially invented as a method to measure (aggregate) market risk, VaR is now being used to set traders certain limits for trading activities and as a common and comparable measure to relate different risky activities. For example assuming a Black-Scholes economy, Ahn et al. (1999) analyzed the problem of an institution optimally managing the market risk of a given exposure by minimizing its Value-at-Risk using options. A major problem arises if a risk manager uses a VaR system for risk control or performance evaluation, because traders try to “game” their VaR measure and evade their risk limits. For example using a standard VaR method for option portfolios, a short at-the-money straddle would actually report zero risk in terms of VaR, because the sensitivity to changes in the underlying of the short put and call positions would offset each other, but the actual downside risk is nearly unlimited. The downfall of Barings Bank was due to a single trader, Nick Leeson, who lost an enormous amount from this kind of derivative trading. The bankruptcy served as a wakeup for financial institutions all over the world (see Rawnsley (1995)).

Beside several shortcomings, the usefulness of VaR depends heavily on the accuracy of the risk measure. The greatest risks together with the use of VaR to measure market risk arise from a misspecification of financial price dynamics. Typically the riskiness of an investment is underestimated, because the misspecified VaR-model reports a Value-at-Risk that is too low. This underestimation can be a result of an underestimation of large events (in case of Gaussian models) or the representative historical period that was chosen doesn't project the current market condition. In practice VaR is often complemented with stress tests; an analysis of portfolio performance and risk in case of exceptional but not impossible changes in market prices.

The challenging statistical problem behind the Value-at-Risk approach and e.g. option pricing is similar. The VaR estimate is a forward-looking measure, because we make a statement about the loss potential of our portfolio value over a certain future time horizon. Equivalently, we determine option prices by assuming a certain process for the input parameters of our valuation model over the lifetime of the option contract. Therefore using the conditional approach we rely on a certain model to describe the historical dynamics of

financial prices (e.g. we derive an estimate for the current volatility of the asset) and we use a method to describe the future dynamics of financial prices (e.g. we have to forecast the volatility of the asset over a certain horizon). The better the description of the *actual* (future) price dynamics, the better the estimated VaR or the option price.

1.4 Model Risk

Financial markets in the past twenty years have seen an enormous emergence of theoretical approaches to valuation of derivative contracts or measurement of risk exposures. The growing derivatives trading volumes can be partly explained by the confidence of traders in using valuation and hedging techniques based on the Black-Scholes models and its extensions. The development of new approaches to risk management also heavily depends on the acceleration of available computing power. The increased volatility in financial markets also triggered the development of more complex financial products, whose valuation and trading was heavily model-driven and requires more sophisticated approaches. One important consequence is that the reliance on models to handle risk carries its own risk: *model risk*.

One major source of model risk is the risk that the stochastic process of the underlying asset is wrongly specified. It is a common approach to assume that prices are log-normally distributed and the stochastic component is time independent and uncorrelated². The Black-Scholes (1973) option pricing model also assumes log-normality. Several generalizations for the return process were recently proposed, but the normality of the stochastic component is still a standard assumption in many financial applications. However, empirically there are more extreme price movements than a lognormal distribution allows for. Additionally for equities extreme downward movements are more likely than extreme upward movements, which cannot be described by a symmetric return distribution. As a result, standard models consistently underestimate the potential risks involved in financial trading. The inaccuracy of a misspecified model cannot be determined, because the true distribution is not known and therefore the model risk characteristics are typically not known.

A very important feature of every risk management strategy or derivatives valuation method is that we have to incorporate expectations about future realizations of input parameters. When parameters are observable, we form a set of postulates about the evolution

² This is consistent with the informationally efficient markets and a major advantage of this assumption is that log-changes of prices are normally distributed and we can use the results from mathematics to study this processes.

of e.g. dividends or interest rates and incorporate these parameter into the model. But parameters like volatility of the underlying asset over the lifetime of the option contract or over the time horizon of the risks measure, like VaR, are not observable. Therefore we firstly have to visualize the parameter that cannot be directly observed and secondly we have to forecast the evolution of e.g. volatility over a certain period. This method leads naturally only to a 'best guess' and model risk characteristics are unknown.

Another important source of risk arises from risk management strategies using options. Under no arbitrage conditions, the market risk exposure of an open position is hedged by taking an offsetting position in options. The so-called options delta indicates how much of the open position is offset by the option's market risk exposure. Proper delta hedging requires that the pricing model is correct. Otherwise the hedge is inaccurate and cannot eliminate price risk, even when the hedge is rebalanced frequently to ensure delta neutrality. For example Bakshi et al. (2000) show that one-dimensional diffusion option pricing models cannot be completely consistent with observed option price dynamics, because empirically, call (put) prices are frequently not positively (negatively) correlated with the underlying. As a result continuous rebalancing may create risk instead of eliminating risk. Using a historical simulation approach, Green and Figlewski (1999) showed that these sources of model error combine to produce substantial risk exposure for an option writer, even when delta hedging of the position is actively followed.

Although all models need actually the future value of some input parameter, the estimate for the value is based on historical data. But history may not necessarily provide a good estimate of future value, because their historical values are unstable and vary strongly with the sample size, frequency or sampling period. Additionally, when a particular phenomenon is not present in the historical data, then the parameters of the model that take account of the particular phenomenon are estimated with a lot of uncertainty, reflected by the standard error of the parameter estimate. Therefore the uncertainty about the parameters, the *estimation risk*, will affect the predictive power of the particular model.

In the light of model risk, the risk of implementing a certain (standard) model should be an important concern of any user of risk management model. Marshall and Siegel (1996) found wide variations in Value-at-Risk results produced by different software packages even using the same model. Identical portfolios were given to leading risk management system' vendors and they were asked to obtain VaR estimates according to the standard RiskMetrics™ method. Additional variations were found related to increasing complexity of asset class.

Market risk and model risk are typically related. Model risk characteristics are typically unknown, because the true data generating process is unknown and therefore it necessarily creates additionally uncertainty about the market risk exposure of financial assets. Additionally model risk leads to an incorrectly specified estimation of market risk and therefore to inappropriate risk management strategies.

Every statistical model to measure the risk exposure of a portfolio of assets, to calculate the price of a derivative contract or to set up hedging strategies relies on certain assumptions about the dynamics of future prices. If these dynamics are incorrectly described, the approach, which was initially meant to manage or minimizes the risks involved, instead creates risks. Relying on a misspecified model makes the risk management tool less effective. Therefore understanding the basic techniques for risk management *and* the advantages and pitfalls of these techniques is of crucial importance.

A part of the dissertation is concerned with an empirical investigation of the highly negative correlation of volatility with changes in market prices. Ignoring this important relationship can involve considerable liquidity risk exposure for an institution writing options. For example it is common practice to trade on the term structure of volatility. The term structure of implied volatility inferred from option prices using the Black-Scholes formula may be upward sloping even when the current volatility level is not expected to increase. A trader using the Black-Scholes option pricing framework would expect that the long-term option are overpriced. He might incorrectly assume that he can make an arbitrage profit from selling and hedging long-term volatility-sensitive options³, but the Black-Scholes hedge does not account for asymmetric GARCH effect. When prices go down slightly, volatility goes up significantly, because volatility is highly negatively correlated with spot returns. The short options position loses considerable value leading to margin calls that have to be met in cash. Even if the institution uses standard delta hedging techniques, this can expose an institution to sudden swings in the values of their positions and to enormous liquidity problems. Long-Term Capital Management finally was a victim of liquidity risk, but it was also reported to have lost a huge amount from short options positions on equity volatility (see e.g. Dunbar (1999)).

³ A short long-term at-the-money straddle.

1.5 The Regulatory Environment

Not only financial institutions are scared by the possibility of financial distress, but also governments are aware of the risk that a single bank failure could lead to a growing instability of the markets and therefore to a crash of the whole banking sector (systemic risk).

The original purpose of the 1988 Bank for International (BIS) Accord was to impose minimum capital requirements for credit risk. In the 1996 Amendment, or “BIS 98”, the Accord was extended and market risk arising from the trading activities of financial institutions was incorporated (Basel Committee on Banking Supervision (1996a)). Market risk is the exposure of the value of a portfolio to changes in financial market prices. A capital requirement was imposed for all positions in the trading book, on as well as off-balance sheet traded instruments. The regulator defined a supervisory framework that also includes a kind of reality check for procedure to measure the risk exposure of trading portfolios. The so-called backtesting approach is a statistical method that consists of verifying over a certain horizon that actual losses of portfolio value are in line with expected losses (Basel Committee on Banking Supervision (1996a)).

The intention of the Amendment was to allow financial institutions within a given methodological framework to use internal models to calculate the capital charge for market risk. Initially the regulator proposed a standardized approach to first determine the capital charge for each risk category. All measures were simply summed up to obtain the global capital charge related to market risk. The approach ignored portfolio effects, because risks are partially offset by diversification across categories. Later on the regulator recognized the superiority of already implemented risk management models compared to the standardized approach. The “BIS 98” in particular urged financial institution to use internal models based on the Value-at-Risk methodology. The 10-day 99% VaR estimate of the financial institution is simply multiplied by 3 or higher, based on the sophistication of the internal model, to derive the regulatory capital charge. Actually the BIS allows the institution to use the 1-day VaR multiplied by the square root of 10 as a proxy for the 10-day VaR. The use of a multiplier and a scaling rule was frequently criticized by institutions and academics, but the regulator considered it as an insurance against various other risks or extreme market movements.

Recently the proposal for a new capital adequacy framework (Basel Committee on Banking Supervision (1999, 2001)) constitutes an even more comprehensive and integrated

approach to risks. The so-called Basel II Accord implements changes with respect to credit risk and operational risk. It can be seen as an improved standardized approach based on the institution's internal ratings, extended by a more sophisticated full modeling approach based on internal models and also applied to the banking book. Additionally the institution is reviewed frequently by the supervisor to ensure that the institution has enough capital to cover various risk types and to become familiar with the increasingly sophisticated techniques developed internally by the institution. The new capital adequacy framework also contains a "pillar 3" on market discipline, information disclosure and transparency. Therefore, the objective of the regulator might be to force institutions to develop an integrated approach to calculating, among others, market risk and credit risk, instead of simply adding a market VaR to credit VaR. This allows risk managers to generate not only economic capital, but also regulatory capital.

Although capital requirements are a less than sophisticated approach, especially in the way they were implemented by the regulator, the Amendment has forced academics and practitioners to rethink common risk management practices. The intention was that internal models to measure market risk exposure become fully integrated in the daily risk management of the institution. However the supervisor should be aware of the fact that changes in regulatory requirements can also be seen as a potential risk factor and create additional risk for financial institutions and financial markets in general.

1.6 Research Objectives

Previous research suggests that model risk in trading and hedging asset and derivative portfolios is quite large⁴. But still methods to measure risk or delta hedging techniques based on valuation models implemented with forecasted volatilities are the only viable trading and risk management strategy for most financial institutions. Of course, this can be expected to involve considerable risk exposure, but the task is to find a statistical description to minimize risk exposures in terms of model risks. Therefore there are two important components of an overall risk management strategy: the use of the best model to describe the dynamics of financial prices and as a result derivative pricing model and the best volatility estimators.

⁴ See e.g. Green and Figlewski (1999).

In this dissertation we concentrate on the measurement and management of market specific risks and extend previous research by improving the description of the conditional behavior of asset returns. We demonstrate the necessity to account for the observed and expected extreme price fluctuations in financial markets and we propose several statistical methods to capture and treat the fat tails in heteroskedastic financial time series. The analysis focuses on single index portfolios or on FX-rates but can be extended to the modeling of asset portfolios in a multivariate setting. We show the great relevance for practical risk management.

This thesis consists of 5 chapters, which focus on different aspects of modeling extreme events and implications for financial risk management. This section briefly describes the content of each of the chapters.

Chapter 2 presents an overview of the theoretical aspects of volatility modeling with conditional leptokurtosis and skewness. We set up an econometric framework for the Value-at-Risk analyses in the following chapters. Volatility is a measure for the degree of dispersion of possible outcomes of a random variable. Traditionally in finance, a normally distributed random variable is considered and the volatility is presented by its standard deviation. The approach has the advantage that we can use standard calculus tools to describe the dynamics of this random variable, for example under the absence of serial correlation, we can use a scaling rule to annualize a standard deviation measured on a daily basis. But returns in financial assets show consistent excess kurtosis, indicating the presence of large fluctuations not predicted by Gaussian models. Assuming a random variable that is distributed according to an asymmetric leptokurtic distribution, volatility is presented by the serially correlated scale parameter of the distribution. We propose a model with time varying scale parameter (GARCH process) with error terms that are distributed according to a skewed leptokurtic distribution. We determine the appropriate GARCH specification for each index series by conducting a specification test based on a generalization of the augmented GARCH process of Duan (1997).

In Chapter 3 we extend the analysis of Chapter 2 and apply it to the analysis of the downside-risk exposure of major index portfolios. Additionally a stochastic process from statistical physics is investigated and applied to financial data. Mandelbrot (1963) first proposed the idea that price changes are distributed according to a Lévy stable law. This

model was frequently criticized, because the tails are now much overestimated and the infinite variance makes it impossible to apply the Central Limit Theorem. In recent studies the truncated Lévy flight has been shown to eliminate these problems and to be very promising for the modeling of financial dynamics. An exponential decay in the tails ensures that all moments are finite and the distribution is fat-tailed for short time scales and it converges to a Gaussian process for increasing time scales. In this chapter we propose a generalization to the popular RiskMetrics™ approach to Value-at Risk. In order to model scale-consistent Value-at Risk (VaR), we propose a model with a time varying scale parameter, with error terms that are truncated Lévy distributed. Lévy flights include a method for scaling up from a single-day volatility to a multi-day volatility. We use this rule to forecast the multiperiod return density, estimate Value-at-Risk (VaR) several days ahead, and compare it to the RiskMetrics™ approach (the exponentially weighted moving average model with the square-root-of-time scaling rule for the volatility), which is a special case of our method. Both the in-sample and the out-of-sample characteristics of our model are quite promising.

In Chapter 4 we take a look at the ability of different models to produce useful Value-at-Risk estimates for exchange rate positions. It's evident that the statistical properties of FX-rates differ from those of equity returns and we restrict our analysis on simpler specifications for the volatility process. We make a distinction between models that include sophisticated tail properties and models that do not. The former type of models often leads to too extreme VaR estimates, whereas the latter type underestimates the risk in case of extreme events. Our analysis shows that it is important to take into account parameter uncertainty, since this leads to uncertainty in the reported VaR. We make this uncertainty in the VaR explicit by means of simulation. Our empirical results suggest that more sophisticated tail-modeling approaches come at the cost of more uncertainty about the VaR estimate itself. In the case of the GARCH(1,1)-Student-t model the average VaR may be adjusted for parameter uncertainty to arrive at levels which are adequate according to out-of-sample tests.

Chapter 5 empirically investigates a method to quantify volatility using the information content of index options. We derive the parameters of a EGARCH option pricing model from the term structure of the observed market smile of DAX 30 index. We find the EGARCH option pricing model performs well in determining the shape of the volatility smile for different maturities in the period of January 2000 to August 2001. Based on the implied EGARCH methodology we use the information in option prices to derive a theoretically

sound 'new' measure for local volatility and analyze how well it explains and forecasts actual realized volatility. The daily realized volatility measure is constructed with 5-minute interval transaction prices in the DAX 30 future. The local volatility measure explains a large part of realized volatility and performs considerably better in one day ahead volatility forecasting than conventional time-series models.

Chapter 6 moves to the effect of sophisticated tail properties in option pricing. We investigate how well the observed market smile can be explained by using the generalized GARCH option valuation model of Duan (1999) with a skewed fat-tailed distribution. We analyze the pricing performance in- and out-of-sample using index options. We compare our model with the closed-form GARCH option pricing model of Heston and Nandi (2000) and the ad-hoc Black-Scholes model of Dumas, Fleming and Whaley (1998). Previous research suggests that models with correlation of volatility with spot returns and the path dependence in volatility together outperform deterministic methods. Additionally to these previous findings we investigate that conditional leptokurtosis and skewness, asymmetry in the volatility process and the asset risk premium together determine the shape of the volatility smile. The improvements in the pricing performance are particularly remarkable for out-of-the-money options.

Chapter 7 concludes the thesis and provides a summary.

CHAPTER 2

Modeling Conditional Leptokurtosis and Skewness in Financial Time Series⁵

2.1 Introduction

In this chapter we provide an econometric framework for the Value-at-Risk analysis in the following part of the dissertation. We consider the price dynamics of a portfolio of stocks traded in financial markets. It is a well known stylized fact that index returns, like the returns on many other financial assets, are non-normal and tend to have fat-tailed distributions. Mandelbrot (1963) first proposed the idea that price changes are distributed according to a distribution with considerably fatter tails than the normal. But at the same time he observed that there is a temporal dependence in asset return, as he wrote

“ ... large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes ... “.

Nowadays it is common belief that tail fatness and volatility clustering are related. Volatility is usually measured as the scale parameter of the distribution; in the case of the normal distribution, it is the standard deviation. Therefore a time varying volatility, a leptokurtic

⁵ This chapter is based on: T. Lehnert, C.C.P. Wolff, 2001, ‘Modeling Scale-consistent Value-at-Risk with the Truncated Levy Flight’, LIFE Working Paper, Maastricht University.

conditional density or both can be an explanation for the excess kurtosis in financial time series.

The class of GARCH models introduced by Bollerslev (1986) has been very successful in modeling the significant volatility clustering and non-i.i.d. nature of the data. Unfortunately, a time varying volatility can often not account for all the excess kurtosis in financial return data. Even in a conditional framework the normality assumption cannot be justified and a quasi-likelihood methodology is typically used to derive consistent parameter estimates. As Blair et al (2001) wrote,

“The parameters are estimated by the usual quasi-likelihood methodology, so that the likelihood function is defined by assuming standardized returns, z_t , have normal distributions even though this assumption is known to be false.”

Additionally to the strong inconsistency with normality, several studies have shown that the standardized residuals from GARCH models also exhibit skewness. Extreme negative market movements can be observed more frequently than extreme positive price changes resulting in a fatter left tail of the return distribution. In addition to the earlier mentioned temporal dependence, financial price fluctuations tend to be negatively correlated with changes in volatility. Black (1976) first mentioned the so-called “leverage effect”. When the firm value falls, a company with debt and equity outstanding becomes more leveraged and as a result equity return volatility increases. Black already argued that the leverage effect alone cannot explain the negative correlation of volatility with spot returns. The work of Black found empirical support by Christie (1982) and Schwert (1989). Therefore an asymmetry in the volatility process, a skewed conditional density or both can be an explanation for the skewness in financial time series.

We will show that an asymmetric time varying volatility process is not enough to capture the excess kurtosis and skewness in the data and that a combination of a skewed fat tailed distribution for the innovations and asymmetry in the volatility process leads to a much better description of the data. Since the use of conditional Student-t-distribution has been very popular in the literature, we consider an asymmetric alternative proposed by Paoletta (1997). For this kind of distribution with power-law decay, the fatness of the tails provides an estimate for the maximally existing moments of the innovation process. Since we are modeling the continuously compounded rate of returns, the conditional expected rate of return is infinite, as the moment generating function of a (asymmetric) Student-t-distributed random

variable does not exist (except for the special case of a normal distribution, where the degrees of freedom are infinite). Because an infinite variance is sometimes a possible solution of the model, we do not model the first and second moment of the return process directly, but the location and scale parameter of the distribution. Therefore, a time varying scale parameter can be interpreted as a time varying volatility in the model. The modeling of the location and scale parameter of the distribution has an important implication, because the evidence is very limited as to which GARCH specification should be used. We cannot necessarily rely on the huge empirical evidence from 'traditional' approaches, describing the first and second moment of the return process. In order to derive an appropriate specification, we develop a specification test based on a generalization of the augmented GARCH process of Duan (1997), which allows us to reject several specifications among the models analyzed.

The outline of this chapter is as follows. Section 2.2 defines the econometric framework for the study and the data is described in Section 2.3. The empirical results are summarized in Section 2.4. Finally, Section 2.5 concludes.

2.2 The Econometric Framework

Volatilities can be described reasonably successful with a parametric model such as GARCH. Traditional GARCH models (with normal-distributed error terms) were designed to capture clustering of large and small innovations, which can be modeled as serially correlated conditional variances when the variance exists (Bollerslev et al. (1992)). In the case of the standard Bollerslev GARCH (p,q) model (Bollerslev (1986)) the conditional variance σ_t^2 can be approximated by

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^p \alpha_i \sigma_{t-i}^2 + \sum_{j=1}^q \beta_j \sigma_{t-j}^2 (\varepsilon_{t-j})^2, \quad (2.1)$$

where ε_t is a realization from a mean zero, variance one density. The analogue of the standard deviation in the family of Student-t distributions is a so-called scale parameter. If we replace the standard deviation by the scale parameter, we allow the scale parameter to be serially correlated, which produces the volatility clustering.

A more general class of models, which is frequently used together with distributions with power-law decay in the tails (see Mittnik et al. (2000a,2000b)) is the power-GARCH process of Ding, Granger and Engle (1993):

$$\sigma_t^\delta = \alpha_0 + \sum_{i=1}^p \alpha_i \sigma_{t-i}^\delta + \sum_{j=1}^q \beta_j \sigma_{t-j}^\delta (|\varepsilon_{t-j}| - \rho \varepsilon_{t-j})^\delta. \quad (2.2)$$

Equation (2.2) is well defined for parameter values for ρ between -1 and 1 . Bollerslev's GARCH model (Bollerslev (1986)) correspond to $\delta=2$ and $\rho=0$. The exponent δ is a parameter of the model and is estimated during the optimization routine. This model specification for $p=q=1$ leads to a different covariance stationary condition. In our notation the condition becomes

$$\alpha_1 + E(|\varepsilon_t| - \rho \varepsilon_t)^\delta \beta_1 < 1. \quad (2.3)$$

Mittnik et al. (2000a) derived a closed form expression for a power-GARCH process driven by an asymmetric Student-t distribution. The motivation for a model with a power transformation of the conditional standard deviation process and the asymmetric absolute residuals lies in the observed autocorrelation functions of $|\varepsilon_t|^\delta$, see Ding, Granger and Engle (1993). The parameter δ can be interpreted as a separate heteroskedasticity parameter (He and Teräsvirta (1999)).

Another feature of financial data is the asymmetric response of volatility to positive and negative shocks, well known as the leverage effect and empirical studies have shown that it is crucial to include leverage parameters (Nelson (1991), Ding, Granger and Engle (1993) and Engle and Ng (1993)). In Equation (2.2) and (2.3) for example, a positive parameter ρ ensures that the impact on volatility is greater for negative shocks compared to positive shocks. Engle and Ng (1993) and Hentschel (1995) studied the asymmetry in terms of a modification of the news impact curve. The news impact on volatility can be modeled either with a rotation parameter (like in the power-GARCH) or with a shift parameter (like in the NGARCH). Also a combination of both is possible, but then they might offset each other. In most empirical papers one or the other asymmetric GARCH specification is used without testing if a shift or a rotation is appropriate.

We use a specification test based on a generalization of the augmented GARCH process of Duan (1997) and determine the type of asymmetry for each data set and each innovations distribution. Additionally there is no evidence that the appropriate GARCH specification for Gaussian innovations is necessarily the appropriate specification for skewed fat tailed innovation. In particular the asymmetric parameter is assumed to capture the skewness in the data if the innovations distribution is symmetric. Since we are dealing with skewed distributions the relationship between asymmetry in the volatility process and skewness for the innovations distribution is crucial. For example allowing for skewness in the innovations distribution of the power GARCH process might lead to a smaller leverage effect and the effect might be different for different methods to model the asymmetry. Therefore, we have to test which specification is appropriate for a particular data set and a particular distribution for the innovations.

Duan (1997) developed the augmented GARCH process, a family of parametric GARCH models containing most of the existing GARCH specifications. In the following we consider the single lag ($p=q=1$) version of the volatility process, which is typically sufficient in practice⁶. The complete generalized augmented GARCH (1,1) model reads:

$$\begin{aligned}
 r_t &= \mu_t + \sigma_t \varepsilon_t, & \varepsilon_t &\sim D(0,1) \\
 \phi_t &= \alpha_0 + \gamma_{1,t-1} \phi_{t-1} + \gamma_{2,t-1} \\
 \sigma_t &= \begin{cases} |\delta \phi_t - \delta + 1|^{1/2\delta} & \text{if } \delta \neq 0 \\ \sqrt{\exp(\phi_t - 1)} & \text{if } \delta = 0 \end{cases} \\
 \gamma_{1,t} &= \alpha_1 + \alpha_2 |\varepsilon_t - b|^\kappa + \alpha_3 \max(0, b - \varepsilon_t)^\kappa \\
 \gamma_{2,t} &= \alpha_4 \frac{|\varepsilon_t - b|^\kappa - 1}{\kappa} + \alpha_5 \frac{\max(0, b - \varepsilon_t)^\kappa - 1}{\kappa},
 \end{aligned} \tag{2.4}$$

where the conditional location parameter μ_t can be specified additionally and the conditional scale parameter σ_t is assumed to vary over time. Our model is not restricted to zero-mean and variance-one distributions, it also allows for location-zero and scale-one continuous distributions $D(0,1)$. The finite stationary scale parameter and the sufficient condition for strict stationarity of the location adjusted augmented GARCH(1,1) process can be expressed in a general condition. The stationary scale parameter is equal to

⁶ See e.g. Duan (1997), McNeil and Frey (2000), Mittnik et al. (2000a, 2000b) and Mittnik et al. (2002).

$$\bar{\phi} = \frac{\alpha_0 + \alpha_4 E \left[\frac{|\varepsilon_t - b|^\kappa - 1}{\kappa} \right] + \alpha_5 E \left[\frac{\max(0, b - \varepsilon_t)^\kappa - 1}{\kappa} \right]}{1 - \alpha_1 - \alpha_2 E[|\varepsilon_t - b|^\kappa] - \alpha_3 E[\max(0, b - \varepsilon_t)^\kappa]} \quad (2.5)$$

$$\bar{\sigma} = \begin{cases} |\delta \bar{\phi} - \delta + 1|^{1/2\delta} & \text{if } \delta \neq 0 \\ \sqrt{\exp(\bar{\phi} - 1)} & \text{if } \delta = 0 \end{cases}$$

Therefore the location-adjusted augmented GARCH(1,1) process is strictly stationary if

$$\alpha_1 + \alpha_2 E[|\varepsilon_t - b|^\kappa] + \alpha_3 E[\max(0, b - \varepsilon_t)^\kappa] \leq 1 \quad \text{if } \delta \neq 0$$

$$\alpha_1 < 1 \quad \text{if } \delta = 0 \quad (2.6)$$

Most GARCH models can be interpreted as special case of the augmented GARCH process. For some specifications there exists a closed form solution of the unconditional scale parameter or the stationary condition, but otherwise they have to be evaluated numerically.

All models are estimated with maximum likelihood. The log-likelihood for a series of observations is equal to the sum of the conditional log-likelihood of each observation in the sample. The log-likelihood is given by

$$\ln L = \sum_t (\ln D(\varepsilon_t, 0, 1) - \ln(\sigma_t)) \quad (2.7)$$

$D(\varepsilon_t, 0, 1)$ is a standardized density function and the term $-\ln(\sigma_t)$ results from taking the log of the Jacobian of the transformation. Maximum likelihood estimators are assumed to follow the standard theory, so they are consistent and asymptotically normal with mean the parameter estimates and a variance-covariance matrix (DuMouchel (1973)). This is typically also being assumed for the (augmented) GARCH process, however a formal prove was not provided yet.

The augmented GARCH specification can be used to construct a specification test. A Lagrange multiplier (LM) test provides an extremely useful class of diagnostic tests. The LM test deals with local alternatives and with the augmented GARCH as a general alternative, it can be used to determine which GARCH specification can be rejected (see Duan (1997) for details). Let θ denote the vector of unknown parameters and let $\ln L(\theta)$ denotes the

associated log-likelihood. In order to test a hypothesis about the Maximum Likelihood Estimate (MLE) θ , we only need to find the restricted MLE $\tilde{\theta}$ and not the unrestricted one $\hat{\theta}$. Therefore $\tilde{\theta}$ is an estimate of θ that maximizes the log likelihood subject to a set of m restrictions on θ . Let $h(\tilde{\theta})$ denote the vector of derivatives evaluated at the restricted estimate $\tilde{\theta}$ given the history of observations. The null hypothesis that the imposed restrictions are true can be tested using the following LM statistic

$$T^{-1} \left[\sum_{t=1}^T h(\tilde{\theta}) \right]' \zeta^{-1} \left[\sum_{t=1}^T h(\tilde{\theta}) \right], \quad (2.8)$$

where T is the sample size and ζ is the information matrix evaluated at the restricted estimate $\tilde{\theta}$, which can be an estimate based on the outer product, the second derivative or as a robust estimate proposed by White (1982). All first and second derivatives are obtained by using numerical methods. The LM statistic is approximately $\chi^2(m)$ distributed if the null hypothesis is true. The LM test provides a particular useful framework for testing the adequacy of an estimated GARCH model, but the approach is not without problems. Lundbergh and Teräsvirta (2002) showed that under particular specifications of their model, a 'smooth transition alternative pose an identification problem'. One example could be testing the hypothesis of constant variance against the GARCH (1,1). Assume the single lag version of the model given in Equation (2.1). If α_1 is set equal to zero, it means that under the null of constant conditional variance, β_1 is an unidentified parameter, and as a result standard asymptotic inference is not applicable for testing the null hypothesis $\alpha_1 = \beta_1 = 0$ implied by the model. However this problem of underidentification does not affect our results. Neither we test the null of no ARCH in the standardized errors, nor for the specifications considered in our analysis, the problem of an unidentified nuisance parameter does influence the results of the specification test.

In the following part we compare, in-sample, two different specifications for the innovations distribution: the Gaussian and a skewed Student-t distribution called t_3 . The GARCH model with normally distributed innovations is nested into the case with time varying volatility and skewed t_3 distributed error terms.

The Student t distribution has often been used to account for the excess kurtosis in asset returns, but it cannot model the skewness. Paoletta (1997) proposed the so-called t_3

distribution, an alternative to the t distribution with skewness. Mittnik et al. (2000) have shown that the skewed distribution outperforms the symmetric distribution in in-sample fit and out-of-sample density predictions.

The standardized density can be written as:

$$D_{t_3}(u_t, 0, 1, v, d, \beta) = \frac{1}{K} \begin{cases} \left[1 + \left(\frac{-u_t^d \beta^d}{v} \right) \right]^{-\left(\frac{1+v}{d}\right)}, & \text{if } u_t < 0 \\ \left[1 + \left(\frac{u_t^d \beta^{-d}}{v} \right) \right]^{-\left(\frac{1+v}{d}\right)}, & \text{if } u_t \geq 0 \end{cases}, \quad (2.9)$$

where $K = (\beta + \beta^{-1})d^{-1}v^{1/d}B(d^{-1}, v)$ and $B(.,.)$ is the beta function. Since the t_3 distribution nests among others the Student-t it can be seen as a generalization of the Student-t distribution. The distribution has power-law tails; therefore vd gives an upper limit for the number of moments that exist for the distribution. The distribution is skewed to the right (left) for $\beta > 1$ ($\beta < 1$). It is also possible to reparametrize the skewed Student-t distribution and express the density in terms of the mean and the variance (see Laurent (2002)).

2.3 Data

In this study we use daily closing prices for some major stock market price indices between May 1992 and April 2000. In particular we examine the S&P500, NASDAQ and FTSE 100 from May 4, 1992, to April 3, 2000. The total number of trading days covered by the data is 2000 (FTSE 100) and 2001 (S&P500 and NASDAQ). The data are obtained from DataStream. We used the percentage daily logarithmic change $100 \cdot \ln(p_t / p_{t-1})$, where p_t is the price index at time t . The calibrations are executed with the well-known BHHH procedure programmed in the statistical software package GAUSS.

The summary statistics in Table 2.1 show that there is skewness and excess kurtosis in the data. The S&P 500, NASDAQ and FTSE 100 log-returns show excess kurtosis related to fatter tails compared to the normal and negative skewness related to a fatter left tail for the return distribution.

Table 2.1: Summary Statistics

Parameter	S&P 500	NASDAQ	FTSE 100
Mean	0.064	0.098	0.044
Standard Dev.	0.907	1.228	0.942
Minimum	-7.113	-8.954	-4.140
Maximum	4.989	5.848	4.345
Skewness	-0.398	-0.676	-0.055
Kurtosis	9.224	7.942	4.748
# of Observations	2001	2001	2000

Notes. Mean, standard deviation, skewness, kurtosis, minimum, maximum and number of observation for the log-index returns for the period May 4, 1992 until April 3, 2000.

2.4 Empirical Results

There is already some evidence that the inclusion of a time-varying asymmetric volatility process (e.g. GARCH) can capture some of the excess kurtosis and skewness in financial data, but it is often found that these models do not adequately account for leptokurtosis, since the residuals display frequently excess kurtosis. As a result several other fat-tail distributions (e.g. generalized error or Student-t) have been employed. But it is still an open question if the asymmetry in the volatility process adequately accounts for skewness. Hentschel (1995) studied extensively the news impact curves of various GARCH models and performed a specification test based on the likelihood ratio. In the standard GARCH process the news impact curve has the form of a parabola with a minimum at zero, therefore negative and positive excess returns have the same impact on volatility. If we shift the news impact curve, excess returns with different signs become different impact on volatility. The above-mentioned negative correlation between returns and volatility can be modeled with a shift of the news impact curve to the right or a clockwise rotation of the news impact curve. Intuitively one might expect that if market participants start selling their stocks and the index drops by a certain amount over the day, the uncertainty in the market in terms of volatility would increase and every forecast about possible future losses should be adjusted upwards. The increase is higher for extreme downward movements compared to moderate downward movements. Therefore both ways of modeling the asymmetry capture this intuition. But for zero or positive returns the economic interpretation of both modifications of the news impact

curve is different. The above mentioned shift of the news impact curve suggests that a zero return over the day still has a positive effect on the volatility, a small positive shock has a moderate or no impact and an extreme positive return shock has a huge impact on the uncertainty in the market. These effects are probably very difficult to justify economically. On the other hand a clockwise rotation of the news impact curve ensures that a zero return has still no impact and positive returns have reasonable positive impact on the volatility; in some situation the impact can even be negative. Therefore, intuitively the model predicts that rising stock prices on average do not necessarily influence the volatility in the market, an implicit assumption of the model that is probably less questionable than the predictions underlying asymmetric GARCH models with a shift of the news impact curve. As a result different modifications of the news impact curve in order to describe the positive impact of negative shocks on volatility also influence the volatility dynamics in general and therefore the quality of the VaR estimates. Already Hentschel (1995) showed that the differences in the conditional volatility estimates could be substantial among the various specifications. In order to compare different models, we have chosen several specifications of the generalized augmented GARCH process, in particular the standard GARCH model of Bollerslev (1986), the power-GARCH process of Ding et al. (1993) with a rotation of the news impact curve, the NGARCH model of Engle and Ng (1993) with a shift, the EGARCH model of Nelson (1991) with a rotation and a power-GARCH process with a shift of the news impact curve. The parameter estimates of the Gaussian model and a model with skewness and conditional leptokurtosis are reported in Tables 2.2 to 2.4. For the estimation the conditional mean μ_t is set equal to the unconditional value. We calibrated the models by restricting the unconditional scale parameter equal to the sample scale parameter (e.g. in the case of the Gaussian model we used the sample variance) and predetermined α_0 . This has a stabilizing effect for the calibration of the various models, in particular for the backtesting exercise. The location parameter is assumed to be equal to the unconditional value. The results show that our choices of model specifications are supported by the data. The significant positive parameters b (NGARCH and power GARCH with shift), α_3 (power GARCH with rotation), α_4, α_5 (EGARCH) and the improved log-likelihood compared to the standard GARCH show generally that there is a leverage effect in index returns. But in particular if we compare the Gaussian and t_3 specifications, the parameters capturing the leverage effect are different and for the t_3 distribution skewness can also be estimated. In general the skewness parameter is reduced if we add an asymmetry parameter to the model. The comparison of the linear GARCH and the

asymmetric GARCH models for the t_3 distribution shows that the asymmetric volatility models filter out some skewness, but there is still skewness in the residuals. For example, if we allow for skewness in the innovation distribution, the shift parameter capturing the leverage effect is reduced compared to the Gaussian model. Modeling the asymmetry with the shift parameter is technically more straightforward, because one parameter can be related directly to the leverage effect. It is more difficult to relate the effects directly to the two rotation parameters, which also interact with each other. Therefore for the rotation parameter we cannot directly observe the 'interaction' between the skewness parameter of the innovations distribution and the parameter measuring the asymmetry in the volatility process. However, we can conclude that besides the well-known leverage effect, skewness in the stochastic component can additionally explain the asymmetry in the volatility process.

Table 2.2: Estimates of the scale parameter equation, S&P 500 returns

Distribution	Parameter	GARCH	PGARCH (rotation)	NGARCH	EGARCH	PGARCH (shift)
<i>Normal</i>	δ	1	0.755	1	0	0.902
	(SE)	-	(0.050)	-	-	(0.056)
	(SE)	-	(0.100)	-	-	(0.108)
	κ	2	1.511	2	1	1.804
	(SE)	-	-	-	-	-
	(SE)	-	-	-	-	-
	α_1	0.940	0.907	0.898	0.984	0.888
	(SE)	(0.006)	(0.011)	(0.008)	(0.002)	(0.010)
	(SE)	(0.020)	(0.032)	(0.042)	(0.008)	(0.044)
	α_2	0.054	0.026	0.064	0	0.083
	(SE)	(0.005)	(0.016)	(0.007)	-	(0.013)
	(SE)	(0.017)	(0.018)	(0.021)	-	(0.037)
	α_3	0	0.149	0	0	0
	(SE)	-	(0.021)	-	-	-
	(SE)	-	(0.055)	-	-	-
	α_4	0	0	0	0.055	0
	(SE)	-	-	-	(0.019)	-
	(SE)	-	-	-	(0.018)	-
	α_5	0	0	0	0.162	0
(SE)	-	-	-	(0.018)	-	
(SE)	-	-	-	(0.055)	-	
b	0	0	0.660	0	0.610	
(SE)	-	-	(0.126)	-	(0.115)	
(SE)	-	-	(0.118)	-	(0.114)	
	Log-Likelihood	-2376.8	-2354.5	-2359.5	-2354.6	-2358.7
<i>Student-t_3</i>	δ	1	0.663	1	0	0.842
	(SE)	-	(0.100)	-	-	(0.105)
	(SE)	-	(0.153)	-	-	(0.153)
	κ	2	1.325	2	1	1.685
	(SE)	-	-	-	-	-
	(SE)	-	-	-	-	-
	α_1	0.953	0.917	0.927	0.991	0.913
	(SE)	(0.008)	(0.019)	(0.010)	(0.003)	(0.017)
	(SE)	(0.015)	(0.025)	(0.025)	(0.006)	(0.030)
	α_2	0.053	0.038	0.062	0	0.092
	(SE)	(0.010)	(0.030)	(0.011)	-	(0.027)
	(SE)	(0.016)	(0.028)	(0.015)	-	(0.038)
	α_3	0	0.157	0	0	0
	(SE)	-	(0.039)	-	-	-
	(SE)	-	(0.049)	-	-	-
	α_4	0	0	0	0.057	0
	(SE)	-	-	-	(0.029)	-
	(SE)	-	-	-	(0.018)	-
	α_5	0	0	0	0.154	0
	(SE)	-	-	-	(0.032)	-
	(SE)	-	-	-	(0.053)	-
b	0	0	0.532	0	0.457	
(SE)	-	-	(0.172)	-	(0.148)	
(SE)	-	-	(0.134)	-	(0.145)	
v	5.566	5.343	6.031	5.263	5.866	
(SE)	(0.171)	(2.447)	(2.637)	(2.435)	(2.677)	
(SE)	(0.178)	(2.534)	(3.168)	(2.595)	(3.093)	
d	1.715	1.733	1.702	1.727	1.709	
(SE)	(0.026)	(0.179)	(0.166)	(0.179)	(0.172)	
(SE)	(0.023)	(0.170)	(0.184)	(0.177)	(0.183)	
β	0.975	1.008	1.003	0.999	1	
(SE)	(0.025)	(0.027)	(0.027)	(0.027)	(0.027)	
(SE)	(0.026)	(0.018)	(0.023)	(0.022)	(0.025)	
	Log-Likelihood	-2316.6	-2303.5	-2309.4	-2304.1	-2308.6

Notes. The Table reports parameter estimates for the Normal and Student- t_3 model and the various GARCH specification. Standard errors and robust standard errors proposed by White (1982) are given within parentheses.

Table 2.3: Estimates of the scale parameter equation, Nasdaq returns

Distribution	Parameter	GARCH	PGARCH (rotation)	NGARCH	EGARCH	PGARCH (shift)
<i>Normal</i>	δ	1	0.675	1	0	0.770
	(SE)	-	(0.083)	-	-	(0.088)
	(SE)	-	(0.080)	-	-	(0.085)
	κ	2	1.350	2	1	1.540
	(SE)	-	-	-	-	-
	(SE)	-	-	-	-	-
	α_1	0.859	0.818	0.828	0.976	0.805
	(SE)	(0.012)	(0.025)	(0.015)	(0.004)	(0.023)
	(SE)	(0.045)	(0.050)	(0.056)	(0.014)	(0.049)
	α_2	0.123	0.138	0.136	0	0.189
	(SE)	(0.010)	(0.034)	(0.014)	-	(0.033)
	(SE)	(0.036)	(0.032)	(0.036)	-	(0.042)
	α_3	0	0.124	0	0	0
	(SE)	-	(0.015)	-	-	-
	(SE)	-	(0.061)	-	-	-
	α_4	0	0	0	0.182	0
	(SE)	-	-	-	(0.024)	-
	(SE)	-	-	-	(0.044)	-
	α_5	0	0	0	0.117	0
(SE)	-	-	-	(0.016)	-	
(SE)	-	-	-	(0.053)	-	
b	0	0	0.311	0	0.291	
(SE)	-	-	(0.047)	-	(0.044)	
(SE)	-	-	(0.097)	-	(0.091)	
	Log-Likelihood	-2888.7	-2872.2	-2875.5	-2876.2	-2872.0
<i>Student-t₃</i>	δ	1	0.590	1	0	0.614
	(SE)	-	(0.119)	-	-	(0.126)
	(SE)	-	(0.086)	-	-	(0.093)
	κ	2	1.180	2	1	1.228
	(SE)	-	-	-	-	-
	(SE)	-	-	-	-	-
	α_1	0.908	0.864	0.905	0.993	0.859
	(SE)	(0.012)	(0.028)	(0.012)	(0.004)	(0.028)
	(SE)	(0.036)	(0.040)	(0.039)	(0.008)	(0.042)
	α_2	0.113	0.161	0.113	0	0.198
	(SE)	(0.015)	(0.048)	(0.017)	-	(0.044)
	(SE)	(0.033)	(0.043)	(0.033)	-	(0.050)
	α_3	0	0.060	0	0	0
	(SE)	-	(0.032)	-	-	-
	(SE)	-	(0.037)	-	-	-
	α_4	0	0	0	0.201	0
	(SE)	-	-	-	(0.042)	-
	(SE)	-	-	-	(0.050)	-
	α_5	0	0	0	0.060	0
	(SE)	-	-	-	(0.035)	-
	(SE)	-	-	-	(0.039)	-
	b	0	0	0.101	0	0.088
	(SE)	-	-	(0.084)	-	(0.068)
(SE)	-	-	(0.088)	-	(0.058)	
v	6.592	5.773	6.260	5.059	6.190	
(SE)	(0.189)	(2.859)	(3.106)	(2.377)	(3.208)	
(SE)	(0.248)	(3.429)	(4.054)	(2.646)	(3.785)	
d	1.850	1.860	1.867	1.933	1.835	
(SE)	(0.021)	(0.194)	(0.191)	(0.207)	(0.190)	
(SE)	(0.020)	(0.237)	(0.267)	(0.227)	(0.229)	
β	0.877	0.889	0.890	0.879	0.883	
(SE)	(0.030)	(0.023)	(0.023)	(0.023)	(0.023)	
(SE)	(0.034)	(0.025)	(0.023)	(0.021)	(0.022)	
	Log-Likelihood	-2838.5	-2830.8	-2837.6	-2833.6	-2831.7

Notes. The Table reports parameter estimates for the Normal and Student-t₃ model and the various GARCH specification. Standard errors and robust standard errors proposed by White (1982) are given within parentheses.

Table 2.4: Estimates of the scale parameter equation, FTSE 100 returns

Distribution	Parameter	GARCH	PGARCH (rotation)	NGARCH	EGARCH	PGARCH (shift)
<i>Normal</i>	δ	1	0.673	1	0	0.756
	(SE)	-	(0.131)	-	-	(0.141)
	(SE)	-	(0.168)	-	-	(0.147)
	κ	2	1.346	2	1	1.512
	(SE)	-	-	-	-	-
	(SE)	-	-	-	-	-
	α_1	0.951	0.930	0.944	0.993	0.921
	(SE)	(0.007)	(0.019)	(0.006)	(0.002)	(0.017)
	(SE)	(0.011)	(0.022)	(0.014)	(0.003)	(0.028)
	α_2	0.044	0.044	0.039	0	0.070
	(SE)	(0.006)	(0.026)	(0.006)	-	(0.026)
	(SE)	(0.009)	(0.022)	(0.008)	-	(0.029)
	α_3	0	0.069	0	0	0
	(SE)	-	(0.020)	-	-	-
	(SE)	-	(0.031)	-	-	-
	α_4	0	0	0	0.056	0
	(SE)	-	-	-	(0.020)	-
(SE)	-	-	-	(0.016)	-	
α_5	0	0	0	0.070	0	
(SE)	-	-	-	(0.018)	-	
(SE)	-	-	-	(0.024)	-	
b	0	0	0.561	0	0.544	
(SE)	-	-	(0.159)	-	(0.152)	
(SE)	-	-	(0.158)	-	(0.151)	
	Log-Likelihood	-2525.2	-2517.3	-2516.6	-2520.0	-2515.2
<i>Student-t_3</i>	δ	1	0.707	1	0	0.690
	(SE)	-	(0.161)	-	-	(0.166)
	(SE)	-	(0.176)	-	-	(0.200)
	κ	2	1.414	2	1	1.380
	(SE)	-	-	-	-	-
	(SE)	-	-	-	-	-
	α_1	0.950	0.944	0.941	0.995	0.923
	(SE)	(0.009)	(0.014)	(0.009)	(0.003)	(0.017)
	(SE)	(0.010)	(0.014)	(0.011)	(0.003)	(0.022)
	α_2	0.082	0.043	0.071	0	0.107
	(SE)	(0.014)	(0.027)	(0.013)	-	(0.028)
	(SE)	(0.015)	(0.021)	(0.012)	-	(0.031)
	α_3	0	0.096	0	0	0
	(SE)	-	(0.027)	-	-	-
	(SE)	-	(0.030)	-	-	-
	α_4	0	0	0	0.069	0
	(SE)	-	-	-	(0.029)	-
	(SE)	-	-	-	(0.022)	-
	α_5	0	0	0	0.108	0
	(SE)	-	-	-	(0.030)	-
	(SE)	-	-	-	(0.031)	-
b	0	0	0.438	0	0.407	
(SE)	-	-	(0.135)	-	(0.116)	
(SE)	-	-	(0.120)	-	(0.074)	
v	4.469	3.638	3.648	3.741	3.551	
(SE)	(0.285)	(1.582)	(1.488)	(1.699)	(1.454)	
(SE)	(0.271)	(1.422)	(1.380)	(1.473)	(1.377)	
d	2.346	2.437	2.471	2.416	2.472	
(SE)	(0.026)	(0.292)	(0.287)	(0.292)	(0.291)	
(SE)	(0.023)	(0.285)	(0.283)	(0.266)	(0.288)	
β	0.938	0.976	0.980	0.969	0.972	
(SE)	(0.035)	(0.028)	(0.028)	(0.027)	(0.028)	
(SE)	(0.034)	(0.030)	(0.028)	(0.027)	(0.030)	
	Log-Likelihood	-2515.0	-2508.0	-2507.1	-2510.0	-2505.7

Notes. The Table reports parameter estimates for the Normal and Student- t_3 model and the various GARCH specification. Standard errors and robust standard errors proposed by White (1982) are given within parentheses.

We conduct a Lagrange multiplier test, using the augmented GARCH (1,1) as the general alternative. The LM statistic and robust LM statistic based on White (1982) of the Gaussian model and a model with skewness and conditional leptokurtosis is presented in Table 2.5. We also report the p-value of the χ^2 distribution with the number of the degrees of freedom equal to the number of parameters for the augmented GARCH model minus the number of parameter for the particular specification. The results underline the usefulness of a GARCH specification test. The choice of a particular distribution for the innovations (normal or skewed leptokurtic) influences the decision of rejecting one or the other specification.

Based on the LM test the standard GARCH model can be rejected for all indexes at the 1 or 5% level, but for example for the NASDAQ returns and based on the Schwarz Bayesian Information Criterion (Schwarz (1978)), we would prefer the symmetric model against 3 out of 4 asymmetric alternatives. For the S&P 500 under the hypothesis of conditional leptokurtosis and skewness and based on the robust LM statistics for example, the LM test rejects all specification except the power-GARCH specification with a rotation of the news impact curve. On the other hand the Schwarz Bayesian Information Criterion suggests that the EGARCH specification is the preferred model. The results suggest that the augmented GARCH-based LM test rejects rather strictly. As a result, the following rule is used to select the appropriate specification: under the hypothesis of conditional leptokurtosis and skewness and among the models that cannot be rejected, we base our decision on the Schwarz Bayesian Information Criterion and select one particular GARCH specification for the remaining analysis. Therefore under the hypothesis of conditional leptokurtosis and skewness we chose the power-GARCH model with rotation for the S&P 500 and NASDAQ returns and the NGARCH model with shift for the FTSE 100 returns.

Table 2.5: Augmented GARCH(1,1) based specification test

Index Innovation		GARCH	PGARCH (rotation)	NGARCH	EGARCH	PGARCH (shift)
S&P 500						
Normal	LM (p-value, df)	rejected (<1%,6)	rejected (<1%,4)	rejected (<10%,5)	rejected (<1%,5)	rejected (<5%,4)
	Robust LM (p-value, df)	rejected (<1%,6)	rejected (<5%,4)	9.0 (0.1075,5)	rejected (<1%,5)	5.5 (0.2362,4)
Student-t ₃	LM (p-value, df)	rejected (<1%,6)	rejected (<5%,4)	rejected (<1%,5)	rejected (<1%,5)	rejected (<1%,4)
	Robust LM (p-value, df)	rejected (<1%,6)	4.9 (0.3030,4)	rejected (<10%,5)	rejected (<10%,5)	rejected (<10%,4)
	SBC	-2339.4	-2333.9	-2336.0	-2330.7	-2339.0
Nasdaq						
Normal	LM (p-value, df)	rejected (<1%,6)	1.2 (0.8825,4)	rejected (<10%,5)	rejected (<1%,5)	0.8 (0.9401,4)
	Robust LM (p-value, df)	rejected (<5%,6)	1.0 (0.9060,4)	rejected (<10%,5)	rejected (<1%,5)	0.7 (0.9472,4)
Student-t ₃	LM (p-value, df)	rejected (<1%,6)	2.1 (0.7260,4)	rejected (<1%,5)	8.3 (0.1410,5)	3.9 (0.4206,4)
	Robust LM (p-value, df)	rejected (<1%,6)	1.6 (0.8054,4)	rejected (<5%,5)	8.1 (0.1482,5)	2.7 (0.6159,4)
	SBC	-2861.3	-2861.2	-2864.2	-2863.2	-2862.1
Ftse 100						
Normal	LM (p-value, df)	rejected (<1%,6)	5.7 (0.2214,4)	3.6 (0.6064,5)	rejected (<1%,5)	4.6 (0.3354,4)
	Robust LM (p-value, df)	rejected (<1%,6)	4.6 (0.3280,4)	3.0 (0.7041,5)	rejected (<1%,5)	4.0 (0.4090,4)
Student-t ₃	LM (p-value, df)	rejected (<1%,6)	rejected (<1%,4)	8.5 (0.1291,5)	rejected (<5%,5)	6.6 (0.1591,4)
	Robust LM (p-value, df)	rejected (<1%,6)	rejected (<5%,4)	6.2 (0.2858,5)	rejected (<5%,5)	5.1 (0.2727,4)
	SBC	-2537.8	-2538.4	-2533.7	-2536.6	-2536.1

Notes. Results of a Lagrange multiplier (LM) test, using the generalized augmented GARCH (1,1) process as the general alternative. Under the hypothesis of a normal or a skewed, leptokurtic distribution for the innovations, the LM statistic and robust LM statistic based on White (1982) suggest which of our models can be rejected (p-value and the number of the degrees of freedom are in parentheses). SBC reports the Schwarz Bayesian Information Criterion for the Student-t₃ model ($SBC=L^*-0.5k\ln(n)$), where L^* denotes the maximized log likelihood, k the number of estimated coefficients and n the number of observations).

Since we model a different scale parameter for each model, we cannot compare them by using likelihood ratio test. But given the particular specification we can analyze the residuals of each model by comparing the QQ-plots. Figure 2.1 and 2.2 show the QQ-plots for the residuals of the different models for the whole sample. For illustration we report the results for the NASDAQ index, but we obtain similar results for the other indices. Under the hypothesis of conditional normality, the first graph shows that the residuals are not normally distributed: there is an underestimation in both tails as well as a deviation from the empirical distribution in the center. Despite the asymmetric volatility parameter in the model there is still skewness in the residuals. The Student- t_3 specification leads to a better fit in both tails compared to the normal and the Student- t distribution (results not reported). Similar results for the NASDAQ are also obtained in Laurent (2002). Nevertheless there is still some residual underestimation in the left tail and the overestimation in the right tail is just reduced. In Chapter 3 we will discuss the truncated Lévy distribution, developed in the physics literature, which is more suitable to describe the behavior in the tails.

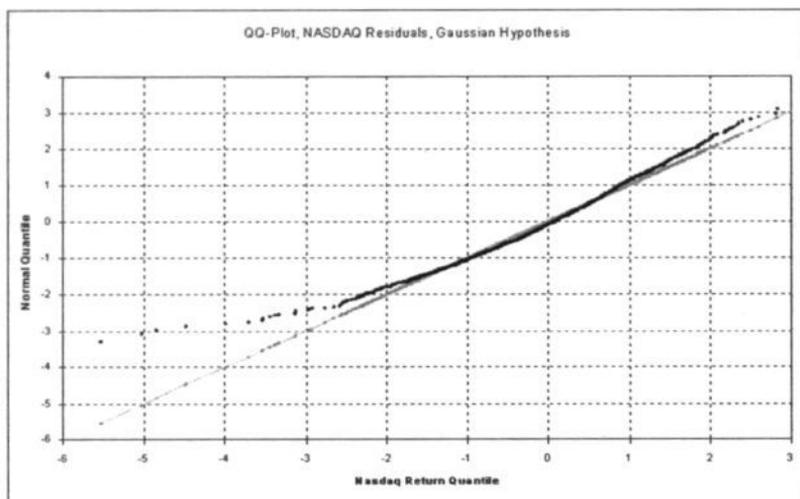


Figure 2.1: Quantile plot daily index returns

The graph depicts how the standardized empirical quantiles compare to the theoretical quantiles of a standard Normal distribution. A deviation from a straight line indicates a misspecification of the model. The underlying model has been estimated using daily returns from May 1992 until April 2000.

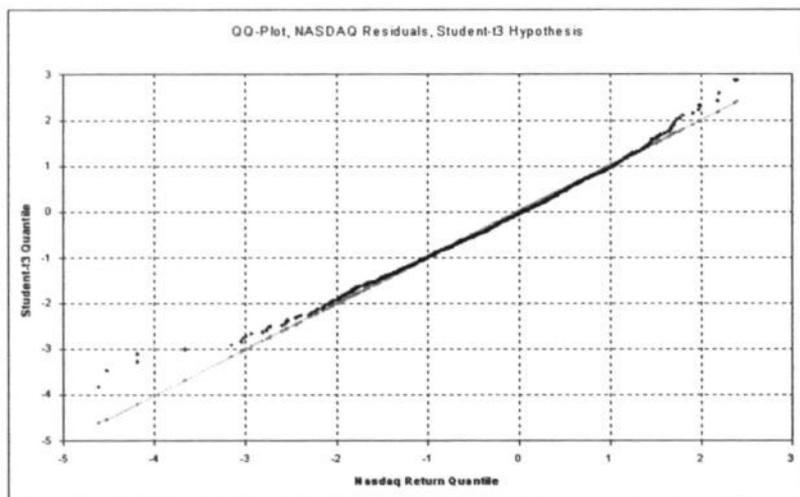


Figure 2.2: Quantile plot daily index returns

The graph depicts how the standardized empirical quantiles compare to the theoretical quantiles of a standard Student- t_3 distribution. A deviation from a straight line indicates a misspecification of the model. The underlying model has been estimated using daily returns from May 1992 until April 2000.

2.4 Concluding Remarks

In this chapter we have defined a very general class of time varying volatility models, the generalized augmented GARCH process, and compared different specifications for the innovations distribution for three equity index series. The augmented GARCH process incorporates two different approaches to model the asymmetry in the process: a shift or a rotation of the news impact curve. The appropriate specification is important, because previous research suggests that the difference in conditional volatility estimates could be substantial among various specifications (Hentschel (1995), Duan (1997)). We also included a model that incorporates conditional leptokurtosis and skewness. Since for most distributions with fat tails the existence of moments cannot be ensured, we generally describe the dynamics of the location and scale parameter of distributions. We used a Lagrange Multiplier test in order to match the appropriate GARCH specification with the conditional density function. The test allows us to reject various models and to find the appropriate specifications.

For particular index portfolios, we found evidence that the inclusion of more sophisticated tail properties for the conditional distribution leads to different conclusions about the appropriate type of asymmetry in the volatility process. In all time series a significant leverage effect is present, but at the same time a skewness parameter for the innovations distribution can be estimated. Therefore we found support for the hypothesis that the inclusion of conditional skewness additionally to the well-known leverage effect can explain the asymmetric response of volatility to the direction of returns and therefore the skewness in financial return data. Under the hypothesis of conditional skewness, models with different modifications of the news impact curve can be rejected for different markets.

However quantile comparison plots show that the asymmetric Student-t distribution is still not flexible enough for capturing the behavior of more volatile processes. In the next chapter we extend the analysis by proposing an alternative distribution that has been often studied in statistical physics.

CHAPTER 3

Modeling Scale-Consistent Value-at-Risk with the Truncated Lévy Flight⁷

3.1 Introduction

In this chapter we empirically investigate the tail characteristics and the so-called scaling behavior of the distributions of equity returns. We provide a conditional approach to the VaR methodology to estimate and forecast the downside risk of a portfolio of stocks considered in the previous chapter. In particular we introduce a stochastic process often used in statistical physics, but in this form never applied to financial data.

3.1.1 Modeling on different Time Scales

While financial applications involve many different time intervals, ranging from a few minutes (intraday) to a number of years, most techniques used in econometrics focus on modeling the fluctuations of price series in a single time interval. But the distribution that successfully explains daily price changes, for example, is typically unable to characterize the nature of hourly price changes. Likewise, the statistical properties of monthly price changes are often not fully covered by a model for daily price changes. In order to describe the statistics of future prices of a financial asset, one needs *a priori* a distribution of price

fluctuations for different time intervals, corresponding to different trading time horizons. In other words, the analysis of financial prices cannot be limited to one particular sampling interval (e.g. daily changes). In order to analyze the dynamics of financial prices and forecast future outcomes, it is appropriate to characterize *a priori* the relationship between return distributions for different sampling intervals (e.g. hourly-, ..., daily-, ..., weekly-, ... , monthly changes).

3.1.2 Physics meets finance

For derivative security pricing and risk management a better description of the large events is crucial. There is a need of an alternative to the Gaussian distribution with more weight in the tails and also considerably more weight around the center. A possible candidate is the Lévy distribution often studied in physics. Fitting the Lévy distribution to the data seems to provide a consistently better representation of returns (for an early reference see Fama (1963,1965)).

The problems with these kinds of distributions are the power law tails, which decay too slowly from a financial modeling perspective, and the fact that the distribution of price changes for larger time intervals converge to a Gaussian. This problem can be overcome by taking the Lévy distribution in the central part and introducing a cutoff in the far tails that is faster than the Lévy power law tails. The Lévy distribution with a cutoff and exponentially declining tails was introduced in the physics literature by Mantegna/Stanley (1994) and is known as a truncated Lévy process (TLP). This cutoff ensures that the variance will be finite and the distribution converges to a Gaussian in the limit. To model financial prices over time the truncated Lévy flight (TLF)⁸ can be constructed by the sum of independent and identically distributed random variables described by a truncated Lévy distribution (TLD).

In contrast to physical systems, hundreds of papers have shown that there is a strong non-i.i.d. clustering effect in financial data. Most of the studies so far performed a summability test of correlated unconditional return data and rejected the Lévy-stable hypothesis. But testing for stable properties requires i.i.d. data. For example Ghose and Kroner (1995) found that summability holds for data generated from an IGARCH process with non-normal innovations. But of course the simulated data is correlated and not i.i.d.

⁷ This chapter is based on: T. Lehnert, C.C.P. Wolff, 2001, 'Modeling Scale-consistent Value-at-Risk with the Truncated Levy Flight', LIFE Working Paper, Maastricht University.

⁸ The truncated Lévy process is the continuous time limit of the truncated Lévy flight.

Fortunately the class of GARCH models has been very successful in modeling the significant volatility clustering and non-i.i.d. nature of the data. More specifically, the standard GARCH model produces a mean reverting time dependent volatility process that “filters out” the correlations in the data and the remaining residuals are assumed to be i.i.d. For the family of Lévy distributions it is also appropriate to model the parameters that directly influence the shape of the distribution. Therefore we do not model the first and second moment of the return process directly, but the location and scale parameter of the distribution. We make use of the econometric framework developed the previous chapter. We show that asymmetry in the volatility process is not enough to capture the skewness in the data and that a combination of a skewed distribution for the innovations and asymmetry in the volatility process leads to a much better fit in-sample and to better out-of-sample performance.

A good parametric description of the distribution of price changes, especially in the tails, is important for VaR. If a distribution like the Lévy distribution overestimates and one like the Gaussian underestimates the probability mass in the tails, we expect the same result for the true VaR at high quantiles. Both situations are undesirable: understatement leads to capital charges for financial institutions which are too low, therefore it leads to increasing systemic risk, while overstatement leads to inefficient use of capital. We use the truncated Lévy flight as a stochastic process for different time intervals and model the time scaling behavior of the quantiles. We use the alpha-root-of-time rule of the Lévy flight for scaling up single day volatility to multi-day volatility and to forecast the multi-day density.

The outline of this chapter is as follows. In Section 3.2 we will describe the family of Lévy processes. Section 3.3 describes the econometric framework for the study. The empirical results are summarized in Section 3.4. Finally, Section 3.5 concludes.

3.2 Truncated Lévy Flights

3.2.1 The Truncated Lévy Distribution

Lévy flights have been observed experimentally in physical systems and have been used very successfully to describe for instance the spectral random walk of a single molecule embedded in a solid. In all these cases an unavoidable cutoff in the tails of the distribution is always present, which ensures the finiteness of the second moment of the process. One possible

cutoff is the exponential function, for which the characteristic function (CF) has been developed (Koponen (1995)). Note the misprint in the original publication (Nakao (2000)), the CF should read:

$$\psi_{TL}(k, \mu, C, \alpha, \lambda, \beta) = \tag{3.1}$$

$$i\mu k - C^\alpha \left\{ \frac{\lambda^\alpha - (k^2 + \lambda^2)^{\alpha/2}}{\cos(\pi\alpha/2)} \cos\left(\alpha \arctan\left(\frac{|k|}{\lambda}\right)\right) \left[1 + i \operatorname{sgn}(k) \beta \tan\left(\alpha \arctan\left(\frac{|k|}{\lambda}\right)\right) \right] \right\},$$

where μ is a location parameter, $C > 0$ is a scale parameter, α is the characteristic exponent determining the shape of the distribution and especially the fatness of the tails ($0 < \alpha \leq 2$, but $\alpha \neq 1$) and λ is the cutoff parameter, which determines the speed of the decay in the tails and as a result the cutoff region. The parameter β ($\beta \in [-1, 1]$) determines the skewness when $\beta \neq 0$, the distribution is skewed to the right when $-1 < \beta < 0$ and skewed to the left when $0 < \beta < 1$. For $\lambda \rightarrow +0$ the TLD reduces to the Lévy distribution. The density function is only known analytically when $\lambda \rightarrow +0$, $\beta=0$, $\alpha=1$, (Cauchy distribution) and $\lambda \rightarrow +0$, $\beta=0$, $\alpha=2$ (Gaussian distribution). However, for the symmetric case the value of the density of the Lévy distribution is known at the origin and in the far tails. For comparison purposes Figure 3.1 shows the density of a truncated Lévy distribution with reasonable parameter values for financial return data and the special case of a Gaussian density. Both densities are standardized, such that the scale parameter C equals one.

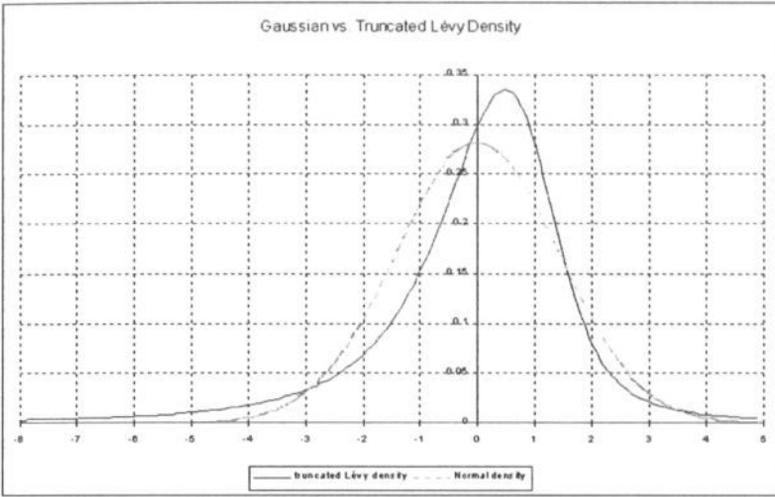


Figure 3.1: Truncated Lévy Density

The graph depicts a comparison of a truncated Lévy density and a normal density. Both densities are standardized, such that the scale parameter C is equal to one.

In contrast to the Lévy distribution the exponential cutoff ensures that all moments exist. Accurate numerical values for the density ψ_L can be calculated by Fourier-transforming the CF and evaluating the integral numerically. We use Romberg integration, which allows ex-ante specification of the tolerated error and in fact a calculation of the density as precise as necessary (see Lambert and Lindsey (1999)). The m -th moment of the distribution can be found from the CF using

$$\langle X^m \rangle = \left(i^m \frac{\partial^m}{\partial k^m} \Psi_{TL}(k) \right) \Bigg|_{k=0}. \quad (3.2)$$

3.2.2 Scaling Properties

In practice, comparing the distributional properties of price increments at various time intervals provides insight into the temporal dependence structure of the time series. It is possible to reconstruct the distributions for different time intervals from the knowledge of the distribution for short time intervals only if we assume independent and identically distributed (i.i.d.) price changes. The normal and the Lévy distributions play a central role in this context, because they are stable under addition: the distribution of the sum of a large number of i.i.d.

random variables belongs to the family of Lévy distributions (generalized Central Limit Theorem) (Gnedenko and Kolmogorov (1954)).

Analyzing the time scaling behavior of financial fluctuations means comparing the increments for shorter time scales τ and for longer time scales $N\tau$. This formally corresponds to summing N random variables. In the case of the Lévy distribution the characteristic function satisfies $N\psi_L(k) = \psi_L(N^{1/\alpha}k)$. The distribution for various time scales for stationary and independent variables is related by a convolution relation: $P_{N\tau} = P_\tau \otimes P_\tau \otimes \dots \otimes P_\tau$ (\otimes indicates a convolution). More generally the distribution $P(x)$ of price changes on a time scale $N\tau$ may be obtained from that of a shorter time scale τ by a rescaling of the variable: $P_L^N(x) = N^{-1/\alpha} P_L^1(N^{-1/\alpha}x)$, where $P^N(x)$ denotes a N -times convoluted distribution of $P^1(x)$.

This relation means that the process $x(t)$ is *self-similar* with a constant self-similar exponent $\frac{1}{\alpha}$. The process is characterized by the scaling behavior of its moments. We can also refer to this as the scaling or fractal property of stochastic processes. The simplest case is the self-similar process with a constant scale factor which is uni-fractal in this sense. Mandelbrot et al. (1997) introduced the idea of a multifractal process with a time depending random scale factor.

The scenario of a *scale invariant* price process is that of a Lévy flight with a characteristic exponent α which is the inverse of the self-similarity exponent. In other words, if the random variable for short time intervals is distributed according to a Lévy distribution, then in the limit the distribution for long time intervals is also a Lévy with the same characteristic exponent α . The infinite variance of the Lévy distribution prevent that it converges to a Normal distribution.

In Mantegna and Stanley (1994) this scale invariant behavior is observed for short time scales with high frequency data. Cont et al. (1997a) showed that it breaks down for longer time scales. These observations have been explained as a structural break in terms of the truncated Lévy Flight. The scaling behavior is also observed for the variance and the kurtosis of price increments (Cont et al. (1997a)). A link can be established between the scaling behavior of the moments and high-order correlation functions of the time series (Cont (1997b)).

Since we introduced a cut-off for the CF of the truncated Lévy $\psi_{TL}(k, \lambda)$, it is no longer self-similar or uni-fractal by the criteria mentioned above, but bi-fractal, the simplest

version of a multi-fractal process (see Nakao (2000)). The convolution of the probability distribution can still be obtained by scaling both x and λ . The CF $\psi_{TL}(k, \lambda)$ satisfies $N\psi_{TL}(k, \lambda) = \psi_{TL}(N^{1/\alpha}k, N^{1/\alpha}\lambda)$ and the N -times convoluted probability distribution satisfies $P_{TL}^N(x, \lambda) = N^{-1/\alpha}P_{TL}^1(N^{-1/\alpha}x, N^{1/\alpha}\lambda)$. For short time scales (daily) the process behaves like a Lévy flight, but converges towards a Gaussian for longer time scales (say monthly) (see Matacz (2000)). The scaling of λ ensures that the process converges towards a Gaussian process instead of staying a Lévy flight. As a result for increasing time scales the cutoff is introduced earlier in the tails, a feature especially interesting for modeling financial fluctuations over time. Typically price changes tend to be leptokurtic and skewed for higher sampling intervals (daily or intra-daily) and nearly normally distributed for lower sampling intervals (monthly).

3.2.3 The Gaussian Process as a Special Case

To illustrate the effect of scaling in the RiskMetricsTM approach, a Gaussian stochastic process $x(t)$ can be defined by the CF

$$\psi_G(k) = -\frac{1}{2}\sigma^2 tk^2, \quad (3.3)$$

where σ is the scale parameter of the process and therefore assumed to grow as \sqrt{t} . This is the scaling rule of a usual Gaussian random walk for the standard deviation and used in the standard RiskMetricsTM approach⁹. Our approach can be interpreted as a generalization of the RiskMetricsTM methodology.

3.3 Model Specifications

3.3.1 The Generalized Augmented GARCH Process

Traditional GARCH models (with normal- or Student-t distributed error terms) were designed to capture clustering of large and small innovations, which can be modeled as serially

⁹ See Diebold et al. (1997) for a critical review.

correlated conditional second moment when the variance exists (Bollerslev et al. (1992)). The analogue of the standard deviation σ in the family of Lévy distributions is the scale parameter C . If we replace the standard deviation σ by the scale parameter C , we allow C_t to be serially correlated, which produces the volatility clustering.

A general class of models that is frequently used together with Lévy-stable distributions (see Mittnik et al. (2000b, 2002)) is the power-GARCH process of Ding, Granger and Engle (1993), discussed in Section 2.2. The model incorporates a power transformation of the conditional standard deviation process (Equation (2.2)). Mittnik et al. first derived a closed form expression for a symmetric power-GARCH process with Lévy-stable innovations. They show that as δ approaches the characteristic exponent of the Lévy distribution α , $E|\varepsilon_t|^\delta$ increases without bounds and leads to an explosive process, which is not covariance stationary. Some GARCH-stable models in the literature are estimated setting δ equal to α and therefore fail to address the correct measure of stationarity¹⁰. The reason for this is that the characteristic exponent α of the Lévy distribution gives a restriction on the number of moments that exist. Since the stationarity condition of an asymmetric power-GARCH process requires the calculation of the fractional moment, $\delta=\alpha$ is the limit case and not defined. Apart from the special case $\lambda \rightarrow 0$, the moments of the truncated Lévy distribution are finite and the process would be covariance stationary for $\delta < \alpha$.

In the following, we consider the generalization of the augmented GARCH process, developed in Section 2.2. The single lag version of the complete augmented GARCH model driven by a truncated Lévy distribution reads:

$$\begin{aligned}
 r_t &= \mu_1 + C_t \varepsilon_t, & \varepsilon_t &\sim \text{truncated Lévy } (0, 1, \alpha, \lambda, \beta) \\
 \phi_t &= \alpha_0 + \gamma_{1,t-1} \phi_{t-1} + \gamma_{2,t-1} \\
 C_t &= \begin{cases} |\delta \phi_t - \delta + 1|^{1/2\delta} & \text{if } \delta \neq 0 \\ \sqrt{\exp(\phi_t - 1)} & \text{if } \delta = 0 \end{cases} \\
 \gamma_{1,t} &= \alpha_1 + \alpha_2 |\varepsilon_t - b|^\kappa + \alpha_3 \max(0, b - \varepsilon_t)^\kappa \\
 \gamma_{2,t} &= \alpha_4 \frac{|\varepsilon_t - b|^\kappa - 1}{\kappa} + \alpha_5 \frac{\max(0, b - \varepsilon_t)^\kappa - 1}{\kappa},
 \end{aligned} \tag{3.4}$$

¹⁰ The paper by Liu and Brorsen (1995) is the most recent example.

where the conditional location parameter μ_t can be specified additionally and the conditional scale parameter C_t is assumed to vary over time. The finite stationary scale parameter and the sufficient condition for strict stationarity of the location adjusted augmented GARCH(1,1) process can be expressed in a general condition, but has to be evaluated numerically (see Equation (2.5)).

For the in-sample analysis, we compare two different specifications for the innovations distribution: the asymmetric Student-t distribution called t_3 , introduced in Chapter 2 and the truncated Lévy distribution. The GARCH model with normally distributed innovations is nested into both cases with time varying volatility and skewed t_3 or truncated Lévy distributed error terms.

GARCH models driven by the truncated Lévy distribution are also estimated with maximum likelihood. The log-likelihood is given by

$$\ln L = \sum_t (\ln TL(\varepsilon_t, 0, 1, \alpha, \lambda, \beta) - \ln(C_t)) \quad (3.5)$$

$TL(\varepsilon_t, 0, 1, \alpha, \lambda, \beta)$ is the standardized truncated Lévy density and the term $-\ln(C_t)$ results from taking the log of the Jacobian of the transformation. The ML estimators are also assumed to follow the standard theory, so they are consistent and asymptotically normal with mean the parameter estimates and a variance-covariance matrix (DuMouchel (1973)).

3.3.2 Density Forecasting and VaR

Once the parameters of the model are estimated, we can forecast the VaR for a particular confidence level and time horizon by determining the multi-period return distribution. The VaR of a position with a confidence level q and for N periods ahead can be defined as:

$$\text{VaR}_{q,N} = W_0 \left\{ 1 - Q_{1-q} \left(\exp \sum_{t=1}^N r_t \right) \right\}, \quad (3.6)$$

where W_0 is the initial value of the position, r_t denotes the continuously compounded return at time t and $Q_{1-q}(\cdot)$ denotes the quantile-operator that calculates the $(1-q)$ -th quantile of the

multi-period return distribution. There are various ways to determine the r_t 's or the sum of r_t 's of the underlying return process. Under the assumption of conditional normality, we can use the stable property of the normal distribution and scale the standard deviation at time t by a square-root-of-time rule to get the multi-day standard deviation. This is also the method used in the standard RiskMetrics™ approach, which is special case of our augmented GARCH model under the hypothesis of Gaussian innovations for $\delta=1$, $\kappa=2$, $\alpha_0=0$, $\alpha_3=0$, $\alpha_4=0$, $\alpha_5=0$, $b=0$, α_1 free and $\alpha_2=1-\alpha_1$. The method is frequently applied in practice, but it has been shown to underestimate the downside risk consistently (see e.g. Pownall and Koedijk (1999)).

Equivalently, under the hypothesis of conditional leptokurtosis and skewness, we can use the characteristics of the truncated Levy flight introduced in Section 3.2. Once the parameters of the model are estimated, we know the probability distribution of one particular time scale, say daily. The N -times convoluted probability distribution can be obtained by scaling the parameter C and the cut-off parameter λ by a α -root-of-time rule to derive the multi-day parameters. Therefore the VaR of a position with a confidence level q and for N periods ahead can be defined as:

$$VaR_{q,N} = W_0 \left(1 - \exp\left(R_N^*\right) \right), \quad (3.7)$$

where $\exp\left(R_N^*\right)$ can be determined by setting R_N^* equal to the $(1-q)$ -th quantile of the truncated Lévy distribution with the characteristic function $\psi_{TL}(\mu_t N, N^{1/\alpha} C_t, N^{1/\alpha} \lambda, \beta)$, where μ_t and C_t are the location and scale parameter, respectively, α is the characteristic exponent of the truncated Lévy distribution and $N^{1/\alpha}$ is the adjustment factor needed to scale the parameter. Equivalently for the RiskMetrics™ EWMA method, $\exp\left(R_N^*\right)$ can be determined by setting R_N^* equal to the $(1-q)$ -th quantile of the normal distribution with the characteristic function $\psi_G(\mu_t N, \sqrt{N} \sigma_t)$, where μ_t and σ_t are the mean and standard deviation, respectively, \sqrt{N} is the adjustment factor needed to scale the standard deviation.

3.3 Empirical Results

3.3.1 In-sample Analysis

The following empirical analysis is based on the financial time series specified in Section 2.3. In Section 2.4, under the hypothesis of conditional leptokurtosis and skewness and among the models that cannot be rejected, we based our decision on the Schwarz Bayesian Criterion (SBC) (Schwarz (1978)) and selected one particular GARCH specification for each return series for the remaining analysis. A Lagrange Multiplier test was also used for the augmented GARCH model driven by truncated Lévy distributed innovations (results not reported) and the same models were preferred. As a result, for the VaR analysis in this section we consider a Power-GARCH process with rotation for the S&P 500 and NASDAQ return series and the NGARCH model for the FTSE 100 return series.

In Table 3.1, the parameter estimates for the GARCH models driven by a truncated Lévy distribution and a Student- t_3 distribution are compared. For the estimation the conditional mean μ_t is set equal to the unconditional value. The comparison shows two effects. First, despite the time varying volatility framework, the conditional distribution has to be leptokurtic in order to capture the excess kurtosis in the data. The estimated parameters for the conditional tail fatness (ν and d for the Student t_3 distribution and α and λ for the truncated Lévy distribution) are significantly different from values for the normal distribution. Second, despite the negative conditional skewness for the innovations distribution (given a value of β lower than 1 for the Student t_3 distribution, except for the S&P 500 return series, and β greater than 0 for the truncated Lévy distribution), the time varying volatility process has to be asymmetric in order to capture the negative skewness in the data. The parameter values suggest that a leverage effect is present in index returns in addition to the conditional skewness. Nevertheless parameter values are not always significant, especially for the Student- t_3 distribution. In particular the parameters ν and d measuring the tail fatness are very volatile suggesting that the Student- t_3 distribution does not successfully model the conditional leptokurtosis and skewness. In contrast, the parameter measuring tail fatness for the truncated Lévy distribution α is significant, however β is not always significantly different from zero.

Table 3.1: Parameter Estimates of the scale parameter equation

Index / Model		Parameter								
δ	κ	α_1	α_2	α_3	α_4	α_5	b	v/α	d/λ	β
S&P500										
Power-GARCH, Student- t_3										
0.663	1.325	0.917	0.038	0.157	0	0	0	5.343	1.733	1.008
(0.100)	-	(0.019)	(0.030)	(0.039)	-	-	-	(2.447)	(0.179)	(0.027)
(0.153)	-	(0.025)	(0.028)	(0.049)	-	-	-	(2.534)	(0.170)	(0.018)
Power-GARCH, truncated Lévy										
0.880	1.760	0.927	0.008	0.072	0	0	0	1.698	0.235	0.276
(0.033)	-	(0.015)	(0.009)	(0.020)	-	-	-	(0.105)	(0.117)	(0.133)
(0.058)	-	(0.018)	(0.007)	(0.026)	-	-	-	(0.185)	(0.153)	(0.248)
NASDAQ										
Power-GARCH, Student- t_3										
0.590	1.180	0.864	0.161	0.060	0	0	0	5.773	1.860	0.889
(0.119)	-	(0.028)	(0.048)	(0.032)	-	-	-	(2.859)	(0.194)	(0.023)
(0.086)	-	(0.040)	(0.043)	(0.037)	-	-	-	(3.429)	-	(0.025)
Power-GARCH, truncated Lévy										
0.791	1.582	0.877	0.040	0.071	0	0	0	1.814	0.089	0.918
(0.052)	-	(0.022)	(0.015)	(0.022)	-	-	-	(0.045)	(0.041)	(0.141)
(0.049)	-	(0.029)	(0.012)	(0.022)	-	-	-	(0.036)	(0.046)	(0.120)
FTSE 100										
NGARCH, Student- t_3										
1	2	0.941	0.071	0	0	0	0.438	3.648	2.471	0.980
-	-	(0.009)	(0.013)	-	-	-	(0.135)	(1.488)	(0.287)	(0.028)
-	-	(0.011)	(0.012)	-	-	-	(0.120)	(1.380)	(0.283)	(0.028)
NGARCH, truncated Lévy										
1	2	0.935	0.019	0	0	0	1.116	1.906	0.164	0.419
-	-	(0.010)	(0.004)	-	-	-	(0.274)	(0.065)	(0.157)	(0.289)
-	-	(0.011)	(0.003)	-	-	-	(0.242)	(0.055)	(0.117)	(0.289)

Notes. The Table reports parameter estimates for the truncated Lévy and Student- t_3 model and the different GARCH specifications. The underlying data set consists of daily observations for the period May 1992 until April 2000. The location parameter is assumed to be equal to the unconditional value. Standard errors and robust standard errors proposed by White (1982) are given within parentheses.

Since we model a different scale parameter for each model, we cannot compare them by using likelihood ratio test. But given the particular specification we can analyze the residuals of each model by comparing the QQ-plots, as we did in Chapter 2. Figure 3.2 and Figure 3.3 shows the QQ-plots for the residuals of the different models for the whole sample. For illustration we report the results for the NASDAQ index, but we obtain similar results for the other indices. As already mentioned the t_3 specification leads to a better fit in both tails compared to the normal, but the underestimation in the left tail and the overestimation in the right tail is just reduced. The QQ-plot for the residuals of the truncated Lévy model shows that the fit in the tails is extremely precise; there is only a slight overestimation in the tails. This result does not hold for the model without cutoff in the tails, as the power law tails of the Lévy distribution lead to an extreme overestimation of the extreme events (results not reported).

The results suggest that the Student- t_3 distribution is not flexible enough to capture the behavior of more volatile processes when compared to an alternative often used in the physics literature. Therefore the in-sample results for the truncated Lévy distribution are quite promising and the following section will show if this superior performance is also evident out-of-sample.

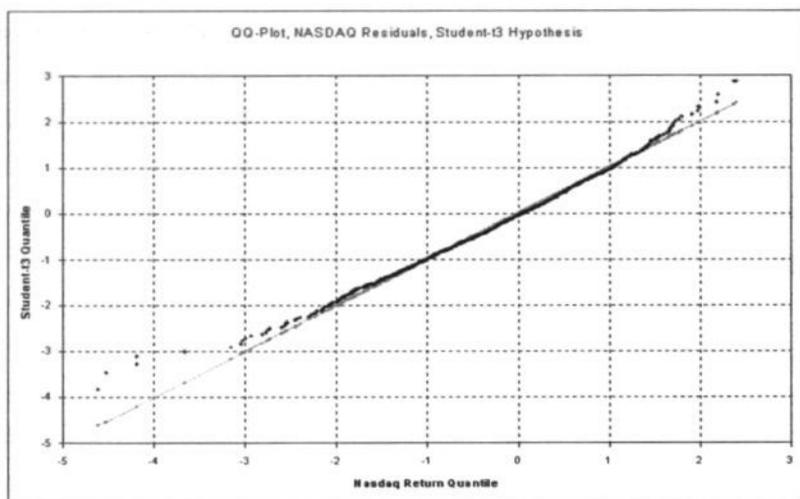


Figure 3.2: Quantile plot daily index returns

The graph depicts how the standardized empirical quantiles compare to the theoretical quantiles of a standard Student- t_3 distribution. A deviation from a straight line indicates a misspecification of the model. The underlying model has been estimated using daily returns from May 1992 until April 2000.

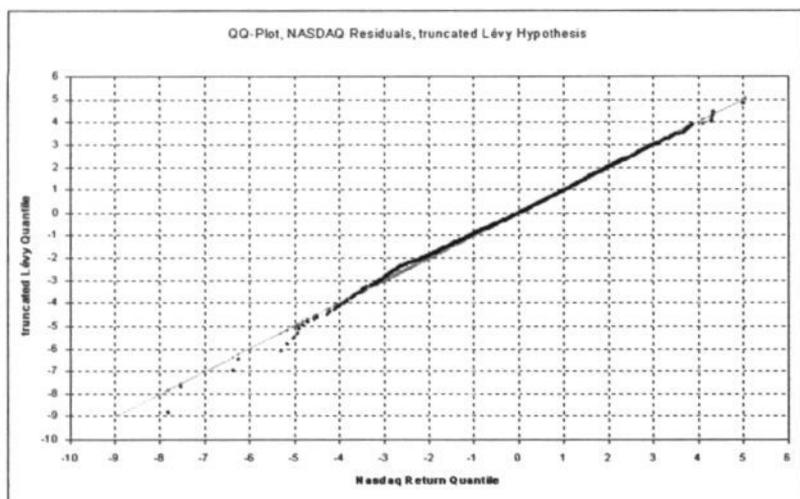


Figure 3.3: Quantile plot daily index returns

The graph depicts how the standardized empirical quantiles compare to the theoretical quantiles of a standard truncated Lévy distribution. A deviation from a straight line indicates a misspecification of the model. The underlying model has been estimated using daily returns from May 1992 until April 2000.

3.3.2 Out-of-Sample Analysis

Since we are interested in the out-of-sample performance, we back-test our model for all indices over a period May 1996 until March 2000 and compare it to the RiskMetrics™ EWMA approach. Every day we estimate the model using the last about 1000 trading days (that means exactly one half of each sample and a moving window) and forecast the 99% (95%) VaR 5, 10 and 20 days ahead. Therefore we obtain a VaR estimate for a certain confidence level each day and compare it to the actual change in the position (given an initial position $W_0=100$), determined by the realized cumulative return over the following 5, 10 and 20 days. Whenever the actual change in the position is greater than the estimated VaR, we notice a so-called violation for that particular day. Given a confidence level of 99% (95%), we assume that on average the estimated VaR is violated 1% (5%) of the time. A higher (lower) violation rate indicates that the model underestimates (overestimates) the actual downside risk. Table 3.2 reports the out-of-sample results.

Table 3.2: Violations of the actual Value-at-Risk

	Horizon (days)	VaR (Conf.Level)	S&P 500	NASDAQ	FTSE 100
<i>EWMA</i>	5	99%	2,6%	3,3%	2,9%
		95%	5,9%	7,6%	6,1%
	10	99%	1,8%	2,7%	1,1%
		95%	5,2%	6,6%	4,9%
	20	99%	0,9%	2,3%	1,3%
		95%	4,6%	7,2%	4,3%
<i>Augmented GARCH Truncated Lévy</i>	5	99%	1,5%	1,2%	1,3%
		95%	5,2%	4,9%	5,3%
	10	99%	1,2%	1,3%	0,8%
		95%	4,3%	5,5%	4,0%
	20	99%	0,7%	1,4%	0,9%
		95%	3,7%	4,6%	3,9%

Notes. The Table reports the percentage violations of the actual Value-at-Risk for the alternative models during the ‘backtesting period’ (second half of the sample). VaR-estimates are calculated given an initial position of $W_0=100$. Given a confidence level of 95% (99%), the model assumes a violation rate of 5% (1%). A higher (lower) violation rate indicates that the model underestimates (overestimates) the actual Value-at-Risk.

The results are very promising: the VaR estimates we obtained by using our method compared to the RiskMetrics™ method produces on average less violations of the expected VaR for all confidence intervals and horizons. The RiskMetrics™ method constantly

underpredicts extreme events and this leads very often to an inappropriate number of violations. This underprediction is slightly reduced for lower confidence intervals or longer forecasting horizons, a well-known result for the Gaussian distribution and the square-root-of-time scaling rule. On the other side the GARCH model with the truncated Lévy distribution and the alpha-root-of-time scaling rule leads to an appropriate number of violations for low *and* high confidence intervals and short *and* long forecasting horizon. This means not only that the fit in the tails of the distribution is very good, but also that the scaling rule captures the scaling behavior of the data very well and shows a convergence from a skewed leptokurtic distribution to a Gaussian for larger sampling intervals. This is actually the unique bi-fractal scaling behavior of the truncated Lévy flight.

3.4 Conclusions

In this chapter, a stochastic process, initially developed in the physics literature, is applied to finance. We analyze the dynamics of financial prices and propose a generalization of the popular RiskMetricsTM approach to calculate VaR for a stock portfolio. The advantage of our approach is that we are able to capture the observed conditional tail fatness and skewness in financial returns by using truncated Lévy flights for the innovations of a GARCH process. We propose a new scaling rule to forecast volatility. Location and scale parameter (volatility) are estimated on one time scale (daily) and the multi-day (weekly, bi-weekly or monthly) values are derived by using the bi-fractal scaling property of truncated Lévy processes. The method has the implied advantage that we are able to identify the relationship between return distributions for different sampling intervals by analyzing the time series of returns on one sampling interval (say daily). Therefore the approach can produce better forecasts compared to methods, which add up one-day forecasts to derive a multi-day forecast.

Both the in-sample performance of the model for VaR calculations and the out-of-sample performance are quite promising. There are plenty of practical benefits of this model, which can be explored in future research.

The first part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow \infty$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow \infty$. The second part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow 0$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow 0$.

The third part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow \infty$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow \infty$.

The fourth part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow 0$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow 0$.

The fifth part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow \infty$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow \infty$.

The sixth part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow 0$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow 0$.

The seventh part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow \infty$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow \infty$.

The eighth part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow 0$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow 0$.

The ninth part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow \infty$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow \infty$.

The tenth part of the paper is devoted to the study of the asymptotic behavior of the solutions of the system (1.1) as $t \rightarrow 0$. It is shown that the solutions of the system (1.1) are bounded and tend to zero as $t \rightarrow 0$.

CHAPTER 4

An Evaluation Framework for Alternative VaR Models¹¹

4.1 Introduction

In this chapter we concentrate our efforts on extreme deviations in exchange rate positions. Different models have been proposed in the literature to model extreme events. These models usually include distribution functions that allow for fat tails, like the Student-t distribution or the stable distribution. Also time-varying volatility models are popular, like GARCH specifications, which we introduced in Chapter 2. The characteristics of the return distribution of FX-rates are typically different to the ones for equity returns. Due to the particular nature of exchange rates, significant skewness and asymmetry is typically not present. In contrast to Chapters 2 and 3, this allows us the adoption of simpler GARCH specifications driven by symmetric distributions. The choice of an adequate distribution function is nevertheless an important one: if a particular distribution function does not allow for an empirical phenomenon, then this also will not appear in future VaR-estimates. This may encourage the adoption of more sophisticated distribution functions, which include all sorts of fat-tail characteristics

In this chapter we argue that this is not necessarily the preferred approach. In VaR-applications not only the distribution function plays an important role, but also the parameter values of these distribution functions. Parameter estimates usually follow from historical data.

¹¹ This chapter is based on: D. Bams, T. Lehnert and C.C.P. Wolff, 2001, 'An Evaluation Framework for alternative VaR models', LIFE Working Paper, Maastricht University.

When a particular phenomenon is not present in the historical data, then the parameters of the distribution function that take account of the particular phenomenon are estimated with a lot of uncertainty, reflected by the standard error of the parameter estimate.

Uncertainty in the parameter estimates implies uncertainty in the underlying distribution function and hence uncertainty in the reported VaR. Especially because VaR-estimates stem from the tail of the distribution function, estimates for the associated tail-parameters follow primarily from extreme events. Usually, not many historical observations of extreme events are available. This may lead to the assertion that models with advanced tail-specifications lead to a better location of the VaR than simple models. At the same time these models imply more uncertainty in the reported VaR, because of the higher level of uncertainty in the parameter estimates. In the empirical part of this chapter we will report VaR-estimates for different models. The VaR-estimates will include an expected VaR and a standard deviation that reflects the uncertainty around this expected VaR. We will show that it is important to take estimation risk into account. Ignoring it may lead to an over- or an under-estimation of the actual VaR. Barberis (2000) also takes parameter uncertainty into account when considering predictability in future stock returns. The trade-off between the location and the preciseness of reported VaR provides a first yardstick on which the adequacy of VaR-models should be evaluated.

A popular method to choose between alternative time series models starts with a simple specification that is subsequently extended to more advanced specifications for volatility and/or tail-behavior. Typically the parameters of the models are estimated using the Maximum Likelihood method. A likelihood ratio test suffices in this particular case to decide which model best describes the data. Note that this type of in-sample comparison only works for models that are nested.

There are at least two reasons why this in-sample selection method may not necessary lead to the optimal VaR-model. First, this method is not applicable for models that are non-nested. Second, the fact that a particular model best fits historical data does not mean that it also provides the best VaR-forecast. We propose an out-of-sample approach in which we split the data sample into two parts. The first part is used to estimate the models. The second part is used to compare the forecasted VaR with a realized change. Repeating this approach for many different sub-samples allows us to test whether the number of realized changes are consistent with the reported VaR at the given confidence level.

The aim of this chapter is to provide an empirical selection approach to arrive at the most suitable VaR-model. The purpose of such a model is extreme loss forecasting. In our

view such an approach should deal with uncertainty in the reported VaR that stems from parameter uncertainty. The out-of-sample selection method as described in the previous chapters is in our view suitable to arrive at the appropriate VaR-model. First, it focuses on the purpose of the model. Second, it allows for a comparison of alternative models, and third it can take into account the uncertainty in the forecasted VaR. Eventually, we propose an adjustment to the ‘best estimate’-reported VaR to account for parameter uncertainty. In order to illustrate our approach, we will focus on VaR estimates in the context of exchange rate positions from the point of view of a currency trader.

In the next section we set up the econometric framework, in Section 4.3 we describe the data, Section 4.4 provides empirical results and Section 4.5 concludes.

4.2 Econometric Framework

4.2.1 Standard Time Series Models

Usually, financial time series are not modeled in terms of prices but in terms of returns. In the empirical part we deal with exchange rates returns. We consider two types of models to describe the return on exchange rates. Also, restricted versions of these two types are taken into consideration. In the first model the returns follow an AR(1) process with error terms that are Student-t distributed and with a degree of volatility that is described by a GARCH (1,1) process, introduced in Chapter 2. The model reads:

$$r_t = \mu + \rho(r_{t-1} - \mu) + \varepsilon_t \quad (4.1)$$

$$\varepsilon_t \sim t(0, \sigma_t^2, \theta) \quad (4.2)$$

$$\sigma_t^2 = \beta_0 + \beta_1 \sigma_{t-1}^2 + \beta_2 \varepsilon_{t-1}^2 \quad (4.3)$$

where r_t denotes the return at time t , μ is the expected return and ρ is the mean-reversion parameter. The error term ε_t is Student-t distributed with volatility σ_t^2 at time t and degrees of freedom θ . Equation (4.3) describes the GARCH(1,1) process. Special cases of this first type of model are the random walk specification ($\rho = 0$), the normal distribution ($\theta \rightarrow \infty$) and constant volatility model ($\beta_1 = \beta_2 = 0$). The log-likelihood function for this model reads

$$\begin{aligned} \ln L = & T \ln \Gamma\left(\frac{\theta+1}{2}\right) - T \ln \Gamma\left(\frac{\theta}{2}\right) - \frac{1}{2} T \ln[\pi(\theta-2)] + \\ & - \frac{1}{2} \sum_{i=1}^T \ln(\sigma_i^2) - \left(\frac{\theta+1}{2}\right) \sum_{i=1}^T \ln \left[1 + \frac{\eta_i^2}{(\theta-2)\sigma_i^2} \right] \end{aligned} \quad (4.4)$$

where $\Gamma(\cdot)$ is the Gamma-function, T is the number of historical observations and r_1, \dots, r_T denote the historical observations.

4.2.2 Modeling Extreme Events with Lévy-Stable Distributions

The second family of models that are taken into consideration is the family of Lévy-stable distributions. The model reads:

$$r_t = \mu + \rho(r_{t-1} - \mu) + \eta_t \quad (4.5)$$

$$\eta_t \sim S(0, c_t, \alpha) \quad (4.6)$$

$$c_t^\delta = \gamma_0 + \gamma_1 \sigma_{t-1}^\delta + \gamma_2 |\eta_{t-1}|^\delta \quad (4.7)$$

The error terms, η_t , follow a Lévy-stable distribution. The volatility is denoted with c_t , which is the scale parameter of the Lévy-stable distribution; the shape in the center of the distribution is determined by the parameter α , which is referred to as the characteristic exponent. Heteroskedasticity is modeled, in Equation (4.7) by a power-GARCH model, in which the parameter δ is estimated along with the other parameter (Ding, Granger and Engle (1993)). The reason for using this particular specification is that the characteristic exponent α of the Lévy-stable distribution gives a restriction on the number of moments that exist (see Section 3.3.1 for a details). The normal distribution is a special case of the Lévy-stable distribution, which follows by setting $\alpha = 2$. The log-likelihood function for the Lévy-stable model reads

$$\ln L = \sum_{i=1}^T \ln S_i(0, 1, \alpha) - \sum_{i=1}^T \ln c_i \quad (4.8)$$

where $S(\cdot)$ denotes the density of the Lévy-stable distribution. Like the truncated Lévy distribution introduced in Chapter 3, the Lévy-stable distribution is not known in analytical

form. Accurate numerical values for the density of the Lévy-stable distributions can also be calculated Fourier-transforming the characteristic function and evaluating the integral numerically (see Section 3.2.1 for details). The characteristic function of the four-parameter Lévy-stable distribution is given by (Mandelbrot (1963), Fama and Roll (1968)):

$$\psi_L(k) = i\mu k - c^\alpha |k|^\alpha \left[1 + i\beta \operatorname{sgn}(k) \tan\left(\frac{1}{2}\pi\alpha\right) \right] \quad 0 < \alpha \leq 2 \quad (4.9)$$

In the empirical application we estimate parameters for the three-parameter Lévy-stable distribution, since we are interested in the symmetric alternative.

The two families are non-nested so an in-sample approach to check which of the two best describes the data is not a viable way to proceed. By choosing the two families of distributions as given above, we seek to demonstrate the strengths and weaknesses of an advanced VaR-model relative to a standard type of model. The same approach is applicable to other types of distributions.

4.2.3 VaR and Parameter Uncertainty

An advanced model allows for more complex shapes of the tail and hence has the advantage of better describing the VaR. At the same time this may lead to more uncertainty in the parameters and hence in the VaR estimate itself. Uncertainty in the parameters may be quantified by the covariance matrix of the parameter estimates. Let p denote the vector of unknown parameters and let $\ln L(p)$ denotes the associated log-likelihood. The covariance matrix of parameter estimates follows as:

$$C = - \left(E \frac{\partial^2 \ln L(p)}{\partial p \partial p'} \right) \Bigg|_{p=\hat{p}}, \quad (4.10)$$

where \hat{p} denotes the point-estimates for the unknown parameters. In the out-of-sample algorithm that we propose, this covariance matrix plays a crucial role since it reflects parameter uncertainty. We will use it to quantify the uncertainty in the VaR.

Denote the vector of parameters of the return model by p . The VaR of a position with a confidence level q and for N periods ahead follows as:

$$VaR_{q,N}(p) = W_0 \left\{ 1 - Q_{1-q} \left(\exp \sum_{t=1}^N r_t \right) \right\} \quad (4.11)$$

where W_0 is the initial value of the position, r_t denotes the continuously compounded return at time t . Because the parameter values in the underlying return generating process (and the uncertainty therein) have an impact on the VaR, we have made this relationship explicit by writing the VaR as a function of p . Finally, $Q_{1-q}(\cdot)$ denotes the quantile-operator that calculates the $(1-q)$ -th quantile of the random variable between parentheses. Usually an analytical expression for the quantile-operator is not available and we have to rely on simulation instead. A sample path of returns is generated by plugging the parameter estimates \hat{p} into the return model and by sampling from the error term. Consider D samples where $r_1^{(i)}, \dots, r_N^{(i)}$ denotes the i -th sample path of returns. Let $C^{(i)} = \exp \sum_{t=1}^N r_t^{(i)}$ denote the associated cumulative return. An estimate for the $(1-q)$ -th quantile follows by first sorting the D cumulative returns in ascending order, and then choosing the $(1-q)D$ -th element of this sorted series, denoted with \hat{Q}_{1-q} . The associated VaR follows easily from Equation (4.11).

Parameter uncertainty may be incorporated by sampling from the parameter distribution. Asymptotic distribution theory leads to the following distribution for the parameter estimates:

$$\hat{p} \sim N(p, C) \quad (4.12)$$

where \hat{p} are the parameter values for which the log-likelihood function is maximal, and C denotes the associated covariance matrix of the parameter estimates. In a Bayesian framework we sample from:

$$p \sim N(\hat{p}, C) \quad (4.13)$$

Consider M samples, which are denoted with $p^{(1)}, \dots, p^{(M)}$. For all these parameter values we calculate the VaR following the procedure outlined before. This leads to M values for the

VaR, denoted with $VaR_{q,N}(p^{(l)}), \dots, VaR_{q,N}(p^{(M)})$. So, instead of arriving at one VaR, we now have an entire sample of VaRs. The uncertainty in the VaR may be quantified by calculating the standard deviation of the VaRs. The expression of the VaR in Equation (4.11) shows that both the return distribution and the parameters may be treated as random variables that have impact on the size and the uncertainty in the reported VaR.

4.2.4 Backtesting

To test the adequacy of the return model for VaR-purposes, we propose an out-of-sample analysis in which the data series is split into two parts. Let T_1 denote the length of the first part. In the k -th sub-sample of the first part we use the observations r_k, \dots, r_{k+T_1} to determine the parameters of the return model, denoted by \hat{p}_k , and the associated covariance matrix, denoted by C_k . These suffice to calculate the magnitude of and the uncertainty in the VaR from the first part, denoted with $VaR_{q,N}^k(p_k^{(l)}), \dots, VaR_{q,N}^k(p_k^{(M)})$. This procedure may be repeated for other sub-samples $k = 1, \dots, K$. Associated with the k -th sub-sample is a second part of the data, which covers the observations r_{T_1+k}, \dots, r_T . From these returns we are able to calculate the actual change of the position in a period of length N , denoted as

$$\Delta_N^k = W_{T_1+k-1} \left(1 - \exp \sum_{i=1}^N r_{T_1+k-1+i} \right) \quad (4.14)$$

By comparison of the reported VaR with the actual change in the position for a particular sub-sample, we may calculate the number of violations of the reported VaR. Because of uncertainty in the reported VaR, the number of violations is different when we use the average reported VaR instead of using some worst-case VaR. Suppose that for all sub-samples we choose the l -th quantile to represent the VaR for the particular sub-sample. Define the series:

$$V_{q,N}^{k,l} = Q_l(VaR_{q,N}^k(p)) \quad k = 1, \dots, K \quad (4.15)$$

where $Q_l(\cdot)$ denotes the quantile operator that determines the l -th quantile from the random variable between parentheses. Let the total number of violations associated with the series in Equation (4.15) be calculated as:

$$v^j \equiv \sum_{k=1}^K I\{\Delta_N^k > V_{q,N}^{k,j}\} \quad (4.16)$$

where $I\{\cdot\}$ is the indicator function. Under the null-hypothesis of a probability of $(1-q)$ that the reported VaR is violated by the actual change in positions, the power of the test follows from

$$f(v^j) = \binom{K}{v^j} (1-q)^{v^j} q^{K-v^j} \quad (4.17)$$

which assumes a binomial distribution for the number of violations. This allows us to test specifically whether a particular quantile of the VaR provides an adequate representation. Also it may lead to an adjustment of the best estimate VaR in order to account for uncertainty in parameters. Two choices are important for the determination of the test-statistic. First, the length of the first sub-sample, T_1 , is arbitrary. A shorter sub-period leads to more parameter uncertainty and hence to more uncertainty in the reported VaR. Second the number of sub-samples that are considered in the moving window procedure, K , is arbitrary. For the sample statistic to converge to the binomial distribution function many sub-samples are required.

4.3 Data

The data set consists of daily prices of foreign currencies in terms of the US dollar. The currencies include the Deutschmark (DM), the British Pound (BP), the Japanese Yen (JY) and the Swiss Franc (SF). The time span includes the period from January 1986 to September 1999, a total of 3,445 observations. The data are obtained from Datastream. The raw exchange rates are transformed into continuously compounded returns, according to:

$$r_t = \ln\left(\frac{W_t}{W_{t-1}}\right) \quad (4.18)$$

where W_t denotes the exchange rate at time t . Summary statistics for the exchange rate returns are given in Table 1. All currencies exhibit fat tails and except for the Japanese Yen no significant skewness as may be observed from the reported kurtosis and skewness estimates.

Table 4.1. Summary Statistics

	DM	BP	JY	SF
Mean	-0.008	-0.004	-0.018	-0.009
Std Dev	0.698	0.669	0.762	0.782
Minimum	-3.505	-4.221	-7.685	-3.707
Maximum	3.457	4.286	4.139	3.779
Skewness	0.003	0.182	-0.698	-0.085
Kurtosis	5.085	6.399	9.869	4.840

Notes. The Table consists of summary statistics for the Deutschmark (DM), the British Pound (BP), the Japanese Yen (JY) and the Swiss Franc (SF). The data set consists of daily observations for the period January 1986 – September 1999. All currencies are expressed in US Dollars.

4.4 Results

4.4.1 Properties of alternative VaR Models

In this section we analyze the properties of the alternative VaR-models that are under consideration. We start with an in-sample analysis. Next, VaR-estimates are calculated given an initial position of $W_0 = 100$ at different confidence levels and for different time-spans. A measure of the reliability of the reported VaR is given with the VaR estimates. Finally, an out-of-sample analysis of the models is carried out, in which the estimated VaR is compared with realized changes in the currency position for different sub-periods. This leads to a test of how the alternative models perform in practice. In Tables 4.2 to 4.5 parameter estimates and associated standard errors are presented for the four models that are under consideration. Average return values (μ) are not significantly different from zero, and also for the autoregressive parameters (ρ) we cannot reject that they are zero. This is, of course, related to a well-known result in the international finance literature, see Meese and Rogoff (1983) and

Wolff (1987). From Tables 4.2 and 4.4 we find that exchange rates exhibit fat tails, since the estimates for the degrees-of-freedom (θ) are low with moderate standard errors. Table 4.5 also includes fat-tail behavior, given the parameter estimates for α . In Table 4.3, 4.4 and 4.5 also estimates for time-varying volatility have been included. Standard errors associated with parameter estimates for β_1 and β_2 suggest that exchange rates exhibit time-varying volatility.

A formal way to compare the in-sample performance of the alternative models is to perform a likelihood ratio test. Two times the difference in loglikelihood is compared with the critical value of a chi-squared distribution with the number of degrees of freedom given by the number of restrictions.

Table 4.2. Parameter Estimates AR(1)-Student-t Model

	DM	BP	JY	SF
μ	-0.001 (0.011)	-0.011 (0.009)	0.008 (0.011)	0.000 (0.012)
ρ	0.018 (0.016)	0.023 (0.015)	0.025 (0.016)	0.023 (0.016)
θ	4.831 (0.298)	3.925 (0.166)	3.850 (0.155)	5.151 (0.351)
LnL	-3518.8	-3277.6	-3668.1	-3932.9

Notes. The Table reports parameter estimates for the Student-t Model with constant volatility. Standard errors are given within parentheses.

Table 4.3. Parameter Estimates AR(1)-GARCH(1,1)-N Model

	DM	BP	JY	SF
μ	-0.003 (0.011)	-0.012 (0.010)	-0.010 (0.012)	-0.003 (0.013)
ρ	0.004 (0.017)	0.000 (0.018)	0.003 (0.018)	0.001 (0.018)
β_1	0.941 (0.010)	0.948 (0.007)	0.933 (0.013)	0.939 (0.012)
β_2	0.042 (0.006)	0.043 (0.008)	0.044 (0.008)	0.038 (0.007)
LnL	-3533.9	-3286.1	-3808.1	-3968.0

Notes. The Table reports parameter estimates for the GARCH(1,1)-N Model. Standard errors are given within parentheses.

Table 4.4. Parameter Estimates AR(1)-GARCH(1,1)-Student-t Model

	DM	BP	JY	SF
μ	0.001 (0.010)	-0.014 (0.009)	0.011 (0.010)	0.005 (0.012)
ρ	0.018 (0.016)	0.025 (0.016)	0.032 (0.016)	0.022 (0.016)
β_1	0.947 (0.011)	0.960 (0.008)	0.925 (0.016)	0.955 (0.010)
β_2	0.039 (0.007)	0.035 (0.006)	0.051 (0.010)	0.032 (0.006)
θ	5.669 (0.455)	4.781 (0.296)	4.259 (0.211)	5.702 (0.452)
lnL	-3450.3	-3148.0	-3587.6	-3882.7

Notes. The Table reports parameter estimates for the GARCH(1,1)-Student-t Model. Standard errors are given within parentheses.

Table 4.5. Parameter Estimates AR(1)-Power-GARCH(1,1)-Lévy-stable Model

	DM	BP	JY	SF
μ	0.000 (0.011)	-0.016 (0.009)	0.011 (0.010)	0.004 (0.012)
ρ	0.016 (0.017)	-0.019 (0.016)	0.032 (0.016)	0.021 (0.017)
β_0	0.006 (0.003)	0.001 (0.001)	0.008 (0.003)	0.005 (0.002)
β_1	0.946 (0.011)	0.960 (0.007)	0.926 (0.016)	0.959 (0.009)
β_2	0.030 (0.007)	0.020 (0.006)	0.031 (0.008)	0.022 (0.005)
α	1.803 (0.029)	1.738 (0.028)	1.723 (0.027)	1.803 (0.027)
δ	1.141 (0.298)	1.475 (0.235)	1.350 (0.236)	1.211 (0.266)
lnL	-3474.7	-3170.4	-3605.6	-3904.1

Notes. The Table reports parameter estimates for the Power-GARCH(1,1)-Lévy-stable Model. Standard errors are given within parentheses.

Comparing the GARCH(1,1)-Student-t model with the Student-t model we reject the restriction of constant volatility and prefer the GARCH(1,1)-t model that allows volatility to vary over time. Testing the GARCH(1,1)-Student-t model against the GARCH(1,1)-N model results in a preference for the former model. Testing the power-GARCH(1,1)-Lévy-stable

model against the GARCH(1,1)-N model results in a preference for the former model. These in-sample model comparisons suggest that the GARCH(1,1)-Student-t and the power-GARCH(1,1)-Lévy-stable models are preferred over restrictive versions of these models. Because the latter two models are non-nested their in-sample performances cannot be compared along these lines.

4.4.2 Uncertainty in the reported VaR

Models that provide the best in-sample fit do not necessarily lead to the best VaR-estimates. In Tables 4.6, 4.7, 4.8 and 4.9 VaR-estimates are given for different forecasting horizons ($N = 5, 10, 20$ days) and for different confidence levels ($q = 90\%, 95\%, 99\%$).

Table 4.6. VaR Results AR(1)-Student-t Model

N	q	DM	BP	JY	SF
5	99%	4.87 (0.33)	5.22 (0.41)	5.82 (0.46)	5.22 (0.34)
	95%	3.23 (0.15)	3.35 (0.19)	3.70 (0.19)	3.50 (0.18)
	90%	2.46 (0.12)	2.52 (0.12)	2.77 (0.14)	2.66 (0.13)
10	99%	6.70 (0.41)	7.04 (0.49)	7.91 (0.65)	7.28 (0.53)
	95%	4.57 (0.21)	4.74 (0.23)	5.24 (0.27)	4.92 (0.26)
	90%	3.54 (0.16)	3.66 (0.19)	3.97 (0.19)	3.79 (0.22)
20	99%	9.22 (0.55)	9.75 (0.58)	10.89 (0.76)	9.96 (0.58)
	95%	6.46 (0.31)	6.77 (0.33)	7.32 (0.41)	6.97 (0.38)
	90%	5.05 (0.27)	5.30 (0.32)	5.64 (0.33)	5.42 (0.33)

Notes. The Table reports VaR-results for different horizons (N in days) and at different confidence levels (q). The expected VaR is reported. The uncertainty in the VaR is reflected by its standard deviation, which is denoted within parentheses.

Table 4.7. VaR Results AR(1)-GARCH(1,1)-N Model

<i>N</i>	<i>q</i>	DM	BP	JY	SF
5	99%	3.66 (0.23)	2.88 (0.16)	4.72 (0.22)	4.17 (0.24)
	95%	2.54 (0.13)	2.01 (0.10)	3.30 (0.15)	2.91 (0.14)
	90%	1.97 (0.10)	1.56 (0.09)	2.57 (0.13)	2.27 (0.12)
10	99%	5.16 (0.31)	4.10 (0.22)	6.62 (0.39)	5.90 (0.36)
	95%	3.61 (0.19)	2.88 (0.18)	4.64 (0.21)	4.09 (0.23)
	90%	2.78 (0.16)	2.25 (0.15)	3.62 (0.21)	3.19 (0.18)
20	99%	7.19 (0.51)	5.90 (0.38)	9.18 (0.50)	8.25 (0.46)
	95%	5.05 (0.33)	4.14 (0.28)	6.42 (0.37)	5.76 (0.36)
	90%	3.93 (0.30)	3.27 (0.24)	5.04 (0.31)	4.49 (0.30)

Notes. The Table reports VaR-results for different horizons (*N* in days) and at different confidence levels (*q*). The expected VaR is reported. The uncertainty in the VaR is reflected by its standard deviation, which is denoted within parentheses.

Table 4.8. VaR Results AR(1)-GARCH(1,1)-Student-t Model

<i>N</i>	<i>q</i>	DM	BP	JY	SF
5	99%	4.86 (0.37)	3.96 (0.34)	7.01 (0.55)	5.37 (0.41)
	95%	3.15 (0.17)	2.53 (0.15)	4.28 (0.23)	3.46 (0.21)
	90%	2.39 (0.12)	1.91 (0.11)	3.17 (0.15)	2.63 (0.15)
10	99%	6.89 (0.56)	5.67 (0.43)	10.11 (0.97)	7.53 (0.57)
	95%	4.55 (0.27)	3.71 (0.21)	6.27 (0.38)	4.99 (0.27)
	90%	3.46 (0.21)	2.85 (0.17)	4.62 (0.28)	3.78 (0.21)
20	99%	10.13 (0.91)	8.63 (0.76)	15.09 (1.60)	10.82 (0.82)
	95%	6.60 (0.43)	5.60 (0.38)	9.24 (0.61)	7.19 (0.45)
	90%	5.05 (0.34)	4.27 (0.29)	6.86 (0.39)	5.47 (0.36)

Notes. The Table reports VaR-results for different horizons (*N* in days) and at different confidence levels (*q*). The expected VaR is reported. The uncertainty in the VaR is reflected by its standard deviation, which is denoted within parentheses.

Table 4.9. VaR Results AR(1)-Power-GARCH(1,1)-Lévy-stable Model

<i>N</i>	<i>q</i>	DM	BP	JY	SF
5	99%	4.77 (1.24)	4.02 (1.00)	6.69 (1.37)	5.12 (1.00)
	95%	2.71 (0.61)	2.07 (0.40)	3.32 (0.42)	2.94 (0.47)
	90%	2.00 (0.45)	1.53 (0.29)	2.39 (0.30)	2.19 (0.34)
10	99%	7.07 (1.93)	5.97 (1.50)	9.95 (2.08)	7.58 (1.53)
	95%	3.97 (0.90)	3.13 (0.60)	4.90 (0.63)	4.31 (0.70)
	90%	2.95 (0.70)	2.32 (0.44)	3.50 (0.44)	3.21 (0.52)
20	99%	10.68 (2.91)	9.10 (2.28)	14.71 (3.03)	11.22 (2.39)
	95%	5.87 (1.39)	4.79 (0.89)	7.22 (1.05)	6.40 (1.05)
	90%	4.38 (1.05)	3.53 (0.64)	5.12 (0.70)	4.76 (0.80)

Notes. The Table reports VaR-results for different horizons (*N* in days) and at different confidence levels (*q*). The expected VaR is reported. The uncertainty in the VaR is reflected by its standard deviation, which is denoted within parentheses.

Parameter uncertainty leads to uncertainty in the reported VaR. This uncertainty is quantified by reporting the standard deviation of the VaR within parentheses. In general the VaR is higher at higher confidence levels and also for longer forecasting periods. Comparison of the reported VaRs that result from the different models shows two effects. First, there is a level effect: because models with Lévy-stable or Student-t distributed error terms include fatter tail-specifications, the reported VaR is also higher. The VaR may even be overstated, because the slow power law decay of the fitted Lévy-stable or Student-t distribution implies extreme values never observed in financial data. Second, because of greater uncertainty in the parameter estimates that account for the fat-tail behavior (θ and α), the uncertainty in the reported VaR is also greater for these models.

For the Power-GARCH(1,1)- Lévy-stable model, an additional source of uncertainty is the δ -exponent in the volatility model. A huge amount of data is necessary to obtain a precise estimate for the δ -exponent. This leads to more parameter uncertainty and hence to more uncertainty in the associated VaR. The empirical results suggest that more sophisticated tail-modeling approaches come at the cost of more uncertainty about the VaR estimate itself.

4.4.3 Out-of-Sample Violations

In the last part of the empirical analysis we focus on the out-of-sample behavior of the alternative VaR-models. The part of the data set that is used to estimate the alternative VaR-models always have the same length of $T_1 = \frac{1}{2}T = 1722$. Given the parameter estimates the VaRs are calculated for out-of-sample periods of $N = 5, 10, 20$ days. Parameter uncertainty is also taken into account: from the different VaR-values that arise because of parameter uncertainty, we save the average VaR, and an upper- and a lower-bound equal to the average VaR plus or minus twice the standard deviation of the VaR, respectively. This procedure is repeated for $K = 1700$ sub-samples that appear as moving windows. In Tables 4.10, 4.11 and 4.12 we report the percentage of violations in the out-of-sample period.

Table 4.10. Out-of-sample Violations AR(1)-GARCH(1,1)-N Model

<i>N</i>	<i>VaR</i>	DM	BP	JY	SF
5	upper	2.1%	1.9%	3.5%	2.5%
	average	1.2%	1.1%	2.1%	1.6%
	lower	0.8%	0.5%	1.6%	1.3%
10	upper	1.4%	1.4%	3.4%	2.8%
	average	0.9%	0.8%	2.6%	1.8%
	lower	0.4%	0.4%	2.1%	0.8%
20	upper	2.2%	1.5%	5.6%	2.6%
	average	1.2%	1.1%	3.7%	1.5%
	lower	0.5%	0.8%	2.5%	0.9%

Notes. The Table reports the percentage of violations of the reported VaR at a 99% confidence level. Different horizons (*N* in days) and different VaR (average, upper bound and lower bound) are reported.

Table 4.11. Out-of-sample Violations AR(1)-GARCH(1,1)-Student-t Model

<i>N</i>	<i>VaR</i>	DM	BP	JY	SF
5	upper	1.3%	1.1%	1.4%	1.4%
	average	0.7%	0.6%	0.9%	0.9%
	lower	0.5%	0.4%	0.5%	0.2%
10	upper	1.5%	1.5%	1.8%	1.1%
	average	0.7%	0.6%	1.1%	0.5%
	lower	0.5%	0.5%	0.5%	0.1%
20	upper	2.1%	1.7%	1.5%	1.4%
	average	1.0%	0.9%	0.0%	0.4%
	lower	0.7%	0.6%	0.0%	0.1%

Notes. The Table reports the percentage of violations of the reported VaR at a 99% confidence level. Different horizons (*N* in days) and different VaR (average, upper bound and lower bound) are reported.

Table 4.12. Out-of-sample Violations AR(1)-Power-GARCH(1,1)- Lévy-stable Model

<i>N</i>	<i>VaR</i>	DM	BP	JY	SF
5	upper	1.9%	1.7%	2.4%	2.2%
	average	0.6%	0.5%	1.1%	0.9%
	lower	0.1%	0.1%	0.4%	0.2%
10	upper	1.8%	1.9%	2.5%	2.2%
	average	0.3%	0.4%	1.3%	0.5%
	lower	0.1%	0.0%	0.0%	0.1%
20	upper	1.9%	2.0%	3.3%	2.2%
	average	0.4%	0.3%	0.1%	0.4%
	lower	0.1%	0.0%	0.0%	0.1%

Notes. The Table reports the percentage of violations of the reported VaR at a 99% confidence level. Different horizons (*N* in days) and different VaR (average, upper bound and lower bound) are reported.

We start with an analysis of the violations without taking parameter uncertainty into account. This means that we only consider the number of violations of the average reported VaR. The GARCH(1,1)-N model underestimates the actual VaR on a 99 % confidence level for the Japanese Yen and the Swiss Franc for weekly, biweekly and monthly forecasts, while it provides adequate results for the Deutschmark and the British Pound. It is well known that the normal distribution underestimates large events and the result is not surprising in the case of the Japanese Yen and the Swiss Franc. The GARCH(1,1)-Student-t model in general overestimates the actual VaR except for the Japanese Yen at weekly and biweekly forecasts. For monthly forecasts it simply reports no violation on average. The overestimation of the GARCH(1,1)-Student-t model for 4 out of 5 exchange rates is a result of the power law tails of the Student-t distribution. For exchange rate data the really extreme events are rare. As a result the degrees of freedom are fitted to the near tails of the empirical distribution. Because of the functional specification, really extreme returns are generated much more frequent than available in the historical data. As a result these extreme returns are the source for the overestimation of the actual VaR especially for monthly forecasts. The same is true for the power-GARCH(1,1)-Lévy-stable model and the results are similar. Based on the average reported VaR we conclude that the GARCH(1,1)-N model leads to an underestimation of the greatest possible loss in the tail because tail behavior is not modeled adequately. The GARCH(1,1)-Student-t and power-GARCH(1,1)-Lévy-stable models overestimate the risk in the tail. This is due to the fact that parameter estimates do not adequately represent the behavior in the far tail due to lack of extreme historical observations.

More parameter uncertainty leads to wider confidence intervals for the reported VaR. In particular the GARCH(1,1)-Student-t and power-GARCH(1,1)-Lévy-stable models lead to wide confidence intervals. See Figure 4.1, 4.2, 4.3 and 4.4 for the back-testing results of the various models. As an example we plot the weekly 99% VaR estimates for the Japanese Yen and the 95% confidence intervals associated with parameter uncertainty.

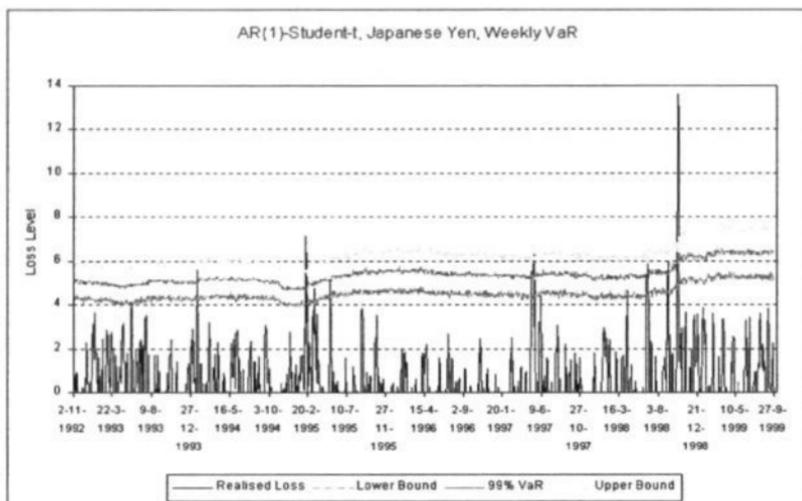


Figure 4.1: AR(1)-Student-t VaR Estimates

The graph depicts how the forecasts of the VaR estimates, using the Student-t distribution, compare to realised losses over a period of 5 trading days. We also plot the 95% confidence level associated with parameter uncertainty. We have used rolling observations of daily data, over the period November 1992 until September 1999, to provide forecasts of the Value-at-Risk at the 99% confidence level.

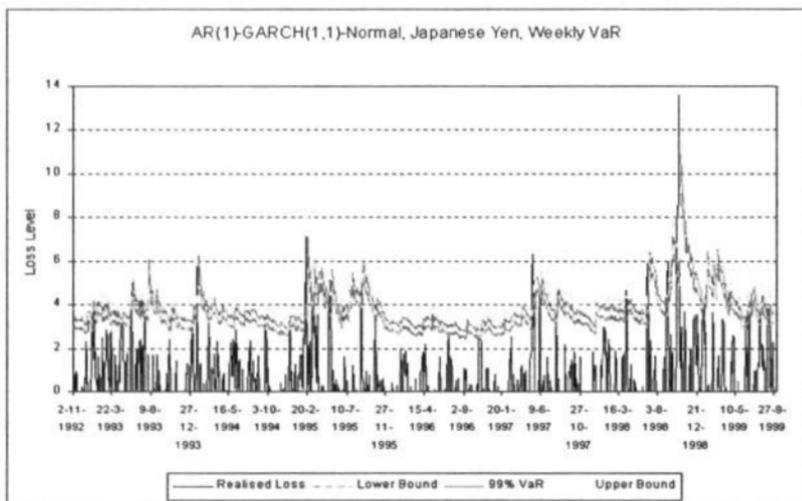


Figure 4.2: AR(1)-GARCH(1,1)-Normal VaR Estimates

The graph depicts how the forecasts of the VaR estimates, using a GARCH model driven by a Normal distribution, compare to realised losses over a period of 5 trading days. We also plot the 95% confidence level associated with parameter uncertainty. We have used rolling observations of daily data, over the period November 1992 until September 1999, to provide forecasts of the Value-at-Risk at the 99% confidence level.

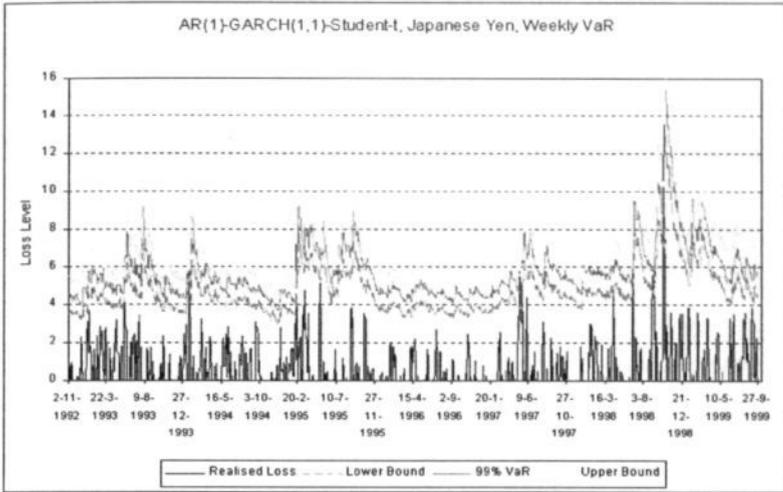


Figure 4.3: AR(1)-GARCH(1,1)-Student-t VaR Estimates

The graph depicts how the forecasts of the VaR estimates, using a GARCH model driven by a Student-t distribution, compare to realised losses over a period of 5 trading days. We also plot the 95% confidence level associated with parameter uncertainty. We have used rolling observations of daily data, over the period November 1992 until September 1999, to provide forecasts of the Value-at-Risk at the 99% confidence level.

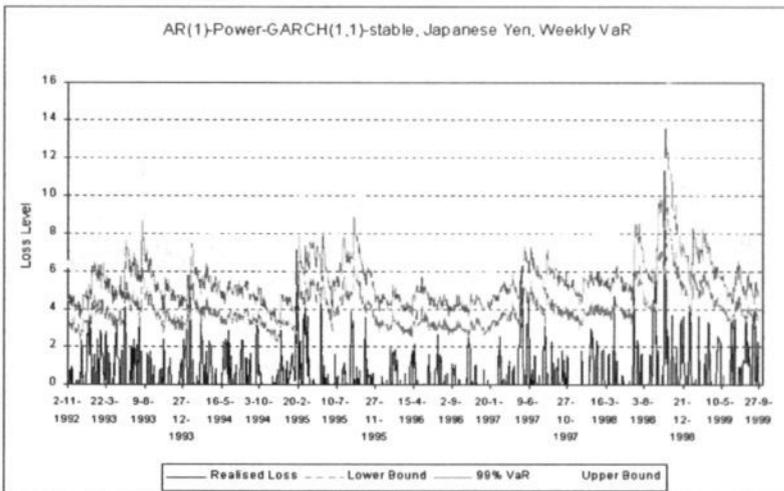


Figure 4.4: AR(1)-Power-GARCH(1,1)-Lévy-stable VaR Estimates

The graph depicts how the forecasts of the VaR estimates, using a GARCH model driven by a Lévy-stable distribution, compare to realised losses over a period of 5 trading days. We also plot the 95% confidence level associated with parameter uncertainty. We have used rolling observations of daily data, over the period November 1992 until September 1999, to provide forecasts of the Value-at-Risk at the 99% confidence level.

The results suggests that it is costly to adopt models with more complex tail-properties, because it leads to more uncertainty in the reported VaR. Parameter uncertainty also explains the overestimation of the reported VaR in case of the GARCH(1,1)-Student-t. Instead of focusing on the average VaR, we may also give more attention to the upper-bound. This has the effect that we acknowledge that tails are less fat than the point estimates would imply. This leads to rejection rates that come closer to the confidence levels of the VaR. In case of the power-GARCH(1,1)-Lévy-stable model there is so much parameter uncertainty that the 95% upper-bound of the reported VaR leads to an underestimation. More sophisticated risk adjustments are required in this case.

4.5 Concluding Remarks

In this chapter we have focused on VaR-model selection for exchange rate positions. On the one hand fat-tail behavior is present in exchange rate returns, but on the other hand there are very few observations far in the tail. Models that take account of tail behavior are required, since otherwise the reported VaR lead to an under-estimation of the risk in the tail. Complex tail models often lead to over-estimation of the VaR, because these models assume more probability mass in the tail of the distribution than is actually present. This is due to the fact that very few extreme observations occur and hence tail behavior is measured with relative great uncertainty. For the GARCH(1,1)-Student-t distribution, taking this adjustment into account leads to rejection rates that are close to the confidence levels of the reported VaR. Among the models analyzed, the GARCH(1,1)-Student-t specification seems to be an adequate model to correctly assess extreme losses for exchange rate positions. Models with more sophisticated tail behavior lead to more parameter uncertainty, which leads to greater uncertainty in the reported VaRs.

1. The first part of the text is a

2. The second part of the text is a

3. The third part of the text is a

4. The fourth part of the text is a

5. The fifth part of the text is a

6. The sixth part of the text is a

7. The seventh part of the text is a

8. The eighth part of the text is a

9. The ninth part of the text is a

10. The tenth part of the text is a

CHAPTER 5

Implied Volatility Forecasting¹²

5.1 Introduction

Recent turbulences in financial markets worldwide have highlighted once again that uncertainty is an unavoidable aspect of financial markets. With good reason the assessment of uncertainty, often measured as volatility, plays a prominent role in all areas of finance, ranging from investment decision making, security valuation and risk management to monetary policy making. Even though the measurement and forecasting of volatility has attracted the interest of many researchers and practitioners, it remains a challenging statistical problem. Not only is it a problem of having a proper volatility model, but also of having a robust volatility forecasting method. The available models, such as GARCH or stochastic volatility, based on historical returns seem to work quite well in-sample but generally perform poorly out-of-sample (Akgiray (1989), Dimson and Marsh (1990), Nelson (1992), Nelson and Foster (1995), Franses and Van Dijk (1995), and Brailsford and Faff (1996)), with a predictive ability often below 12%.

¹² This chapter is based on: C. de Jong, T. Lehnert, 2001, 'Implied Volatility Forecasting', LIFE Working Paper, Maastricht University.

5.1.1 The Information Content of Implied Volatilities

In recent years there is some support for the hypothesis that the information provided by implied volatilities from daily option prices is more relevant in forecasting volatility than the volatility information provided by historical returns (Jorion (1995), Christensen and Prabhala (1998), Fleming (1998), and Blair et al. (2001)). If option markets are efficient, option prices contain information about the expected price process of the underlying asset over the lifetime of the option. For the purpose of volatility forecasting this comes down to extracting the expectation of market participants about the development of future volatility. Therefore, the volatility measure derived from option prices is a forward looking (risk-neutral) measure and eventually different from a measure based on historical return data. In some sense option prices make the unobservable expected volatility observable and also eliminate the choice of a particular historical sample period. As a result, the information in traded options may result in better volatility forecasts.

Although the idea of option-implied volatility estimates is relatively simple, there is not one straightforward method to extract the information. Every proposed method relies on a number of assumptions regarding the model underlying option prices. In line with the large number of option pricing models, academics and practitioners have applied a large number of methods to extract option-implied volatilities. After the crash in 1987 practitioners seem to adjust the volatility for moneyness and maturity before plugging it into the Black-Scholes model. As a result, we can observe an implied volatility pattern that cannot be fully explained by the recent option pricing literature using historical returns of the underlying. It is known that implied volatility covaries with realized volatility, but the major difficulty is to back out volatility information of the underlying from observed option prices. Since the assumptions of the Black-Scholes option-pricing framework are usually violated, it is a challenge to select the appropriate implied volatility. Previous studies try to explore information from (Black-Scholes) implied volatilities of traded options to estimate and forecast future volatility without explicitly modeling the underlying return process. A number of researchers have been extensively investigating the optimal weighting scheme for the different implied volatilities at different strikes¹³. A measure often applied is VIX, which is an S&P100 volatility index that combines a number of close-to-at-the-money implied volatilities into a single measure (see Fleming et al. (1995) for a description). The index is constructed in such a way that it represents the implied volatility of a hypothetical at-the-money option with 22 days to

maturity, and is therefore unable to capture any smile or term structure effects of the whole implied volatility surface. Those methods that adhere to the Black-Scholes implied volatility framework are not only arbitrary and theoretically questionable, but also the result often turns out to be a severely biased predictor of future volatility (Fleming (1998) and Blair et al. (2001)).

5.1.2 A 'New' Measure of Local Volatility

In this chapter we present a new method to make volatility forecasts, which is based on a relatively recent set of option-pricing models that applies the time-series GARCH-methodology to option pricing (see Duan (1995), Kallsen and Taqqu (1998), Ritchken and Trevor (1999), Bauwens and Lubrano (2000) and Heston and Nandi (2000)). We construct the expected future price process by deriving the relevant parameters of the GARCH option-pricing model from prices of traded options with different strike and maturity. In contrast to methods for estimating and forecasting volatility that use past index returns, this method derives anticipated parameters of a GARCH process and therefore market expectations about the future price process. Since in most finance applications of volatility we need a forward looking measure, the characteristics of the future price process should be more informative than the ones of the historical price process.

There are now several GARCH option pricing models available in the literature, but a very flexible one is the GARCH option-pricing model of Duan (1995). It has shown some empirical success and it is appropriate for our study. We estimate the parameters of the model by minimizing the relative pricing error between the market prices and the theoretical option prices. Once we have the GARCH parameter estimates we can use Monte Carlo simulations to make volatility forecasts some periods ahead. As a comparison, we construct a second volatility estimate using the same GARCH specification, but now calibrated with historical returns. In an out-of-sample analysis we compare our 'implied GARCH' model with the 'historical GARCH' and conclude which measure is superior in making one-day ahead forecasts of the volatility of a market index. In line with recent literature on volatility measurement we use intraday data to calculate the daily realized volatility that serves as a benchmark for our forecasts.

¹³ See Bates (1996) for a review.

We make several contributions to the existing volatility modeling and forecasting literature. First, we use a completely new measure based on the informational content of option prices. With the ever-increasing trading volumes in derivative markets, we believe the information in derivatives will become the standard for making volatility forecasts. Second, we use a data set (DAX index options) that has deserved relatively limited attention by researchers but has become large in trading volume. This limited attention for our dataset has the advantage that it is largely independent of the previously done volatility research that was mainly directed to S&P100 index options.

In the next section we set up the econometric framework; in Section 5.3 we describe the data; Section 5.4 provides empirical results and Section 5.5 concludes.

5.2 Model specifications

5.2.1 The GARCH Option Pricing Model

In this section we focus on the GARCH option pricing framework developed in Duan (1995) and implement a model based on the exponential GARCH (EGARCH) of Nelson (1991). Apart from the well documented GARCH effects (see Bollerslev et al. (1992)), this process is also able to model the well-known leverage effect of stock market returns (Nelson, 1991). Volatility tends to rise in response to bad news (negative excess returns) and to fall in response to good news (positive excess returns). The form of the EGARCH specification is comparable to the non-linear asymmetric GARCH process of Engle and Ng (1993), the GJR-GARCH model of Glosten et al. (1993) and the power GARCH of Ding et al. (1993). Those studies show it is crucial to include the asymmetric term in financial time series models.

The EGARCH option pricing model assumes the risk-neutral valuation principle, and the standard Black and Scholes (1973) model can be derived as a special case. Kallsen and Taquq (1998) develop a continuous-time version of the model and show that the same pricing results can be derived via an arbitrage-free argument. Heston and Nandi (2000) develop a closed form solution of a GARCH option-pricing model. There are many alternative option-pricing models available in the literature. Examples are the bivariate diffusion model of Hull and White (1987), the jump-diffusion model of Naik and Lee (1990), the variance-gamma

model of Madan and Milne (1991) and the stochastic volatility models of Stein and Stein (1991), Wiggins (1991) and Heston (1993).

Our choice of the GARCH option pricing model of Duan (1995) was motivated by its flexibility, the recent empirical successes of the model (see among others Amin and Ng (1994), Heynen et al. (1994), Duan (1996), and Ritchken and Trevor (1999)) and the emerging availability of numerical methods for this class of option pricing models (see Hanke (1997), Duan et al. (1998), Ritchken and Trevor (1999), Duan and Simonato (1998a&b), Heston and Nandi (2000), and Duan et al. (2001)).

For this kind of derivative valuation models with a high degree of path dependency, computationally demanding Monte Carlo simulations are commonly used for valuing derivative securities. We use the recently proposed simulation adjustment method, the empirical martingale simulation (EMS) of Duan and Simonato (1998a), which has been shown to substantially accelerate the convergence of Monte Carlo price estimates and to reduce the so called 'simulation error'. The EMS reproduces the martingale property for the simulated sample, a characteristic of all derivative pricing models.

As a first step in our empirical analysis we derive the spot rate $X_{t,0}$ from futures prices with different maturities using spot-futures parity. Given a futures price $F_{t,T}$ with time to maturity $T-t$, we use the Equation

$$F_{t,T} = e^{r_{t,T}(T-t)} X_{t,T} \quad (5.1)$$

to derive the corresponding spot rate $X_{t,T}$. We use the transaction prices of the two futures contracts with the shortest time to maturity and determine the price of a theoretical futures contract with a remaining lifetime of 0 days, which is assumed to be the underlying index level¹⁴.

In a discrete-time economy the value of the index at time t can be assumed to follow the following dynamics:

¹⁴ See e.g. the appendix in Poteshman (2001) for details. A similar approach is used by Deutsche Börse within their VDAX framework (see Deutsche Börse (2002), p.13).

$$r_t = \ln\left(\frac{X_t}{X_{t-1}}\right) = \mu_t + \sigma_t e_t$$

$$e_t | \Omega_{t-1} \sim N(0,1) \text{ under probability measure P} \quad (5.2)$$

$$\ln(\sigma_t^2) = \omega + \alpha \ln(\sigma_{t-1}^2) + \beta(|e_{t-1}| - \gamma e_{t-1})$$

where μ_t represents the conditional mean; Ω_{t-1} is the information set in period t-1 and the combination of β and γ captures the leverage effect. Daily returns of financial time series may exhibit non-zero autocorrelation. One can account for this effect by specifying the conditional mean as an autoregressive process (Bauwens and Lubrano (2000) and Hafner and Herwatz (2001)) discuss how this affects option prices) or by allowing for a risk premium attached to time-varying volatility. The specification for the conditional mean we selected includes the second alternative advocated by Duan (1995, 1999) and Heston and Nandi (2000):

$$\mu_t = r_f + \lambda \sigma_t - \frac{1}{2} \sigma_t^2, \quad (5.3)$$

where the risk premium λ is a constant parameter, and the term $-\frac{1}{2} \sigma_t^2$ gives additional control for the conditional mean. This specification completes the baseline EGARCH model that we

used for the analysis. The long run stationary volatility is equal to $\sqrt{\exp\left(\frac{\omega + \beta \sqrt{\frac{2}{\pi}}}{1 - \alpha}\right)}$. The

parameter α measures the degree of mean reversion in that $\alpha=1$ implies that the variance process is integrated. We also tried alternative specifications for the volatility dynamics, but for the 'implied' GARCH calibration we experienced frequent violations of the covariance stationary condition and if we control for covariance stationarity the fit was sometimes extremely bad. In contrast, restriction of the mean reversion parameter α in the EGARCH specification to values below 1 did not cause notable problems in the estimation process.

Duan (1995) shows that under the Local Risk Neutral Valuation Relationship (LRNVR) the conditional variance remains unchanged, but under the pricing measure Q the conditional expectation of r_t is equal to the risk free rate r_f . Therefore, risk neutralization transforms the error term in the following way:

$$r_t = r_f - \frac{1}{2}\sigma_t^2 + \sigma_t \varepsilon_t$$

$$\varepsilon_t | \Omega_{t-1} \sim N(0,1) \text{ under risk-neutralized probability measure } Q \quad (5.4)$$

$$\ln(\sigma_t^2) = \omega + \alpha \ln(\sigma_{t-1}^2) + \beta(|\varepsilon_{t-1} - \lambda| - \gamma(\varepsilon_{t-1} - \lambda))$$

In the Equations above ε_t is not necessarily normal, but to include the Black-Scholes model as a special case we assume that ε_t is a Gaussian random variable. The shift of the error term can be interpreted as an additional modification of the news impact curve, therefore also modifies the asymmetry in the volatility process. The long run stationary volatility level can be shown to be equal to $\sqrt{\exp\left(\frac{\omega + \beta E[|\varepsilon - \lambda| + \gamma(\varepsilon - \lambda)]}{1 - \alpha}\right)}$, in which the expected value should be evaluated numerically.

A European call option with exercise price K and maturity T has at time t price equal to:

$$c_t = e^{-r_f(T-t)} E_t^Q [\max(X_T - K, 0) | \Omega_t] \quad (5.5)$$

We rely on Monte Carlo simulations to evaluate the option numerically. Given the value of the index X_t , we generate N standard normal random numbers to advance the dynamics one period ahead and then make the empirical martingale adjustment. We repeat this procedure $T-t$ times until maturity and arrive at N simulated prices X_T . We calculate each of the N option payoffs, take the average and discount them back to period t of option valuation. Using this procedure we compute the value of an option for all exercise prices and all maturities.

5.2.2 Model Calibration

We calibrate the parameters of the EGARCH option-pricing model in (5.5) by minimizing the square root of the mean squared relative pricing error between the market prices and the theoretical call and put option prices:

$$\text{e.g. } RMSE = \sqrt{\frac{1}{M} \min \sum_{i=1}^n \sum_{j=1}^{m_i} \left(\frac{\hat{c}_{i,j} - c_{i,j}}{c_{i,j}} \right)^2}, \quad (5.6)$$

where M is the total number of call options evaluated, the subscript i refers to the n different maturities and subscript j to the m_i different strike prices in a particular maturity series i . We use relative instead of absolute pricing errors in order to give options with different levels of moneyness equal weight in the calibration process.

As starting values for the calibration we use the time-series estimates from the EGARCH model using approximately three years (752 trading days) of historical returns. Before we run the time-series calibrations, we first restrict one of the model parameters to ensure that the long run stationary volatility level $\bar{\sigma}$ will be equal to the relatively stable 5-year historical standard deviation:

$$\omega = (1 - \alpha) \ln(\hat{\sigma}^2) - \beta \sqrt{\frac{2}{\pi}} \quad \text{‘historical’ calibration} \quad (5.7)$$

After we calibrated the time-series estimates, we use two time-series parameter estimates for the option calibration: the long run volatility $\bar{\sigma}$ and the risk premium parameter λ . We do this, because the large number of parameters to be estimated makes the calibration process unstable: especially the joint identification of λ and γ is cumbersome, since both parameters control for asymmetry in the news impact curve. Our choice for fixing the stationary volatility is slightly different from that of Duan (1996, 1999), and Heston and Nandi (2000). They perform a constrained calibration in which the parameters λ and the local volatility are restricted to the time-series GARCH-estimates. We apply slightly different calibration restrictions in that we do not restrict the local volatility to its time-series equivalent, but the long run volatility. For the latter we use again the 5-year historical standard deviation. Since the long run stationary volatility ($\bar{\sigma}$) is not an explicit parameter in the model but can be derived directly from the other parameter estimates and vice versa, we constrain ω to:

$$\omega = (1 - \alpha) \ln(\bar{\sigma}^2) - \beta E[|\varepsilon_i - \lambda| + \gamma(\varepsilon_i - \lambda)] \quad \text{‘implied’ calibration} \quad (5.8)$$

We derive the local volatility from option prices, because time-series models have most problems in accurately forecasting short-run volatility fluctuations, whereas option prices can reflect new information instantaneously. In return, we constrain the long-run volatility to its time-series estimate because it turned out to be very unstable if estimated from option prices.

Our approach can now be summarized as follows. First, we use five years of historical returns to estimate long run volatility $\bar{\sigma}$. Then, we use three years of historical returns to estimate the time-series GARCH process in which the long run volatility is restricted to the five-year standard deviation. Finally, we use option prices to carry out a constrained calibration by restricting the risk premium parameter λ and the long run stationary volatility level $\bar{\sigma}$ to the estimates derived from the history of asset prices. The final calibration yields estimates of the parameters α , β , γ and local volatility σ_t .

Splitting the estimation of the parameters of our model in a ‘historical’ and an ‘implied’ part has some advantages: it is more likely that option prices contain information about the future, but for risk management purposes it would be misleading to ignore all the information contained in the history of asset prices. Therefore, the method readily exploits the combination of information about the times series (the volatility risk premium and the long run volatility) and the information about the volatility dynamics contained in option prices.

Given the parameters from the ‘historical’ calibration we use the EGARCH model under probability measure P in (5.2) to derive a volatility forecast $\hat{\sigma}_{t+1}^{TS}$. The estimated local volatility level $\hat{\sigma}_{t+1}^{Opt}$ from the ‘implied’ calibration is a one-day ahead forecast, which contains information about the expectation of market participants about tomorrow’s volatility.

Since actual volatility is a latent variable, we have to construct an accurate measure to evaluate our ‘historical’ and ‘implied’ forecasts. Because of its unobservability, different methods have been proposed to compute ex post estimates of realized volatility. The simplest and most common one is the square of realized return over the data. However, as Andersen and Bollerslev (1998) clearly point out theoretically, this measure produces very noisy estimates of the actual volatility because of the randomness in the return process. By sampling more frequently the randomness effect can be reduced. Theoretically, the realized volatility is a better estimate of volatility. Empirically, this is confirmed in our dataset, where we use 5-minute intraday returns to construct a volatility estimate, denoted by σ_t^{real} . This estimate is much more stable than squared returns (see also Andersen et al. (2001a and 2001b)). It is based on the 5-minute log-returns of the closest-to-maturity index future. We use a 5-minute interval, because that yields a relatively large number of returns per day without notable bid-ask bounce problems, and because it is the frequency that Andersen and Bollerslev (1998) propose. Our first index level is the opening price. All subsequent levels are the latest level closest to the 5-minute mark, ending with the closing price. Contrary to Blair et al. (2001) we

do not include any overnight return, because we believe that it is rather arbitrary what exact weight this return should get and then rather leave it out completely.

5.3 Data & Methodology

We use closing prices of DAX index options and transaction prices of DAX futures for a period from January 2000 until August 2001 and closing prices of DAX index futures for a period from January 1997 until August 2001. For the time-series analysis we need daily DAX index levels (derived from futures prices) for a period from January 1995 until August 2001. The raw data set is directly obtained from the EUREX, European Futures and Options Exchange¹⁵. A fully electronic exchange, EUREX was formally established in 1998 following the merger of DTB Deutsche Terminbörse (German Options and Futures Exchange) and SOFFEX (Swiss Options and Financial Futures Exchange). In June 2000 EUREX introduced a new software version of the electronic trading system and the trading hours were extended to 8:00 p.m. CET. The market for DAX index options and futures is the most active index options and futures market in Europe.

For index options the expiration months are the three nearest calendar months, the three following months within the cycle March, June, September and December, as well as the two following months of the cycle June, December. For index futures the expiration months are the three nearest calendar months within the cycle March, June, September and December. The last trading day is the third Friday of the expiration month, if that is an exchange trading day; otherwise it is on the exchange-trading day immediately prior to that Friday. Figure 5.1 shows the daily development of the DAX index futures from January 1997 till August 2001.

¹⁵ The author gratefully acknowledges Deutsche Börse Group for providing the data.

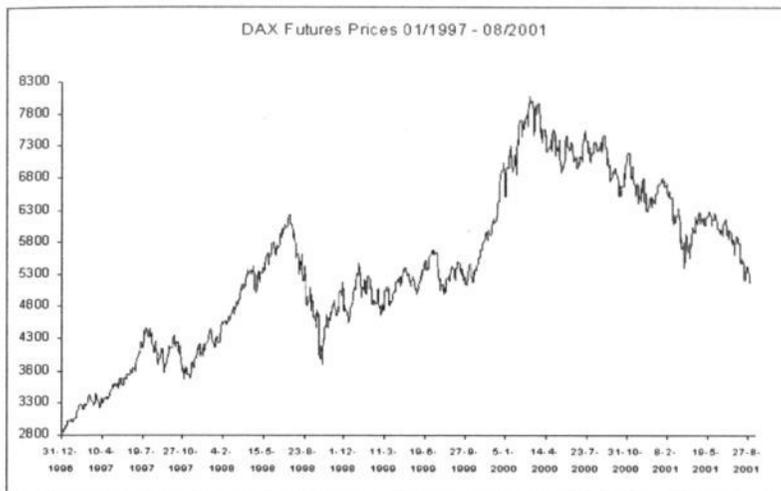


Figure 5.1: DAX Futures Prices

The graph shows the closest-to-maturity DAX index futures prices for the period January 1997 until August 2001.

The DAX index calculation is based on the assumption that the cash dividend payments are reinvested. Therefore theoretically we don't have to adjust the index level for the fact that the stock price drops on the ex-dividend date when we calculate option prices. But since the cash dividend payments are taxed the reinvestment does not fully compensate for the decrease in the stock price and empirically we observe a different implied underlying for different maturities. Therefore we always work with the adjusted underlying implied out from futures or option prices.

Before we run the 'implied' calibration, we compute the implied interest rates and implied index rates from the observed put and call option prices using essentially the method of Shimko (1993) based on put-call parity. We modified the standard put-call parity regression to ensure that the implied index value is a non-increasing function of the maturity of the option (see Duan (1996)). It can be shown that the put-call parity holds sufficiently well for our data. Alternatively, for the option calculations we could have chosen index levels from DAX index futures closing prices. Both methods would be equally safe since both markets are closely integrated.

For the 'implied' calibration, we estimate our model using the closing prices of traded call and put options every day in our sample period. We exclude options with less than 5 and more than 75 trading days until maturity. Furthermore, instead of using a static rule and exclude options with absolute moneyness of more than 10% (see Dumas et al., 1998), we use

a volume rule and exclude options with a daily Euro turnover of less than 10,000 Euro and a price of less than 2 Euro. The choice of these particular filter rules can be motivated as follows. Since we are interested in short term volatility forecasting of the underlying index, we are interested in the information content of short term options. Second, we exclude options with less than 5 trading days to avoid liquidity-related biases. We furthermore exclude options with a price of less than 2 Euro to avoid problems due to stale prices. Finally, we don't automatically eliminate options whose absolute moneyness is greater than 10%, because deep in- and out-of-the-money options still contain useful information when they are actively traded. That's why we control for active trading by only using those options with a certain trading volume over the day. Other studies in contrast, that use a moneyness rule, run the risk of excluding options that are actually actively traded and contain information on volatilities.

Imposition of the aforementioned filtering rules reduces the number of options per trading day (puts and calls) in our dataset to around 22% of the original more than 580 options. Since our 'implied' analysis covers a period of 424 trading days, on average we have 124 options (puts and calls) per trading day that meet the criteria, with a minimum of 56 and a maximum of 226 options, which is more than sufficient for a reliable estimation of the four parameters. On every trading day we have at least two and at most 4 maturities with liquid options. Together with the long run stationary volatility level that we estimate from the time-series of five-year historical returns, the minimum of two maturities ensures a sensible estimation of the volatility term structure. Furthermore, in order to be able to estimate some shape of the volatility 'smile' we ensured that we had at least 10 different strikes in every maturity series. In the year 2000 nearly 32 million DAX option contracts were traded with a trading volume of more than €1.1 trillion: a clear indication that the DAX index options are a heavily traded option series. Therefore, we expect that they contain a lot of predictive value for future volatility.

The 'implied' calibration is executed with the well-known Newton-Raphson procedure programmed in the statistical software package GAUSS. We use the time-series estimates as starting values in the estimation and simulate ten thousand price paths. Although this large number of simulations caused the procedure to need sometimes over one hour of computation time per trading day, it appeared to be necessary to ensure stable option values. Local optima are always possible with this type of large-scale optimizations, but convergence was nevertheless hardly ever a problem.

As a benchmark for our volatility forecasts we use a volatility metric based on 5-minute intra-day returns. These returns are constructed from the contemporaneous index

futures transaction prices when DAX options are traded. Every day in our sample there are always several traded futures, each with a different time-to-maturity. We select the future with the shortest maturity. Since we are using transaction prices, negative autocorrelation may be present due to the bid-ask bounce. However, the bid-ask spread in the DAX future is minimal, normally a fraction of a percentage, and we do not find evidence for negative autocorrelation. Option and futures are traded from 9:00 till 17:30 in the first part of our sample (January - June 2000), and till 20:00 in the rest of the sample. The first option trade takes place after the first futures trade between around 9:09 and 9:15, but no later than 9:15. Therefore, for the calculation of realized volatility we take the difference between the opening price and the 9:15 trading price as the first return. For the rest of the day we take the returns from 9:15-9:20, 9:20-9:25 and so on. Since there are trading prices available just before and just after every 5-minute stamp (maximal a few seconds), we use the average of the prices preceding and following it. The realized volatility is computed as the square root of the sum of the 100 or 130 squared intra-day returns. We tested that other methods for defining the 5-minute interval return yield similar realized volatility estimates. Moreover, we did not use any overnight returns, because we believe it is rather arbitrary what weight such a return would deserve. Obviously, overnight returns are on average quite a bit larger in absolute terms than intra-day 5-minute returns, so the larger the weight it is given, the higher will the realized volatility be on average. Our realized volatility σ_i^{real} is a strict intra-day metric and that's why we expect it to be slightly biased downwards compared to average squared returns. Yet, this is not the case (Figure 5.2).

The annualized volatility measure based on squared daily returns is on average lower than the realized volatility measure (18.2% compared to 19.3%), but also much more volatile. Daily squared returns are a noisy measure for volatility and using it as a benchmark for the forecasting exercise in Section 5.4 would be inappropriate (see Blair (2001)).

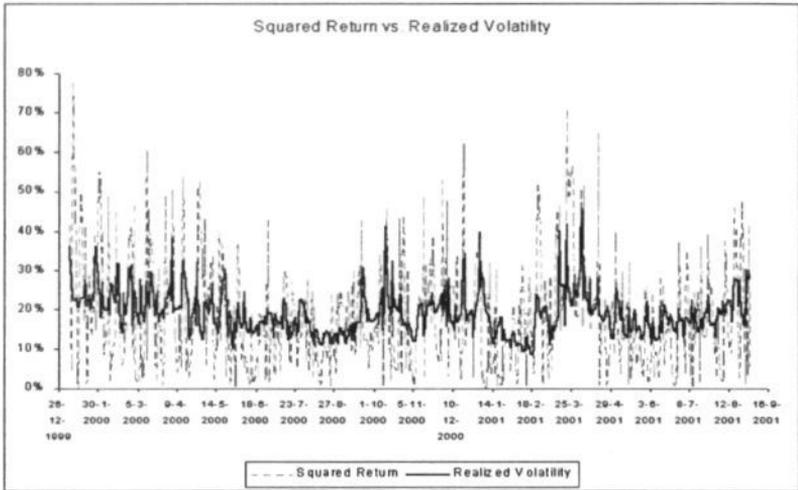


Figure 5.2: Squared Return vs. Realized Volatility

The graph shows a comparison of a volatility measure based on daily squared returns and a volatility measure constructed using intra-day returns for the period January 2000 until August 2001.

5.4 Empirical Results

5.4.1 In-sample Valuation Error

This section describes the parameter estimates, in- and out-of-sample comparisons of the 'implied' GARCH model with the historical' GARCH model. The 'historical' and 'implied' GARCH models are estimated daily using the implied DAX levels from futures prices and DAX option prices, respectively. For both calibrations we assume that the long run volatility is equal to the sample standard deviation, therefore ω is equal to $(1-\alpha)\ln(\hat{\sigma}^2) - \beta\sqrt{\frac{2}{\pi}}$ and $(1-\alpha)\ln(\hat{\sigma}^2) - \beta E[|\varepsilon_{t-1} - \lambda| + \gamma(\varepsilon_{t-1} - \lambda)]$, respectively.

As a benchmark for the in-sample fit of our GARCH option-pricing model, we use the ad-hoc Black-Scholes model of Dumas, Fleming and Whaley (1998). We allow each option to have its own Black-Scholes implied volatility depending on the exercise price K and time to maturity T and use the following functional form for σ_{ij} :

$$\sigma_{ij} = \omega_0 + \omega_1 M_i + \omega_2 M_i^2 + \omega_3 T_j + \omega_4 T_j^2 + \omega_5 M_i T_j, \quad (5.9)$$

where σ_{ij} denotes the implied volatility and M_i the moneyness, $\frac{K_i}{F_j}$ (F_j is the forward price) of a option for the i -th exercise price and j -th maturity. For every exercise price and maturity we compute the implied volatility and derive option prices using the Black-Scholes model.

The mean and standard deviation of parameter estimates from the daily maximum likelihood or the least squares estimation of the various models defined in Section 2 for the period January 3, 2000 until August 31, 2001 and the in-sample valuation errors appear in Table 5.1. The results show that the parameters of the model vary over time, but all parameters are relatively stable. It can be shown that pre-specifying the mean in the mean reverting volatility model stabilizes the estimation process and therefore the estimates. For the time series GARCH calibration the effect is known and the results are not surprising, but for the option implied GARCH calibration the method might be appropriate to estimate the parameters more efficiently and to avoid local maxima. The evaluation criteria for the option pricing models, the root mean squared error (RMSE) defined earlier is on average lower for the ad-hoc Black-Scholes model. However, the GARCH option pricing model with the constraint that the long run volatility is equal to the sample standard deviation has a competitive fit in-sample. The results of Heston and Nandi (2000) show that the ad-hoc Black-Scholes model might achieve better in-sample fit only by overfitting the data, but underperforms GARCH option pricing models out-of-sample. An out-of-sample pricing analysis is beyond the scope of this chapter, but we can conclude that the pricing performance of our method is reasonably accurate.

Table 5.1: In-Sample Analysis: Parameter Estimates

'Historical' GARCH			'Implied' GARCH			Ad-hoc Black-Scholes		
Parameter	Mean	(SD)	Parameter	Mean	(SD)	Parameter	Mean	(SD)
λ	0.067	(0.044)	λ	0.067	(0.044)	ω_0	1.615	(0.504)
ω	-0.409	(0.041)	ω	-0.556	(0.562)	ω_1	-2.318	(0.931)
α	0.965	(0.004)	α	0.958	(0.054)	ω_2	0.936	(0.456)
β	0.144	(0.013)	β	0.240	(0.144)	ω_3	-1.202	(1.482)
γ	0.349	(0.060)	γ	0.586	(0.420)	ω_4	0.735	(2.929)
						ω_5	1.007	(1.051)
			RMSE	0.0786	(0.0434)	RMSE	0.0664	(0.0335)

Notes: Mean and standard deviations of the parameter estimates from the daily Maximum Likelihood or least squares estimation of the GARCH models or the ad-hoc Black-Scholes model for the 'out-of-sample period' (January 2000 – August 2001) using the implied DAX® levels from the futures prices or DAX® option prices. For both GARCH models, ω is actually a result of setting the long run volatility level equal to the sample standard deviation (last 5 years) and therefore updated but not estimated. Note that for the GARCH models the conditional variance (not reported) is estimated simultaneously with the other parameters in the optimization procedure. 'RMSE' is the root mean squared error. Number of observations for the time series calibration = 753 every day.

5.4.2 'Historical' and 'Implied' Volatility Forecasts

The resulting one day ahead volatility forecast produced by the 'historical' and 'implied' GARCH models are presented in Figure 5.3.

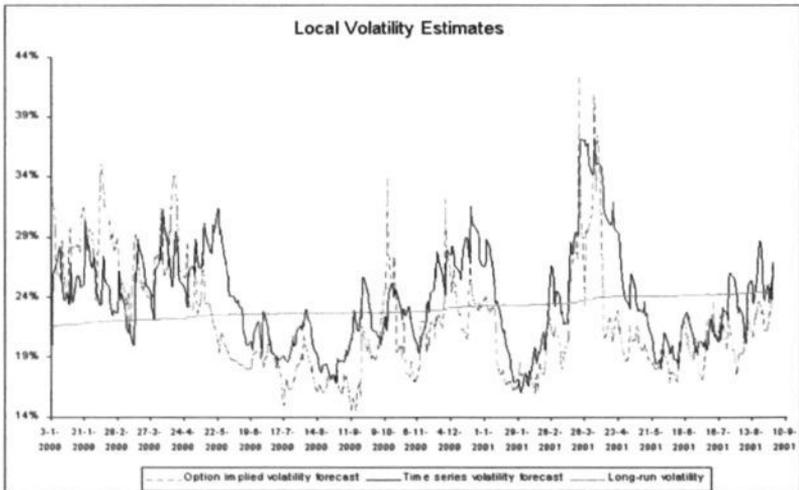


Figure 5.3: Option Implied Volatility Forecast vs. Time Series Volatility Forecast

The graph shows the annualized one-day ahead volatility forecasts of the option implied GARCH model and the time series GARCH model for the period January 3, 2000 until August 31, 2001. 'Long run volatility' is the annualized standard deviation each day for the last 5 years of daily returns.

It can be seen from the graph that in general both methods lead to different volatility estimates. Most of the time the 'implied' GARCH volatility forecast is lower compared to the 'historical' GARCH forecast, but during some periods the 'implied' exceeds the 'historical' GARCH forecast. Visual inspection suggests that whenever news is entering the market leading to a rising volatility estimate for the 'historical' GARCH model over the following days, this news is already incorporated in the 'implied' volatility forecast and the 'implied' forecast suddenly exceeds the 'historical' one. Therefore a positive or negative jump in 'implied' volatility forecast seems to indicate that there is new information in the market, but the 'historical' GARCH model needs some days to update the volatility estimate. As a result the time series of both local volatility estimates suggests that the 'implied' forecast is leading the 'historical' forecast.

5.4.3 Granger Causality Test

We want to test this hypothesis by conducting a Granger Causality test (Granger (1969)). The method determines the causal directions between two variables by indicating if changes in one variable induce changes in the other variable or both variables are jointly determined. Under the hypothesis of one variable not Granger-causing the other variable, the test statistic has the F-distribution and a rejection of the null hypothesis indicates causality. Table 5.2 reports the results for the Granger causality test.

Table 5.2: Granger Causality Test

Null Hypothesis	Lag = 1		Lag = 3		Lag = 5	
	F-stat	Prob.	F-stat	Prob.	F-stat	Prob.
'Historical' Forecast does not Granger Cause 'Implied' Forecast	1.22	0.271	1.81	0.150	2.20	0.054
'Implied' Forecast does not Granger Cause 'Historical' Forecast	96.3	0.000	56.2	0.000	33.4	0.000

Notes: This Table presents the results of a Granger Causality Test for 'historical' and 'implied' volatility forecasts from January 2000 to August 2001. 'Lag' is the number of lags used in the regression. 'F-stat' is the Wald F-statistic and 'Prob.' is the corresponding p-value. The test is based on 424 forecasts.

The test strongly suggest that our conjecture about a causality is statistically significant: The hypothesis that the option implied forecast does not cause the time series forecast can be rejected on all reasonable significant levels, while we cannot reject the hypothesis that the

'historical' forecast does not cause the 'implied' forecast on a 5% significance level for lags equal 1, 3 and 5. In particular we are interested in a lag of 1 and the results are very promising. Therefore we can conclude so far that the volatility estimate derived from option prices causes the volatility estimate using historical return data and as a result we can assume that the options market is highly informative when forecasting volatility.

5.4.4 Out-of-Sample Accuracy of the Volatility Forecasts

In the following, the out-of-sample accuracy of volatility forecast is compared from January 3, 2000 until August 31, 2001. Given 'historical' volatility forecast $x_{TS,t+1}$ and 'implied' volatility forecast $x_{OPT,t+1}$ made at time t of the realized volatility y_{t+1} known at time $t+1$, we can evaluate both models by comparing the multiple R^2 statistics from the regression

$$y_{t+1} = \alpha + \beta_{TS}x_{TS,t+1} + \beta_{OPT}x_{OPT,t+1} + u_{t+1}. \quad (5.10)$$

The multiple R^2 statistics can be interpreted as a measure of information content of the mixture of forecasts, which have more predictive power than univariate forecasts (Day and Lewis (1992)). Figures 5.4 and 5.5 plot the realized volatility estimate against the 'implied' and 'historical' forecasts, respectively and Table 5.3 reports the (multiple) squared correlation R^2 from regressions of realized volatility on one or two volatility forecasts.

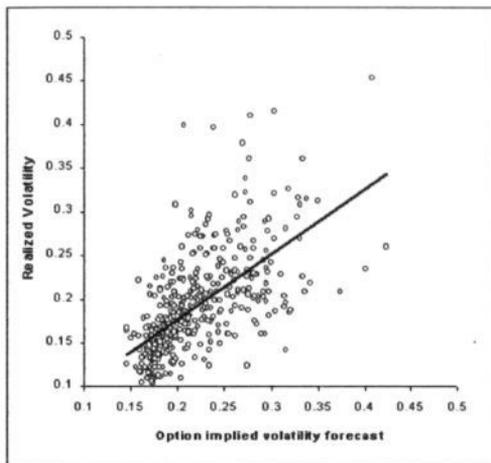


Figure 5.4: Forecast Evaluation ('Implied' Forecast)

Scatter plot of actual values and the 'implied' forecasts of realized volatility over 1 day (annualized), measured using intraday returns from January 2000 until August 2001.

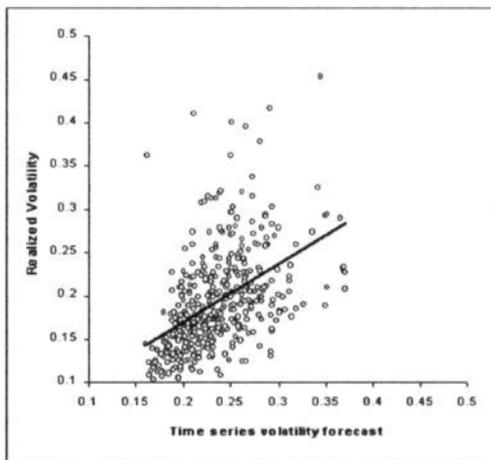


Figure 5.5: Forecast Evaluation ('Historical' Forecast)

Scatter plot of actual values and the 'historical' forecasts of realized volatility over 1 day (annualized), measured using intraday returns from January 2000 until August 2001.

Table 5.3: Regression Results

Forecasting Method	α^2	α^{2*}	β_{GARCH}	β_{GARCH^*}	$\beta_{GARCH} - \beta_{GARCH^*}$	R^2
Historical Forecast	0.095	2%	0.688	0.12		0.238
Implied Forecast	0.098	2%		0.578	0.00	0.279
Historical + Implied Forecast	0.098	0%	0.158	2%	0.680	0.298

Note: All data report parameter estimates and the correlation and simple correlation coefficients of realized volatility on volatility forecasts from historical GARCH and implied GARCH models from January 2000 to August 2001. Values of R^2 are calculated from the forecast.

It is seen first that the 'historical' forecast is less informative than the 'implied' forecast. Second, linear combinations of 'historical' and 'implied' forecasts have a higher statistic than univariate forecasts. Third, the incremental information is utilizing the option-implied information: the value of R^2 increases by 0.160 when the realized volatility measure is regressed on the 'historical' and 'implied' forecasts instead of 'historical' forecast alone, while the increase is only 0.001 when the realized volatility measure is regressed on the 'historical' and 'implied' forecasts instead of 'implied' forecast alone. Therefore more or less all relevant information is provided by the option implied measure.

5.5 Concluding Remarks

This chapter is concerned with short-term volatility forecasting. We compare and combine the information in historical returns and option prices to investigate what source contains most valuable information in forecasting DAX future volatility in the period of January 2000 till August 2001. More particularly, we compare the forecasts of a time-series GARCH model to the forecasts of a GARCH model whose main parameters are derived from contemporaneous option prices. We use the Duan (1995) option pricing model to identify an option-implied GARCH process and the corresponding 1-day ahead volatility forecast. A large number of simulations and optimizations is required to identify the parameters of the 'implied' model. Those parameters vary over time, but are relatively stable and provide an accurate fit to the option prices.

Our results indicate the following. First, the implied volatility forecast is leading the historical forecast. The first forecast Granger-causes the second, but not the other way round. This is an indication that information is more quickly compounded in option prices than in the most recent returns. This intuition is confirmed in the out-of-sample 1-day ahead volatility

forecasts. We test the implied and historical forecast accuracy against a realized volatility measure that is constructed with 5-minute intra-day returns. This is a more reliable measure of actual volatility than the very noisy squared returns. The historical forecast is able to explain around 23% of the variation in realized volatility, whereas the same ability of the implied measure is more than 1.5 times larger at 38%. Both measures appear to be slightly biased predictors of realized volatility, a phenomenon that is very common for option-implied forecasts.

Our result of options containing more (if not all) information than the time-series of historical returns is not new. Our main contribution is that we introduce a very different and theoretically sound methodology to extract the information from the option prices. Previous studies try to explore information from Black-Scholes implied volatilities of traded options to estimate and forecast future volatility without explicitly modeling the underlying return process. In those studies various optimal weighting schemes are being proposed for the different implied volatilities at different strikes. In doing so, those methods ignore the information contained in the volatility smile pattern, and probably more importantly, in the volatility term structure.

Apart from their theoretical drawbacks, we believe the ignorance of smile and term structure information by Black-Scholes-based methods may hurt their forecasting ability, especially at longer horizons. Therefore, in further research we will extend the forecasting to longer horizons of up to several days ahead. This opens up new interesting fields of research that are not solely focused on the second moment. For example, as an application of return distribution forecasting, we plan to calculate Value-at-Risk measures and analyze their accuracy. Since a longer time-series is required for a sensible Value-at-Risk comparison, we will then extend the analysis to a second dataset that spans a longer time-period and contains FTSE-100 index options from 1995-2000.

CHAPTER 6

Explaining Volatility Smiles¹⁶

6.1 Introduction

Before the crash in 1987 the Black-Scholes model was applied mechanically in its original form; there was a nearly constant relationship between the implied volatility and the exercise price of the option. After the crash practitioners seem to adjust the volatility for moneyness and maturity before being plugged into the Black-Scholes model. There are many alternative option-pricing models available in the literature, which try to explain the observed pattern. The most recent ones are the bivariate diffusion model of Hull and White (1987), the jump-diffusion model of Naik and Lee (1990) the variance-gamma model of Madan and Milne (1991) and the continuous-time stochastic volatility models of Stein and Stein (1991), Wiggins (1991) and Heston (1993), among others. A related class of volatility models is the discrete-time GARCH option-pricing model of Duan (1995). Kallsen and Taqqu (1998) developed a continuous-time version of the model and show that the same pricing results can be derived via an arbitrage-free argument. Heston and Nandi (2000) developed a closed form solution of a GARCH option-pricing model. They show that the single lag version of their model contains Heston's (1993) model as a continuous-time limit, but the discrete-time counterpart is much easier to apply with available data. The choice of the GARCH option pricing model was motivated by its flexibility, the recent empirical successes of the model

¹⁶ This chapter is based on: T. Lehnert, 2001, 'Explaining Smiles: GARCH Option Pricing with Conditional Leptokurtosis and Skewness', LIFE Working Paper, Maastricht University.

(see among others Amin and Ng (1994), Heynen, et al (1994), Duan (1996), Heston and Nandi (2000), and Ritchken and Trevor (1999)) and the emerging availability of numerical methods for this class of option pricing models (see Hanke (1997), Ritchken and Trevor (1999), Duan and Simonato (1998a/b), Duan et al. (1998), Duan et al. (2001), Bauwens and Lubrano (2000) and Heston and Nandi (2000)). In contrast to previous studies, our GARCH option-pricing model is driven by a skewed leptokurtic distribution. In particular we make use of the generalized GARCH option-pricing model of Duan (1999) where the innovations follow the skewed generalized error distribution of Theodossiou (2001).

Usually the parameters of the valuation formula are estimated from the discrete observations of asset prices. It has been shown that all option pricing models consistently exhibit systematic pricing errors when compared with market prices. They are not able to price the observed market smile, the convex relationship of the implied volatility to the exercise price of the option (for equities skewed towards the in-the-money call options) and the models cannot price the observed 'term structure of the smile', the smile seems to disappear for longer maturities. An obvious solution from the empirical point of view is to fit the GARCH option-pricing model directly to the observed market prices for all exercise prices and maturities. Instead of using historical returns of the underlying asset over a certain period in the past, we use every day the observed call and put prices of European DAX index options for different exercise prices and maturities from the EUREX derivative exchange. We estimate the parameters of the model by minimizing the relative root mean squared error (RMSE) between the market prices and the theoretical GARCH option prices. Given the parameters at a particular day we price traded in-, at- and out-of-the-money options during the following day using intra day options and futures data from the EUREX derivative exchange. Previous studies analyzed the pricing performance on a weekly basis. We extend previous research by analyzing the pricing performance for different model on a daily basis and we test empirically if conditional leptokurtosis and skewness for the asset return innovation can explain the shape of the volatility smile.

In an in- and out-of-sample analysis we use the RMSE criterion and compare our model with two challenging benchmarks: the closed-form option valuation model of Heston and Nandi (2000) and the ad-hoc Black-Scholes model of Dumas, Fleming and Whaley (1998) (DFW) and conclude which method is superior in describing the observed market smile in- and out-of-sample. DFW have shown that their ad-hoc Black-Scholes model outperform the implied binomial tree or deterministic volatility models in out-of-sample options valuation errors in the S&P 500 index options market. In particular DFW have shown

that misspecified models achieve good in-sample results by overfitting the data, but have less predictive power for out-of-sample option valuation. Heston and Nandi showed that their GARCH option valuation model outperforms the ad-hoc Black-Scholes model in out-of-sample pricing errors.

The chapter is arranged as follows. Section 2 describes the GARCH option-pricing framework. Section 3 discusses the data and methodology. In Section 4, we provide the empirical results and finally Section 5 concludes.

6.2 An Option Valuation Model

6.2.1 Option Pricing under GARCH

In this section, we focus on the generalized GARCH option-pricing framework developed in Duan (1999). We implement a model based on the exponential GARCH of Nelson (1991). Despite the well documented GARCH effects (see Bollerslev, et al. (1992)), this process is also able to model the well-known leverage effect of stock market returns: volatility tends to rise in response to bad news (negative excess returns) and to fall in response to good news (positive excess returns) (Nelson (1991)). The form of EGARCH specification is comparable to the non-linear asymmetric GARCH process of Engle and Ng (1993) or the power GARCH of Ding et al. (1993). The studies also show that it is crucial to include the asymmetric term in financial time series models. Because volatility and shocks to returns are negatively correlated, there is negative skewness in multiperiod index returns, which can produce the particular smile effect.

6.2.2 Conditional Leptokurtosis and Skewness

Previous studies suggest that conditional modeling with asymmetry in the volatility process cannot capture all the skewness and excess kurtosis in financial return distributions (see Chapter 2 and 3). Therefore it is crucial to model the additional skewness with a skewed leptokurtic innovations distribution, which supports the effect of asymmetry in the volatility process and improves the pricing of the smile in index options.

In contrast to previous studies our GARCH option-pricing model is driven by skewed leptokurtic innovations, we make use of the skewed generalized error distribution (GED) of Theodossiou (2001). The density function of a skewed GED random variable y_t is

$$g(y_t | \mu, \sigma, \nu, \beta) = \begin{cases} \frac{C}{\sigma} \exp\left(-\frac{\varphi^\nu}{(1-\beta)^\nu \sigma^\nu} |y_t - (\mu + \delta\sigma)|^\nu\right) & \text{for } y_t < \mu + \delta\sigma \\ \frac{C}{\sigma} \exp\left(-\frac{\varphi^\nu}{(1+\beta)^\nu \sigma^\nu} |y_t - (\mu + \delta\sigma)|^\nu\right) & \text{for } y_t \geq \mu + \delta\sigma \end{cases}, \quad (6.1)$$

where

$$\begin{aligned} A &= \frac{\Gamma(\frac{2}{\nu})}{\sqrt{\Gamma(\frac{1}{\nu})}\sqrt{\Gamma(\frac{3}{\nu})}} \\ B &= \sqrt{1 + 3\beta^2 - 4A^2\beta^2} \\ \varphi &= \frac{B\sqrt{\Gamma(\frac{3}{\nu})}}{\sqrt{\Gamma(\frac{1}{\nu})}} \\ \delta &= \frac{-2\beta A}{B} \\ C &= \frac{\frac{1}{2}\nu\varphi}{\Gamma(\frac{1}{\nu})} \end{aligned}$$

and μ is the expected value and σ the standard deviation of the random variable y_t , ν determines the tail-fatness ($\nu > 0$), β is a skewness parameter ($-1 < \beta < 1$) and $\Gamma(\cdot)$ is the gamma function. The random variable y_t can be normalized to have zero mean and unit variance. The use of the skewed generalized error distribution has one major advantage compared to leptokurtic alternatives. For this kind of distribution with exponential decay in the tails (GED), the fatness of the tails does not restrict the number of maximally existing moments of the innovation process. Since we are modeling the continuously compounded rate of returns, the conditional expected rate of return is always finite, as the moment generating function of the skewed generalized error-distributed random variable does always exist.

6.2.3 Local Risk Neutralization

The generalized GARCH option-pricing model also assumes the risk-neutral valuation principle and the standard Black-Scholes (1973) model can be derived as a special case. For this kind of derivative contracts with a high degree of path dependency, computationally

demanding Monte Carlo simulations are commonly used for valuing derivative securities, when a closed-form is not available. There are several variance reduction methods available to reduce the so-called 'simulation error'. We use the recently proposed simulation adjustment method, the empirical martingale simulation (EMS) (Duan and Simonato (1998a)), which has been shown to substantially accelerate the convergence of Monte Carlo price estimates. The EMS reproduces the martingale property for the simulated sample, a characteristic of all derivative pricing models.

In a Gaussian discrete-time economy the value of the index at time t , X_t can be assumed to follow the following dynamics¹⁷:

$$\ln\left(\frac{X_t}{X_{t-1}}\right) = r_f - d_t + \lambda\sigma_t - \frac{1}{2}\sigma_t^2 + \sigma_t e_t$$

$$e_t \mid \Omega_{t-1} \sim N(0,1) \text{ under probability measure } P$$

$$\ln(\sigma_t^2) = \alpha_0 + \alpha_1 \ln(\sigma_{t-1}^2) + \alpha_2 (|e_{t-1}| - \gamma e_{t-1}),$$
(6.2)

where r_f is the risk-free interest rate; d_t is the dividend yield of the index portfolio; λ can be interpreted as the constant unit risk premium; Ω_{t-1} is the information set in period $t-1$ and the combination of α_2 and γ captures the leverage effect. In the following we make use of the generalized local risk-neutral valuation relationship (GLRNVR) derived by Duan (see Duan (1999) for details). For the special case of Gaussian innovations the GLRNVR implies that $e_t = \varepsilon_t - \lambda$, where ε_t is a random variable under a risk-neutralized probability measure Q . ε_t is not necessarily normal, but to include the Black-Scholes model as a special case we assume that ε_t is a Gaussian random variable. We can substitute e_t into the conditional volatility process and derive an asset price process, which takes the following form:

$$\ln\left(\frac{X_t}{X_{t-1}}\right) + d_t = r_f - \frac{1}{2}\sigma_t^2 + \sigma_t \varepsilon_t$$

$$\varepsilon_t \mid \Omega_{t-1} \sim N(0,1) \text{ under risk-neutralized probability measure } Q.$$

$$\ln(\sigma_t^2) = \alpha_0 + \alpha_1 \ln(\sigma_{t-1}^2) + \alpha_2 (|\varepsilon_{t-1} - \lambda| - \gamma(\varepsilon_{t-1} - \lambda)),$$
(6.3)

If ε_t , conditional on Ω_{t-1} and with respect to measure P , is a skewed generalized error distributed random variable, then

$$\ln\left(\frac{X_t}{X_{t-1}}\right) + d_t = r_f - \ln\left(E^Q(\exp(\sigma_t \eta_t) | \Omega_{t-1})\right) + \sigma_t \eta_t,$$

$$\varepsilon_t | \Omega_{t-1} \sim N(0, 1) \text{ under risk-neutralized probability measure } Q. \quad (6.4)$$

$$\eta_t = G^{-1}(\phi(\varepsilon_t - \lambda); \nu, \beta)$$

$$\ln(\sigma_t^2) = \alpha_0 + \alpha_1 \ln(\sigma_{t-1}^2) + \alpha_2 (|\eta_{t-1}| - \gamma \eta_{t-1}),$$

where $\phi[\cdot]$ stands for the standard normal cumulative distribution and $G^{-1}[\cdot]$ stands for the inverse skewed GED cumulative distribution with tail parameter ν and skewness parameter β . $E^Q(\cdot | \Omega)$ can be evaluated numerically, but the approximation is extremely time demanding and therefore the estimation is carried out by restricting a part of the conditional mean $\ln[E^Q(\cdot | \Omega)]$ to a constant ω . If $\alpha_1 = \alpha_2 = \gamma = 0$, $\nu = 2$ and $\beta = 0$ the model reduces to the standard Black-Scholes model.

A European call option with exercise price K and maturity T has at time t price equal to:

$$c = e^{-r_t(T-t)} E_t^Q[\max(X_T - K, 0) | \Omega_t]. \quad (6.5)$$

In our model this price can be determined numerically. Given the value of the index X_t and the variance σ_{t+1}^2 , we generate N skewed generalized error distributed random numbers to advance the dynamics one period ahead and then make the empirical martingale adjustment (see Duan and Simonato (1998a)). Since we are dealing with a skewed fat-tailed distribution we are using $N=50,000$ price sample paths. For Gaussian models, Monte Carlo simulations with $N=10,000$ realizations per period are typically sufficient in practice, but the fatter tails of the skewed generalized error distribution require more sample paths. We repeat this procedure $T-t$ times until maturity and arrive at N simulated prices X_T . We calculate each of the N option payoffs, take the average and discount them back to period t of option valuation. Using this procedure we compute the value of an option for all exercise prices and all maturities.

¹⁷ The spot rate X_t is again determined from futures transaction prices; see section 5.2.1. for details.

6.2.4 Alternative Option Valuation Methods

Based on the GARCH option-pricing model of Duan (1995), Heston and Nandi (2000) derived a closed-form solution. In contrast to our model the asymmetry in the volatility process is modeled by a shift instead of a rotation of the news impact curve. In our notation the single lag version of the model reads

$$\ln\left(\frac{X_t}{X_{t-1}} + d_t\right) = r_f - \frac{1}{2}\sigma_t^2 + \sigma_t\varepsilon_t. \quad (6.6)$$

$\varepsilon_t \mid \Omega_{t-1} \sim N(0,1)$ under risk-neutralized probability measure Q .

$$\sigma_t^2 = \alpha_0 + \alpha_1\sigma_{t-1}^2 + \alpha_2(\varepsilon_{t-1} - (\gamma + \lambda + \frac{1}{2})\sigma_{t-1})^2,$$

The authors developed the generating function of the GARCH process and use it to price options. A European call option with exercise price K and maturity T has at time t price equal to:

$$c = \frac{1}{2}X_t + \frac{1}{\pi}e^{-r_f(T-t)} \int_0^{\infty} \text{Re}\left[\frac{K^{-i\theta} f^*(i\theta+1)}{i\theta}\right] d\theta - Ke^{-r_f(T-t)} \left(\frac{1}{2} + \frac{1}{\pi} \int_0^{\infty} \text{Re}\left[\frac{K^{-i\theta} f^*(i\theta)}{i\theta}\right] d\theta\right), \quad (6.7)$$

where $\text{Re}[\cdot]$ denotes the real part of a complex number and $f^*(i\theta)$ is the characteristic function of the risk-neutral process in Equation (6.6). Accurate numerical values for option prices can be calculated Fourier-transforming the characteristic functions and evaluating the integrals numerically. We use Romberg integration, which allows the prespecification of the tolerated error and in fact a calculation of the integral as precise as necessary (see Lambert and Lindsey (1999)).

We compare both GARCH option-pricing models with the ad-hoc Black-Scholes model of Dumas, Fleming and Whaley (1998). We allow each option to have its own Black-Scholes implied volatility depending on the exercise price K and time to maturity T and use the following functional form for σ :

$$\sigma_{ij} = \omega_0 + \omega_1 M_i + \omega_2 M_i^2 + \omega_3 T_j + \omega_4 T_j^2 + \omega_5 M_i T_j, \quad (6.8)$$

where σ_{ij} denotes the implied volatility and M_i the 'moneyness', $\frac{K_i}{F_j}$, (F_j is the forward price) of an option for the i -th exercise price and j -th maturity. For every exercise price and maturity we can compute the implied volatility and derive option prices using the Black-Scholes model.

6.3 Data & Methodology

6.3.1 DAX 30 Index Options

We use closing and intra-day DAX 30 index options and futures prices for a period from August 2000 until March 2001. The raw data set is directly obtained from EUREX, European Futures and Options Exchange¹⁸. Section 5.3 provides a detailed description of the data.

For the estimation of the model, we exclude options with less than 5 and more than 75 trading days until maturity and a price of less than 2 Euro to avoid liquidity-related biases and because of less useful information on volatilities. Instead of using a static rule and exclude options with absolute moneyness $|K/F-1|$ of more than 10% (see DFW), we exclude options with a daily Euro trading volume in terms of contract size of less than 10000 Euro. This rule was applied after carefully analyzing the particular data set. Among others DFW argue that options with absolute moneyness of more than 10% are not actively traded and therefore contain no information on volatilities. But this is also true for options within the 10% absolute moneyness interval. Therefore an obvious solution is to filter the available option prices and include all options that are actively traded, inside or outside the 10% absolute moneyness interval. As a result, on each day we determine a certain interval of exercise prices for the calibration, which we also use for the out-of-sample pricing exercise. The calibrations are carried out with the well-known Newton-Raphson procedure programmed in the statistical software package GAUSS.

6.3.2 Methodology

In particular we are using the following procedure for one particular day to price options during the following trading day:

First we compute the implied interest rates and implied index rates from the observed put and call option prices using essentially the method of Shimko (1993) based on the put-call parity (see Section 5.3 for details).

Second we estimate the parameters of the EGARCH option pricing model, the closed-form GARCH option valuation model or the ad-hoc Black-Scholes model by minimizing the relative mean squared error between the market prices and the theoretical option prices for calls and puts predicted by the different models. Given starting values obtained from estimating the model using historical returns, we price European calls and puts with exercise price K_i and maturity T_j . We repeat this procedure with the usual optimization method (Newton-Raphson method) and obtain the parameter estimates (e.g. $\hat{\omega}, \hat{\alpha}_0, \hat{\alpha}_1, \hat{\alpha}_2, \hat{\gamma}, \hat{\nu}, \hat{\beta}, \hat{\sigma}_{t+1}$ for the EGARCH model) that minimize

$$\frac{1}{M} \sum_{i=1}^n \sum_{j=1}^{m_i} \left(\frac{\hat{c}_{i,j} - c_{i,j}}{c_{i,j}} \right)^2 \quad \text{and} \quad \frac{1}{M} \sum_{i=1}^n \sum_{j=1}^{m_i} \left(\frac{\hat{p}_{i,j} - p_{i,j}}{p_{i,j}} \right)^2, \quad \text{where } n \text{ is the number of}$$

exercise prices and m_i represents the number of prices available out of all maturities for the i -

th exercise price. The total number of options used in the calibration is $M = 2 \sum_{i=1}^n m_i$. The

goodness of fit measure for the optimization is the root mean squared error criterion. In contrast to Heston and Nandi and DFW we use relative pricing errors instead of absolute pricing errors and put and call prices instead of simply call prices. The intuition behind this is based on the following observation: when we use absolute pricing errors and minimize the difference between the market prices for calls and the theoretical prices (predicted by the various models) the optimization procedure is mainly minimizing the difference of deep in-the-money calls. It happens frequently that for deep in-the-money calls, the minimized pricing error is e.g. 5€, but simultaneously deep out-of-the-money calls have theoretical prices of 0€. Since we are interested in the pricing performance for all options, the relative pricing errors is more appropriate.

Third we use the parameter estimates to price put and call options during the following day. We filter the available option prices according to our criteria for the in-sample calibration on the previous day. The futures market is the most liquid market and the options and the futures market are closely integrated, therefore it can be assumed that the futures price is more informative for option pricing than just using the value of the index. For every observed futures price we can derive the implied DAX index level and evaluate the option traded

¹⁸ The author gratefully acknowledges Deutsche Börse Group for providing the data.

directly after the futures trade. Given a futures price $F_{t,T}$ with time to maturity $T-t$, we use the Equation

$$F_{t,T} = e^{r_c(T-t)} X_{t,T}, \quad (6.9)$$

to derive the corresponding spot rate $X_{t,T}$. Additionally we use the method of Poteshman (2001) to match the maturity of the futures contract with the maturity of the option contract. This method allows us to perfectly match the observed option price and the underlying spot rate. Given our pricing model we compare the theoretical option prices and the market prices by using the RMSE criterion. We repeat this procedure every day for every traded option from 9.00 hours a.m. until 4.00 hours p.m. We stop the procedure at 4.00 hours p.m., because of the fact that Wall Street opened and probably new information is entering the market, which would have an effect on the pricing performance of the models that is not comparable to the pricing performance of the morning or early afternoon sessions. We repeat the procedure for all models and conclude which model is superior in evaluating options out-of-sample.

6.4 Pricing Results

6.4.1 In-Sample Analysis

The results suggest that there is a smile effect in DAX index options. For short time to maturity options the smile is skewed towards in-the-money call options and the smile seems to flatten for longer maturities. Both GARCH option-pricing model can be fitted very precisely to the observed market smile. The number of iterations in the optimization routine is similar to the standard GARCH calibration with return data for reasonable starting values, but the calibration is more time consuming.

The mean and standard deviation of the parameter estimates for the daily estimation (170 days) of the various models are given in Table 6.1. The results suggest that the parameters of all models are changing over time, but the estimates for the EGARCH model seem to be rather stable over time. The parameter estimates for the GARCH option-pricing models suggest a strong mean reversion of the volatility process, as measured by $\alpha_1 + \alpha_2 \gamma^2$ or α_1 , respectively. The implied risk-neutral stationary volatility e.g. for closed-form NGARCH

as given by $\sqrt{\frac{\alpha_0 + \alpha_2}{1 - \alpha_1 - \alpha_2 \gamma^2}}$ is usually greater than the historical standard deviation. The particular specification of the GARCH model allows α_0 to be negative as long as $\alpha_0 + \alpha_2 > 0$.

Table 6.1: In-Sample Analysis: Parameter Estimates

Ad-hoc Black-Scholes			Closed-form GARCH			EGARCH-skewed GED		
Parameter	Mean	(SD)	Parameter	Mean	(SD)	Parameter	Mean	(SD)
ω_0	1.633	(0.593)	α_0	-1.9e-05	(2.0E-05)	ω	3.7E-04	(8.7E-07)
ω_1	-2.456	(1.102)	α_1	0.780	(0.262)	α_0	-0.452	(0.341)
ω_2	1.050	(0.520)	α_2	4.7e-05	(4.2E-05)	α_1	0.961	(0.032)
ω_3	-1.072	(1.046)	γ	39.4	(15.520)	α_2	0.145	(0.096)
ω_4	0.822	(1.426)				γ	0.537	(0.181)
ω_5	0.826	(0.979)				ν	0.829	(0.262)
						β	-0.087	(0.021)

Notes: This Table reports the mean and standard deviation of all parameter estimates from the daily estimation of the various models (for the period from August 2000 until March 2001). Note that for the GARCH models the conditional variance (not reported) is estimated simultaneously with the other parameters in the optimization procedure.

Volatility and shocks to returns are negatively correlated as indicated by a γ being positive and significantly different from zero. Because of this negative correlation, there is negative skewness in the simulated multi-period index returns. Therefore the parameter γ produces the particular smile effect. In the case of the EGARCH option-pricing model this effect is supported by a negative skewness in the innovations distribution. The parameter estimates for the EGARCH option-pricing model driven by skewed GED innovations suggest that the particular choice of the distribution is motivated by the data: the innovations distribution is not normal. Estimates for the tail-fatness parameter ν are different from 2, usually even smaller than 1, and estimates for the skewness parameter β are different from 0, usually around -0.1 , which indicates excess kurtosis and negative skewness for the innovations distribution. Therefore we expect the tails of the multi-period return distribution to be fatter and the whole distribution to be more left-skewed than the one from the closed-form NGARCH model. As a result for option prices, the implied volatilities of out-of-the money puts and calls might be greater for the EGARCH-skewed GED than the ones for the closed-form NGARCH model, a feature that makes our model very flexible and interesting for pricing the observed market smiles for short-term options. Additionally there are (nearly) no parameter restrictions for the EGARCH model, while we experienced some possible

violations of the stationarity condition for the NGARCH option-pricing model of Heston and Nandi. Other characteristics of the time series dynamics of volatility are quite similar across the two GARCH models. The in-sample relative pricing error, the root mean squared error (RMSE) defined earlier, is usually in the range of 2.5% – 9% of option prices; usually slightly smaller for our model (on average ~3.5%) and closed-form GARCH (on average ~5.1%) option valuation model compared to the ad-hoc Black-Scholes model (on average ~6.7%). Table 6.2 reports the results for every month and for various models.

Table 6.2: In-Sample Analysis: Relative Valuation Errors

Month	Ad-hoc Black-Scholes		Closed-form GARCH		EGARCH-skew. GED	
	Mean	(SD)	Mean	(SD)	Mean	(SD)
August 2000	0.0503	(0.0231)	0.0458	(0.0247)	0.0314	(0.0154)
September 2000	0.0603	(0.0187)	0.0548	(0.0354)	0.0378	(0.0223)
October 2000	0.0801	(0.0316)	0.0569	(0.0311)	0.0393	(0.0196)
November 2000	0.0939	(0.0469)	0.0784	(0.0321)	0.0541	(0.0202)
December 2000	0.0655	(0.0299)	0.0466	(0.0205)	0.0321	(0.0129)
January 2001	0.0484	(0.0241)	0.0367	(0.0222)	0.0253	(0.0140)
February 2001	0.0658	(0.0242)	0.0499	(0.0229)	0.0344	(0.0144)
March 2001	0.0734	(0.0287)	0.0471	(0.0205)	0.0325	(0.0129)
Total	0.0671	(0.0324)	0.0513	(0.0282)	0.0354	(0.0178)

Notes: This Table reports the average percentage in-sample valuation errors for all months by various models from minimizing the sum of squared relative errors between model option values and market option prices. All models are estimated using daily closing prices from August 2000 until March 2001. 'Mean' and 'SD' are the mean and standard deviation of the valuation errors (the root mean squared relative pricing error (RMSE)).

Of course, the pricing performance of all models varies over time. During volatile periods (e.g. October/November 2000 or March 2001) the RMSE of all models increases (Mean) and there is more variation from one trading day to the other within the month (SD). But we do not confirm the results of Heston and Nandi: in particular the ad-hoc Black-Scholes model cannot outperform the closed-form NGARCH in-sample. Results suggest that the EGARCH-skewed GED model additionally improves the in-sample fit consistently. The different kind of modeling the asymmetry in the volatility process (a rotation of the news impact curve instead of a shift) seems to support the effect. Also the variations of pricing errors from one day to the

other are lower for the GARCH models (especially for the EGARCH specification) compared to the ad-hoc Black Scholes model.

6.4.2 Out-of-Sample Valuation Errors

The second part analyzes if the out-of-sample valuation errors of the EGARCH model are also lower than the closed-form GARCH model or the ad-hoc Black-Scholes model. It is still possible that our more complex model is overfitting the data, but has less predictive power for out-of-sample option valuation. We start by characterizing the traded options used for the out-of-sample pricing exercise. Table 6.3 shows the number of traded options across option moneyness and maturity categories.

Table 6.3: Out-of-sample Analysis: Number of Observations

Moneyness		Days to Expiration									Total
		< 21			[21,42]			> 42			
		Total	Mean	Max	Total	Mean	Max	Total	Mean	Max	
< 0.95	C	18994	151	991	21999	133	922	12064	77	524	53057
	P	25896	202	617	28674	173	901	15185	96	456	69755
[0.95,0.98]	C	30330	237	469	14078	85	497	3475	22	140	47883
	P	25611	200	396	13917	84	676	4309	26	176	43837
[0.98,1.02]	C	34050	266	556	11109	67	333	3061	19	101	48220
	P	32718	256	796	14015	84	1175	4779	30	260	51512
(1.02,1.05]	C	2744	21	167	1218	9	64	518	4	32	4480
	P	4762	37	260	3015	19	275	1418	12	97	9195
> 1.05	C	980	8	32	675	6	32	494	6	17	2149
	P	3595	32	405	3089	22	299	1858	20	200	8542
Total	C	87098	680	1463	49079	296	1775	19612	120	784	155789
	P	92582	723	1856	62710	378	3147	27549	169	1113	182841

Notes: This Table reports the average number of observations (number of options priced) for different moneyness and maturity during the period August 2000 until March 2001. 'Moneyness' is defined in the following way: a call option is said to be at the money if $F/K \in [0.98, 1.02]$, (deep) out of the money if $F/K \in [0.95, 0.98]$ ($F/K < 0.95$) and (deep) in the money if $F/K \in (1.02, 1.05]$ ($F/K > 1.05$), where K is the strike price and F is the forward price. Similar terminology is defined for puts by replacing F/K by K/F . 'C' stands for call options and 'P' stands for put options. 'Days to Expiration' is the number of trading days until maturity. 'Total' is the total number of options priced during the period August 2000 until March 2001 of the particular maturity and/or exercise price. 'Mean' and 'Max' are the average and the maximum number of options priced during one trading day of the particular moneyness and/or maturity.

In total we used prices of 155789 calls and 182841 puts to evaluate the pricing performance of the various models. We define 3 different categories for the maturity and 5 different categories for the moneyness of the options: Options with less than 1 month to maturity (< 21 trading days, short term options), 1-2 months to maturity (≥ 21 or ≤ 42 trading days, short/medium term) and more than 2 months to maturity (> 42 and < 75 trading days, medium term options) are analyzed separately. 'Moneyness' categories are defined in the following way: a call option is said to be at the money (ATM) if $F/K \in [0.98, 1.02]$, (deep) out of the money (OTM) if $F/K \in [0.95, 0.98)$ ($F/K < 0.95$) and (deep) in the money (ITM) if $F/K \in (1.02, 1.05]$ ($F/K > 1.05$), where K is the strike price and F is the forward price. Similar terminology is defined for puts by replacing F/K by K/F . In each given moneyness category, short term options are on average the most actively traded, followed by short/medium term and medium term options. In the short-term maturity category, ATM options have on average the highest trading volume, followed by OTM and deep OTM options. In the short/medium- and medium-term maturity category, deep OTM options are the most actively traded options. For all maturities, deep OTM options are the most actively traded call and puts, followed by ATM and OTM calls and puts. Next to this 'average trading activities', on some days especially deep OTM calls and ATM, deep OTM and deep ITM puts are traded in huge volumes for all maturities.

The out-of-sample valuation errors appear in Table 6.4 and 6.5. Results strongly suggest that the findings of Heston and Nandi are also valid for DAX options: the out-of-sample valuation errors of the closed-form NGARCH model are on average lower compared to the ad-hoc Black-Scholes model of DFW. But the EGARCH-skewed GED option-pricing model outperforms both models. The out-of-sample valuation errors for different moneyness and maturity are reported in Table 6.4. In each maturity category and across models the prediction of (deep) OTM options is worst, followed by ATM and ITM options. In each moneyness category the fit for longer maturities is better than the one for shorter maturities. This is a general result of option pricing studies, because of the extreme skewness in the smile for shorter maturities. After the crash 1987 traders always price short-term out-of-the-money puts with a huge premium, because there is always uncertainty in the market that investors will panic and prices go down dramatically. As a result the probability of extreme downward movements is assumed to be much greater than the probability of extreme upward movement and therefore the implied multi-period return distribution is skewed to the left. In terms of pricing this is difficult to model by describing the underlying return process. This is extremely

important for shorter maturities, because prices tend to mean revert in the long run and therefore market participants are more concerned about the short run. Therefore all models exhibit the largest valuation errors for short-term OTM options. In general we find that the GARCH option pricing models are superior to the ad-hoc Black Scholes model especially for ATM and OTM options and for all maturities. Therefore we confirm the findings of Heston and Nandi for DAX index options.

Table 6.4: Out-of-Sample Analysis: Average Valuation Errors across moneyness and maturity

Moneyness		Ad-hoc Black-Scholes			Closed-form GARCH			EGARCH-skewed GED		
		Days to Expiration			Days to Expiration			Days to Expiration		
		< 21	[21,42]	>42	< 21	[21,42]	>42	< 21	[21,42]	>42
< 0.95	C	0.2485	0.1324	0.1213	0.2189	0.1218	0.0792	0.1580	0.1114	0.0660
	P	0.2701	0.1231	0.0953	0.2647	0.1269	0.0836	0.1464	0.0942	0.0936
[0.95,0.98]	C	0.1356	0.0559	0.0464	0.0983	0.0550	0.0416	0.0973	0.0466	0.0347
	P	0.1281	0.0559	0.0528	0.1176	0.0586	0.0836	0.0932	0.0359	0.0520
[0.98,1.02]	C	0.0865	0.0396	0.0325	0.0591	0.0376	0.0280	0.0603	0.0309	0.0262
	P	0.0905	0.0392	0.0383	0.0585	0.0364	0.0311	0.0587	0.0265	0.0310
(1.02,1.05]	C	0.0273	0.0194	0.0195	0.0235	0.0196	0.0176	0.0193	0.0161	0.0218
	P	0.0297	0.0225	0.0274	0.0216	0.0205	0.0253	0.0218	0.0180	0.0189
> 1.05	C	0.0106	0.0130	0.0129	0.0103	0.0131	0.0121	0.0088	0.0099	0.0137
	P	0.0127	0.0111	0.0290	0.0115	0.0104	0.0278	0.0102	0.0103	0.0243

Notes: This Table reports the average relative pricing errors of the alternative models for different moneyness and maturity during the period August 2000 until March 2001. 'Moneyness' is defined in the following way: a call option is said to be at the money if $F/K \in [0.98,1.02]$, (deep) out of the money if $F/K \in [0.95,0.98]$ ($F/K < 0.95$) and (deep) in the money if $F/K \in (1.02,1.05]$ ($F/K > 1.05$), where K is the strike price and F is the forward price. Similar terminology is defined for puts by replacing F/K by K/F . 'C' stands for call options and 'P' stands for put options. 'Days to Expiration' is the number of trading days until maturity.

Comparing both GARCH option-pricing models, we find that the EGARCH model substantially outperforms the other model for all maturities and for different moneyness. The EGARCH model is in particular able to value (deep) OTM options much better for all maturities than the Heston specification. The pricing performance is especially remarkable for puts. Consequently it is the best model in- and out-of-sample. Table 6.5 reports the average valuation errors and the daily standard deviation for the different months. The number of observations and the average pricing errors are quite different across the different months.

The results suggest that the pricing error and the number of traded options during the day are positively correlated. This is counter-intuitive, but empirically established in several other studies¹⁹.

Table 6.5: Out-of-sample Analysis: Average Valuation Errors by Months

Month		Total	Ad-hoc Black-Scholes		Closed-form GARCH		EGARCH-skew.GED	
			Mean	(SD)	Mean	(SD)	Mean	(SD)
August 2000	C	18762	0.0925	(0.0349)	0.0742	(0.0261)	0.0675	(0.0222)
	P	20334	0.1043	(0.0309)	0.1050	(0.0402)	0.0756	(0.0281)
September 2000	C	22125	0.1249	(0.0380)	0.1058	(0.0432)	0.0963	(0.0367)
	P	27764	0.1226	(0.0333)	0.1319	(0.0555)	0.0949	(0.0389)
October 2000	C	19982	0.1539	(0.0548)	0.1092	(0.0642)	0.0993	(0.0546)
	P	25367	0.1457	(0.0615)	0.1227	(0.0495)	0.0884	(0.0346)
November 2000	C	21402	0.1353	(0.0437)	0.1063	(0.0423)	0.0967	(0.0360)
	P	23369	0.1480	(0.0694)	0.1387	(0.0661)	0.0998	(0.0463)
December 2000	C	16654	0.1256	(0.0551)	0.1041	(0.0542)	0.0948	(0.0461)
	P	17830	0.1416	(0.0582)	0.1382	(0.0631)	0.0995	(0.0441)
January 2001	C	16348	0.0897	(0.0329)	0.0746	(0.0327)	0.0679	(0.0278)
	P	18230	0.1101	(0.0326)	0.0995	(0.0415)	0.0716	(0.0290)
February 2001	C	15771	0.1350	(0.0564)	0.1109	(0.0537)	0.1009	(0.0457)
	P	17975	0.1408	(0.0573)	0.1387	(0.0681)	0.0998	(0.0477)
March 2001	C	24745	0.1706	(0.0741)	0.1605	(0.0760)	0.1376	(0.0646)
	P	31972	0.1557	(0.0639)	0.1406	(0.0692)	0.1013	(0.0484)
Total	C	155789	0.1281	(0.0558)	0.1055	(0.0563)	0.0960	(0.0478)
	P	182841	0.1332	(0.0547)	0.1264	(0.0583)	0.0910	(0.0408)

Notes: This Table reports the out-of-sample relative valuation errors by various models from pricing traded options during the period August 2000 until March 2001. All models are estimated using the closing prices of the previous day and the parameters are used to price all traded options during the following day. 'C' stands for call options and 'P' stands for put options. 'Days to Expiration' is the number of trading days until maturity. 'Total' is the total number of options priced during one particular month or for all months. 'Mean' and 'SD' are the mean and standard deviation of all valuation errors during the particular period (the root mean squared relative pricing error (RMSE)).

¹⁹ Similar results are also obtained in Heston and Nandi (2000).

Again it can be explained with the extreme skewness in the smile for shorter maturities. Investors always buy short-term protection of their portfolios through out-of-the money put options. The higher demand for these options make them relatively more expensive compared to at-the-money options. Increasing uncertainty in the market leads additionally to an increased demand of short-term protection and to a rise in the skew of the smile. The period August 2000 until March 2001 was an extremely volatile period, which additionally explains the huge pricing errors for short-term out-of-the-money puts. The valuation problems for these options are typically greater when the pricing model tries to describe the whole smile and not only a part of it. But also note that the valuation errors for the GARCH models and the ad-hoc Black-Scholes model do not necessarily covary.

The average pricing errors for the EGARCH model are consistently lower and remarkable stable for all months. Evaluated on a monthly basis the EGARCH model is able to value traded options much better than the ad-hoc Black-Scholes model and the closed-form NGARCH model. Additionally the variations of the RMSE from one day to the other are substantially lower. The results also strongly suggest that the EGARCH specification driven by a skewed leptokurtic distribution is superior to the NGARCH-Normal model in replicating the future data generating process underlying option prices.

6.5 Conclusions

This chapter presents an alternative solution to option evaluation. We consider a GARCH option-pricing model driven by a skewed leptokurtic distribution and allow the variance of the process to be correlated with asset returns. In an in-sample analysis we show that the skewness and leptokurtosis seem to add more flexibility to the GARCH option-pricing model.

Empirical analysis on the DAX index options data shows out-of-sample valuation improvements over the ad-hoc Black-Scholes model of DFW and the closed-form GARCH model of Heston and Nandi. The particular shape of the innovations distribution results in substantial valuation improvements for short term out-of-the money put options. From the analysis we can conclude that skewness and conditional leptokurtosis for the asset return innovation and the correlation of volatility with spot returns together can better explain the shape of the volatility smile.

CHAPTER 7

Summary and Concluding Remarks

7.1 Introduction

This study provided a set of empirical methods to model extreme events in financial data and implications for financial risk management have been investigated. In this concluding chapter we will outline the results of each of the preceding chapters and formulate some general conclusions.

The literature on the management and measurement of financial risk is huge and deals with many interesting issues. Every statistical approach relies on certain assumptions about the price process of the underlying asset and naturally the risk arises that the stochastic process is wrongly specified. Various methods are proposed in the literature, but commonly asset prices are assumed to be log-normally distributed. Several generalizations for the return process are frequently considered, but the stochastic component is usually assumed to be normal. However, our results suggest that empirically, there are more extreme price movements than a lognormal distribution allows for and in order to describe the dynamics of more volatile processes, a precise description of the stochastic component is necessary. Additionally the asymmetric response of volatility to changes in the asset value is empirically investigated. Besides the well-known leverage effect, we found that skewness in the stochastic component can additionally explain the asymmetry in the volatility process. We

showed that the assumption of a non-normal stochastic component could be justified not only for methods to quantify financial risk exposure, but also for option valuation methods.

The choice of a more sophisticated distribution with particular tail properties to describe extreme fluctuations is sometimes problematic, because we observe these extreme changes only occasionally and therefore the number of observations in the tail of the return distribution is limited. There is a lot of uncertainty about the actual fatness of the tails and about the parameters measuring the fatness. We show explicitly the impact of estimation risk on the measurement of market risk in term of a Value-at-Risk estimate. We have also been concerned with the information content of implied index volatility in the context of forecasting market risks. Previous studies have produced conflicting conclusions about the informational efficiency of option markets. Using a new approach, we found evidence that there appears to be incremental forecasting information in option prices and only minor information in historical index returns.

The following section summarizes all results in detail.

7.2 Summary

In Chapter 2, we considered a general class of time varying volatility model, the generalized augmented GARCH process, to model the dynamics of financial prices over time and compared different specifications for the stochastic component. Previous research (Hentschel (1995)) suggests that the difference in conditional volatility estimates could be substantial among GARCH specifications, where the leverage effect is either modeled with a shift or a rotation of the news impact curve. The most sophisticated model among the models analyzed also incorporates conditional leptokurtosis and skewness. Allowing for a leptokurtic stochastic component for the return process means that a priori the existence of moments cannot be ensured. Therefore, instead of modeling the first and second moment of the return process, we generally describe the dynamics of the location and scale parameter of distributions. A Lagrange Multiplier based on the augmented GARCH process was developed that allows us to reject various models. For particular index portfolios, we found evidence that the inclusion of more sophisticated tail properties for the conditional distribution, leads to different conclusions about the appropriate type of asymmetry in the volatility process. In all time series the observed skewness in the data is a result of a significant leverage effect *and*

skewness in the stochastic component. Under the hypothesis of conditional skewness, models with different modifications of the news impact curve can be rejected for different markets. However quantile comparison plots show that the asymmetric Student-t distribution is still not flexible enough for capturing the behavior of more volatile processes.

In the third chapter, we propose an alternative distribution for the stochastic component, which also incorporates a method to derive the multi-day return distribution from the estimated single-day distribution. A stochastic process, initially developed in the physics literature, is successfully applied to finance. We propose a generalization of the popular RiskMetricsTM approach to Value at Risk for a stock portfolio. We investigate a new approach to optimally describe the price process and propose a new scaling rule to forecast volatility: location and scale parameter (volatility) are estimated on one time scale (daily) and the multi-day (weekly, bi-weekly or monthly) values are derived by using the bi-fractal scaling property of truncated Lévy processes. The method has the implied advantage that we are able to identify the relationship between return distributions for different sampling intervals by analyzing the time series of returns on one sampling interval (say daily). One interesting empirical feature of financial data is that typically annualized volatility for daily sampling interval appears to be substantially lower than for monthly sampling interval. This may be a result of short-term positive serial correlation in the daily data series. If price changes are not independent over time, estimated volatility is affected. Positive autocorrelation, which occurs when observed prices adjust to new information with a lag over short intervals, reduces estimated volatility. This problem largely disappears with longer differencing intervals, which is the reason to consider the return distribution for monthly rather than daily data for longer horizon forecasting. Therefore our approach might produce better forecasts compared to methods, which add up one-day forecasts to derive a multi-day forecast. Both the in-sample performance of the model for VaR calculations and the out-of-sample performance are quite promising. There are plenty of practical benefits of this model, which can be explored in future research.

In the fourth chapter we have focused on VaR-model selection for exchange rate positions. We obtained several interesting results: The exchange rate fluctuations under consideration exhibit fat-tail behavior, but the number of extreme observations is limited. Models that take account of tail behavior are required, since otherwise the reported VaR leads to an under-estimation of the downside risk of the currency portfolio. On the other hand more

sophisticated approaches using fat tail distribution with slow power-law decay in the tails assume more probability mass in the tail of the distribution than is actually present. Because of a very limited number of extreme observations, the actual tail behavior is measured with relative great uncertainty. In general, models with more sophisticated tail behavior lead to more parameter uncertainty, which leads to greater uncertainty in the reported VaRs. If we adjust our VaR measure for the uncertainty of the parameter estimates, we obtain e.g. for the GARCH(1,1)-Student-t model rejection rates that are close to the confidence levels of the reported VaR. Among the models analyzed, the GARCH(1,1)-Student-t specification with adjustment seems to be an adequate model to correctly assess extreme losses for exchange rate positions.

In the fifth chapter, we combine the information in historical returns and option prices to investigate what source contains most valuable information in forecasting DAX future volatility. More particularly, we compare the forecasts of a time-series GARCH model to the forecasts of a GARCH model whose main parameters are derived from contemporaneous option prices. Our results indicate the following. First, the implied volatility forecast Granger-causes the time-series volatility forecast. Second, the option implied method outperforms the time-series method in out-of-sample forecasting accuracy when compared against a realized volatility measure that is constructed with 5-minute intra-day returns. Third, both measures appear to be slightly biased predictors of realized volatility. The main contribution in this chapter is that we introduce a very different and theoretically sound methodology to extract the information from the option prices. We showed that option prices contain more information about future volatility, in particular more information about the future price process. Therefore option prices should also be considered for the measurement of market specific risks, because all practical measures, like VaR, are forward looking measures and option prices might be more informative. Previous studies try to explore information from an implied volatility index based on Black-Scholes implied volatilities of traded options to estimate and forecast future volatility without explicitly modeling the underlying return process. Apart from their theoretical drawbacks, we believe the ignorance of smile and term structure information by Black-Scholes-based methods may hurt their forecasting ability, especially at longer horizons.

The last chapter presents an alternative solution to option evaluation. The hypothesis of a skewed, leptokurtic stochastic component in the return process is tested in a conditional

option-pricing framework. In an in-sample analysis we show that the conditional skewness and leptokurtosis seem to add more flexibility to the GARCH option-pricing model. Empirical analysis on the DAX index options data shows out-of-sample valuation improvements over the ad-hoc Black-Scholes model of DFW and the closed-form GARCH model of Heston and Nandi. There is already some evidence that existing option-pricing models fail to predict the observed skew in the implied volatility smile, the greatest valuation errors occur for (deep) out-of-the money options. The particular shape of the innovations distribution we considered, results in substantial valuation improvements for short term out-of-the money put options. From our analysis we can conclude that conditional skewness and leptokurtosis for the asset return innovation and the correlation of volatility with spot returns together can better explain the shape of the volatility smile. This is in line with the results obtained in Chapter 2 and 3.

7.3 Directions for further research

The stochastic process used in Chapter 3 is only one example how (price) fluctuations are modeled in the physics literature. The idea of describing the scaling behavior of the moments of a return distribution for different time scales together with the availability of high frequency data is in particular interesting for financial applications where e.g. price changes are typically serially correlated on high frequencies. Apart from a parametric approach, the idea of an empirical scaling function for the moments is extremely promising (see e.g. Calvet and Fisher (1999)). This approach also allows modeling higher moments like skewness and kurtosis. Together with the concept of entropy densities, a lot of financial applications of the method are possible, which can be explored in future research.

In Chapter 2, 3 and 4, we applied our method to index portfolios or FX-rates and discussed the results considering a univariate setting. Of course the analyses can be extended to a multivariate case, looking at e.g. volatility spillovers and the asymmetric response of volatility in one market to price fluctuations in another market. A practical application could also be the measurement of the risks involved in trading a portfolio of stocks and options, considering the dynamic variances and covariances of the price changes. An interesting framework for that kind of analysis would be the family of multivariate GARCH models, for instance the asymmetric dynamic covariance (ADC) matrix model of Kroner and Ng (1998). In general assuming non-normality for the stochastic component of the dynamics of changes

can also be discussed in a multivariate setting assuming that the GARCH model is driven by skewed leptokurtic shocks. Bauwens and Laurent (2002) developed a new class of multivariate skew densities, for which each marginal may have different asymmetry, a feature especially useful for modeling stock returns. Another multivariate application could be the pricing of options on a basket of stocks.

Bibliography

- Ahn, D.-H., J. Boudoukh, M. Richardson and R.F. Whitelaw 1999, "Optimal Risk Management Using Options", *Journal of Finance*, 54, 359-375.
- Akgiray, V. (1989), "Conditional heteroskedasticity in time series of stock returns: evidence and forecasts", *Journal of Business*, 62, 55-80.
- Amin, K. and V. Ng (1994): "A Comparison of Predictable Volatility Models Using Option Data", University of Michigan, Working Paper.
- Andersen, T.G. and T. Bollerslev (1998), "Answering the skeptics: yes standard volatility models do provide accurate forecasts", *International Economic Review*, 39, 885-905.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and H. Ebens (2001a), "The distribution of realized stock return volatility", *Journal of Financial Economics*, 61, 43-76.
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2001b), "The distribution of exchange rate volatility", *Journal of the American Statistical Association*, 96, 42-55.
- Bakshi, G., C. Cao, Z. Chen (2000), "Do Call Prices and the Underlying Stock Always Move in the Same Direction", *The Review of Financial Studies*, 13, 549-584.
- Barberis, N., (2000), "Investing for the Long Run when Returns Are Predictable", *Journal of Finance*, 55-1, 225-264
- Basel Committee on Banking Supervision, 1996a, *Amendment to the Basel Capital Accord to Incorporate Market Risk*, BIS, Basel, Switzerland.

- Basel Committee on Banking Supervision, 1996b, *Supervisory Framework for the Use of "Backtesting" in Conjunction with the Internal Models Approach to Market Risk Capital Requirements*, BIS, Basel, Switzerland.
- Basel Committee on Banking Supervision, 1999, *A New Capital Adequacy Framework*, BIS, Basel, Switzerland.
- Basel Committee on Banking Supervision, 2001, *The New Basel Capital Accord*, BIS, Basel, Switzerland.
- Bates, D.S. (1996), "Testing option pricing models", in Maddala, G.S. and C.R. Rao (eds.), *Handbook of statistics 1.4: Statistical methods in finance*, Elsevier Science, North-Holland, 567-611.
- Bauwens, L. and M. Lubrano (2000), "Bayesian option pricing using asymmetric GARCH models", CORE, Université catholique de Louvain, Working Paper.
- Bauwens, L. and S. Laurent (2002), "A new class of multivariate skew densities, with application to GARCH models", CORE, Université catholique de Louvain, Working Paper.
- Black, F. and M. Scholes (1973), "The pricing of options and corporate liabilities", *Journal of Political Economy*, 81, 637-659.
- Black, F. (1976), "Studies of Stock Price Volatility Changes", *Proceedings from the American Statistical Association, Business and Economic Statistics Section*, 177-181.
- Blair, B.J., S.-H. Poon and S.J. Taylor (2001), "Forecasting S&P100 volatility: The incremental information content of implied volatilities and high frequency index returns", *Journal of Econometrics*, 105, 5-26.
- Bollerslev, T., (1986), "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, 31, 307-327.
- Bollerslev, T., R.Y. Chou and K.F. Kroner (1992), "ARCH Modeling in Finance: A Selective Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 5-59.

- Bollerslev, T., R.F. Engle and D.B. Nelson (1994), "ARCH Models," in R.F. Engle and D. McFadden (eds.), *Handbook of Econometrics, Volume IV*, 2959-3038. Amsterdam: North-Holland.
- Bouchaud, J.-P. and M. Potters (2000), "*Theory of Financial Risk*", Cambridge University Press.
- Brailsford, T.J. and R.W. Faff (1996), "An evaluation of volatility forecasting techniques", *Journal of Banking and Finance*, 20, 419-438.
- Calvet, L. and A. Fisher (1999), "Multifractality in asset returns: Theory and evidence", Working paper.
- Christensen, B.J. and N.R. Prabhala (1998), "The relation between implied and realized volatility", *Journal of Financial Economics*, 50, 125-150.
- Christie, A.A. (1982), "The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects", *Journal of Financial Economics*, 10, 407-432.
- Cont, R., M. Potters and J.-P. Bouchaud (1997a), "Scaling in stock market data: stable laws and beyond", Science & Finance Working Paper (cond-mat/970587).
- Cont, R., (1997b), "Scaling and correlation in financial time series", Science & Finance Working Paper 97-01 (cond-mat/970575).
- Danielsson, J. and C. de Vries (1997), "Value-at-Risk and Extreme Returns", Manuscript, London School of Economics.
- Day, T.E. and C.M. Lewis (1992), "Stock market volatility and informational content of stock index options", *Journal of Econometrics*, 52, 267-287.
- Derman, E. (2001), "The Great Pretender", *Derivatives Strategy Magazine*, 6, 1.
- Deutsche Börse (2002), "Guide to the Volatility Indices of Deutsche Börse", Frankfurt am Main.
- Diebold, F.X., A. Hickman, A. Inoue, T. Schuermann (1997), "Converting 1-Day Volatility to h-Day Volatility: Scaling by $h^{1/2}$ is Worse than You Think", University of Pennsylvania, Working Paper.

- Diebold, F.X., P. Christoffersen, T. Schuermann, (1998a) "Horizon Problems and Extreme Events in Financial Risk Management", University of Pennsylvania, Working Paper.
- Diebold, F.X., T. Schuermann and A. Inoue (1998b), "Scale Models," *RISK*, 11, 104-107.
- Dimson, E. and M. Marsh (1990), "Volatility forecasting without data-snooping", *Journal of Banking and Finance* 14, 399-421.
- Ding, Z., C.W.J. Granger, and R.F. Engle (1993), "A Long Memory Property of Stock Market Returns and a New Model," *Journal of Empirical Finance*, 1, 83-106.
- Duan, J.-C. (1995), "The GARCH option pricing model", *Mathematical Finance*, 5, 13-32.
- Duan, J.-C. (1996), "Cracking the smile", *RISK*, 9, 55-59.
- Duan, J. D. (1997), "Augmented GARCH (p,q) process and its diffusion limit", *Journal of Econometrics*, 79, 97-127.
- Duan, J.-C. (1999), "Conditional Fat-Tailed Distributions and the Volatility Smile in Options", Rotman School of Management, University of Toronto, Working Paper.
- Duan, J.-C. and J.-G. Simonato (1998a), "Empirical martingale simulation for asset prices", *Management Science*, 44, 1218-1233.
- Duan, J.-C. and J.-G. Simonato (1998b), "American option pricing under GARCH by a Markov chain approximation", Rotman School of Management, University of Toronto, Working Paper.
- Duan, J.-C. Gauthier G. and J.-G. Simonato (1998), "An Analytical Approximation for the GARCH Option Pricing Model", Rotman School of Management, University of Toronto, Working Paper.
- Duan, J.-C. Gauthier G. and J.-G. Simonato (2001), "Asymptotic Distribution of the Empirical Martingale Simulation Option Price Estimator", to appear in *Management Science*.
- Dumas, B., J. Fleming and R.E. Whaley (1998), "Implied volatility functions: empirical tests", *The Journal of Finance*, 53, 2059-2106.

- DuMouchel, W., (1973), "On the asymptotic normality of the maximum-likelihood estimate when sampling from a stable distribution", *Annals of Statistics* 1, 948-957.
- Dunbar, N. (1999), "LTCM and the Dangers of Marking to Market", *Risk*, 12, 6-7.
- Embrechts, P., S. Resnick and G. Samorodnitsky, (1998), "*Extreme Value Theory as a Risk Management Tool*", North American Actuarial Journal, to appear.
- Engle, R.F. and V. Ng (1993), "Measuring and testing the impact of news on volatility", *Journal of Finance*, 48/5: 1749-78.
- Fama, E.F., (1963), "Mandelbrot and the stable Paretian hypothesis", *Journal of Business*, 36, 420-429.
- Fama, E.F., (1965), "The behavior of stock market prices", *Journal of Business*, 38, 34-105.
- Fama, E.F. and R. Roll (1968), "Some Properties of Symmetric Stable Distribution", *Journal of the American Statistical Association*, 63.
- Fleming, J. (1998), "The quality of market volatility forecasts implied by S&P100 index option prices", *Journal of Empirical Finance*, 5, 317-345.
- Fleming, J., B. Ostdiek and R.E. Whaley (1995), "Predicting stock market volatility: a new measure", *Journal of Futures Markets*, 15, 229-235.
- Franses, P.H. and D. van Dijk (1995), "Forecasting stock market volatility using (non-linear) GARCH models", *Journal of Forecasting*, 15, 229-235.
- Ghose, D. and K. F. Kroner (1995), "The relationship between GARCH and symmetric stable processes: Finding the source of fat tails in financial data", *Journal of Empirical Finance*, 2, 225-251.
- Glosten, L.R., R. Jagannathan and D.E. Runkle (1993), "On the relation between the expected value and the volatility of the nominal excess return on stocks", *Journal of Finance*, 48, 1779-1801.
- Gnedenko, B.V. & A.N. Kolmogorov, "*Limit Distributions for Sums of Independent Random Variables*", Addison Wesley, 1954.

- Granger, C.W.J. (1969), "Investigating Causal Relations by Econometric Models and Cross-spectral Models", *Econometrica*, 37, 424-438.
- Green, T.C. and S. Figlewski (1999), "Market Risk and Model Risk for a Financial Institution Writing Options", *Journal of Finance*, 54, 1465-1499.
- Hafner, C.M. and H. Herwartz (2001), "Option pricing under linear autoregressive dynamics, heteroskedasticity, and conditional leptokurtosis", *Journal of Empirical Finance*, 8, 1-34.
- Hanke, M. (1997), "Neural network approximation of option pricing formulas for analytically intractable option pricing models", *Journal of Computational Intelligence in Finance*, 5, 20-27.
- Hentschel, L. (1995), "All in the Family: Nesting Symmetric and Asymmetric GARCH Model", *Journal of Financial Economics*, 39, 71-104.
- Heston, S.L. (1993), "A closed-form solution for options with stochastic volatility with applications to bond and currency options", *The Review of Financial Studies*, 6, 327-343.
- Heston, S.L. and S. Nandi (2000): "A Closed-Form GARCH Option Valuation Model", *The Review of Financial Studies* 3, 585-625.
- Heynen, R., A. Kemma and T. Vorst (1994), "Analysis of the term structure of implied volatilities", *Journal of Financial and Quantitative Analysis*, 29, 31-56.
- Hull, J. and A. White (1987), "The pricing of options on assets with stochastic volatilities", *Journal of Finance* 42, 281-300.
- Jorion, P. (1995), "Predicting volatility in the foreign exchange market", *Journal of Finance*, 50, 507-528.
- Kallsen, J. and M. Taqqu (1998), "Option pricing in ARCH-type models", *Mathematical Finance*, 8, 13-26.
- Koponen, (1995), "Analytic approach to the problem of convergence of truncated Lévy flights towards the Gaussian stochastic process", *Physical Review E*, Volume 52, Number 1, 1197-1199.

- Kroner, K. and V. K. Ng (1998), "Modeling Asymmetric Comovements of Asset Returns," *The Review of Financial Studies*, 11, 817-844.
- Laurent, S. (2002), "Asymmetry and fat-tails in financial time series", PhD thesis, Maastricht University, The Netherlands.
- Lambert P. and J.K. Lindsey (1999), "Analysing financial returns by using regression models based on non-symmetric stable distributions", *Journal of Applied Statistics*, 48, 409-424.
- Liu, S. and Brorsen, B.W. (1995), "Maximum Likelihood Estimation of a GARCH-Stable Model", *Journal of Applied Econometrics*, 10, 273-285.
- Longin, F., (1997), "From value-at-risk to stress testing, the extreme value approach", Discussion Paper 97-004, CERESSEC.
- Lundberg, S. and T. Teräsvirta (2002), "Evaluating GARCH models", *Journal of Econometrics*, 110, 2, 417-435.
- Madan, D. and F. Milne (1991), "Option pricing with VG martingale components", *Mathematical Finance*, 1, 39-56.
- Mandelbrot, B. (1963), "The Variation of Certain Speculative Prices," *Journal of Business*, 36, 394-419.
- Mandelbrot, B. (1997), "A multifractal model of asset returns," Cowles Foundation Discussion Paper #1164, Department of Economics, Yale University.
- Mantegna, R. N. and H. E. Stanley (1994), "Stochastic process with ultraslow convergence to a Gaussian: the truncated Lévy flight", *Phys. Rev. Lett.*, 73, 2946-2949.
- Marshall, C. and M. Siegel (1996), "Value-at-Risk: Implementing a Risk Measurement Standard", Working Paper, The Wharton School, University of Pennsylvania
- Matacz, A. (2000), "Financial Modeling and Option Theory with the Truncated Lévy Process", *International Journal of Theoretical and Applied Finance*, 3, 143.
- McNeil A.J. and R. Frey (2000), "Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach", *Journal Of Empirical Finance* (7), 3-4, 271-300.

- McNeil A.J., (1998), "Calculating Quantile Risk Measures for Financial Return Series using Extreme Value Theory", Working Paper, ETH Zürich.
- Meese, R.A. and K. Rogoff (1983), "Empirical exchange rate models of the seventies: Are any fit to survive?", *Journal of International Economics*, **14**, 3-24.
- Mittnik, S. and M. S. Paoella, (2000a), "Conditional Density and Value-at-Risk Prediction of Asian Currency Exchange Rates", *Journal of Forecasting*, **19**, 313-333.
- Mittnik, S. Paoella, M.S., Rachev, S.T. (2000b), "Diagnosing and treating the fat tails in financial return data", *Journal of Empirical Finance* (7), 3-4, 389-416.
- Mittnik, S. Paoella, M.S., Rachev, S.T. (2002), "Stationarity of stable power-GARCH processes", *Journal of Econometrics*, 106, 97-107.
- Naik, V. and M.H. Lee (1990), "General equilibrium pricing of options on the market portfolio with discontinuous returns", *Review of Financial Studies*, 3, 493-522.
- Nakao, H., (2000), "Multi-scaling properties of truncated Lévy flights", Manuscript, Graduate School of Mathematical Sciences, University of Tokyo.
- Nelson, D. (1991), "Conditional heteroskedasticity in asset returns: a new approach", *Econometrica*, 59, 347-370.
- Nelson, D. (1992), "Filtering and forecasting with misspecified ARCH models I: getting the right variance with the wrong model", *Journal of Econometrics*, 52, 61-90.
- Nelson, D. and D.P. Foster (1995), "Filtering and forecasting with misspecified ARCH models II: making the right forecast with the wrong model", *Journal of Econometrics*, 67, 303-335.
- Paoella, M., (1997), "Tail Estimation and Conditional Modeling of Heteroscedastic Time-Series", PhD thesis, Institute of Statistics and Econometrics, Christian Albrechts University, Kiel, Germany.
- Poteshman, A.M. (2001), "Underreaction, Overreaction, and Increasing Misreaction to Information in the Option Market", *Journal of Finance*, 56, 3, 851-876.

- Pownall, R. A.J. and Koedijk, K.G. (1999), "Capturing downside risk in financial markets: the case of the Asian Crisis, *Journal of International Money and Finance*, **18**, 853-870.
- Rawnsley, J. (1995), "*Total Risk: Nick Leeson and the Fall of Barings Bank*", Harper, New York.
- RiskMetrics™ (1996), J.P. Morgan, *Technical document*, fourth edition.
- Risk Standards Working Group (1996), *Risk Standards for Institutional Investment Managers and Institutional Investors*.
- Ritchken, P. and R. Trevor (1999), "Pricing options under generalized GARCH and stochastic volatility processes", *Journal of Finance*, **54**, 377-402.
- Schwarz, G. (1978), "Estimating the dimension of a model", *Annals of Statistic*, **6**, 2, 461-464.
- Schwert, G.W. (1989), "Business Cycles, Financial Crises, and Stock Volatility", *Carnegie-Rochester Conference Series on Public Policy*, **39**, 83-126.
- Shimko, D. (1993), "Bounds of probability", *RISK*, **6**, 33-37.
- Stein, E. and J. Stein (1991), "Stock price distributions with stochastic volatility: an analytic approach", *Review of Financial Studies*, **4**, 727-752.
- Theodossiou, P. (2001): "Distribution of Financial Asset Prices, the Skewed Generalized Error Distribution, and the Pricing of Options", School of Business, Rutgers University, Working Paper.
- White, H. (1982), "Maximum likelihood estimation of misspecified models, *Econometrica* **50**, 1-25.
- Wiggins, J. (1987), "Option values under stochastic volatility: theory and empirical estimates", *Journal of Financial Economics*, **19**, 351-372.
- Wolff, C.C.P. (1987), "Time-varying parameters and the out-of-sample forecasting performance of structural exchange rate models", *Journal of Business and Economic Statistics*, **5**, 87-97.

Nederlandse samenvatting

Dit onderzoek reikt een aantal empirische methodes aan waarmee extreme gebeurtenissen in financiële markten gemodelleerd kunnen worden. De implicaties voor financieel risicobeheer worden hierbij onderzocht. De hoeveelheid literatuur over het beheren en kwantificeren van financieel risico is gigantisch en beschouwt vele interessante onderwerpen. Iedere statistische benadering beroept zich op een aantal assumpties over het prijsproces van de onderliggende waarde, waarbij men natuurlijkerwijs het risico loopt dat het stochastisch proces verkeerd gespecificeerd is. Verscheidene generalisaties voor het proces van rendementen worden veelvoudig beschouwd, doch de stochastische component wordt meestal normaal verondersteld. Echter onze empirische resultaten wijzen erop dat er meer extreme prijsschommelingen zijn dan verklaard kunnen worden door een lognormale verdeling. Zodoende is het van belang dat de stochastische component correct gedefinieerd wordt alvorens men de dynamiek van het volatiliteitsproces kan beschouwen. Daarbij wordt de asymmetrische uitwerking van volatiliteit op veranderingen in de onderliggende waarde empirisch beschouwd. Naast het bekende hefboomeffect vinden wij dat scheefheid in de stochastische component de asymmetrie in het volatiliteitsproces verder kan verklaren. Wij tonen aan dat de aanname van een niet-normale stochastische component niet enkel relevant is voor het kwantificeren van financieel risico, maar ook toepasbaar is voor optie waarderingsmethodes.

De keuze voor een meer verfijnde verdeling met bijzondere staarteigenschappen, die de extreme fluctuaties beschrijven, is dikwijls problematisch, aangezien we deze extreme veranderingen zelden observeren. Daarom is het aantal observaties in de staart van de rendementsverdeling beperkt. Er is veel onzekerheid over de werkelijke dikheid van de staarten en de parameters die deze dikstaartigheid meten. We laten de impact van schattingsrisico op de meting van marktrisico expliciet zien in termen van een Value-at-Risk schatting. Verder hebben we ons toegelegd op de informatie waarde van impliciete

indexvolatiliteit in de context van het voorspellen van markt risico's. Voorgaande studies trokken tegensprekende conclusies over de informatiele efficiëntie van optie markten. Door een nieuwe aanpak te volgen vinden wij een bewijs dat er meer voorspellingsinformatie is in optieprijsen en slechts weinig informatie in historische indexrendementen.

In hoofdstuk 2 beschouwen wij een algemene klasse van tijdsvariërende volatiliteitsmodellen, de "generalized augmented GARCH" processen, om de dynamiek van financiële prijzen te modelleren en wij vergelijken de verschillende specificaties van de stochastische component. Voorgaand onderzoek suggereert dat de verschillen in conditionele volatiliteitsschattingen omvangrijk zijn tussen de verschillende GARCH-specificaties, waar het hefboomeffect gemodelleerd wordt door óf een verschuiving óf een rotatie van de "news impact curve". Het meest uitgebreide geanalyseerde model houdt ook rekening met conditionele dikstaartigheid en scheefheid. Door het toelaten van een dikstaartige stochastische component voor het rendementsproces is men a-priori onzeker over het bestaan van de momenten. Daarom richten we ons tot het beschrijven van de dynamiek van de locaties en de schaalparameter van de verdelingen, in plaats van het modelleren van eerste en tweede momenten van het rendementsproces. Een Lagrange Multiplier gebaseerd op het "augmented GARCH" proces was ontwikkeld om zo verschillende modellen te kunnen verwerpen. Voor bepaalde index portefeuilles vonden wij dat de toevoeging van meer geraffineerde staarteigenschappen voor de conditionele verdeling tot verschillende conclusies over de correcte vorm van asymmetrie in het volatiliteitsproces leidde. In alle tijdreeksen is de geobserveerde scheefheid in de data een gevolg van een significant hefboomeffect *en* scheefheid in de stochastische component. Onder de hypothese van conditionele scheefheid worden modellen met verschillende aanpassingen van de "news impact curve" verworpen voor verschillende markten. Desalniettemin laten grafieken waar we quantielen vergelijken zien, dat de asymmetrische Student-t verdeling nog niet flexibel genoeg is om het gedrag van volatiele processen te beschrijven.

In het derde hoofdstuk reiken we een alternatieve verdeling voor de stochastische component aan, die ook een methode omvat waarmee meerdaagse rendementsverdelingen kunnen worden afgeleid uit de eendaagse verdeling. Een stochastisch proces dat oorspronkelijk ontwikkeld is in de natuurkunde literatuur wordt succesvol toegepast op het gebied van financiering. We stellen een generalisatie van het bekende RiskMetricsTM aanpak voor, voor de Value-at-Risk van een aandelenportefeuille. We onderzoeken een nieuwe

benadering om het prijsproces optimaal te beschrijven en stellen een nieuwe schalingsregel voor om volatiliteit te voorspellen: de locatie- en schaalparameter (volatiliteit) worden geschat op een tijdschaal (dagelijks), en de meerdaagse (wekelijkse, tweewekelijkse of maandelijkse) waarden worden afgeleid door de bi-fractionele schalingseigenschappen van de truncated Lévy processen. Deze methode heeft het voordeel dat we in staat zijn om de relatie tussen rendementsverdelingen voor verschillende tijdschalen te identificeren, door het analyseren van de tijdreeks van rendementen op een bepaald tijdsinterval (zeg dagelijks). Een interessante empirische eigenschap van financiële data is dat jaarlijkse volatiliteit op basis van dagelijkse intervallen beduidend lager is dan voor maandelijkse intervallen. Dit kan een gevolg zijn van korte termijn positieve seriële correlatie in de dagelijkse data tijdreeks. Als prijsveranderingen niet onafhankelijk zijn door de tijd, dan wordt de geschatte volatiliteit beïnvloedt. Positieve autocorrelatie, dat zich manifesteert als geobserveerde prijzen de nieuwe informatie langzaam absorberen, leidt tot een lagere geschatte volatiliteit. Dit probleem verdwijnt grotendeels als de waarnemingsintervallen toenemen, hetgeen de reden is om de rendementsverdelingen op maandelijkse basis in plaats van dagelijkse basis te beschouwen als de voorspellingshorizonnen toenemen. Hierdoor kan onze methode waarschijnlijk betere voorspellingen maken dan methodes waarbij eendaagse voorspellingen worden opgeteld om meerdaagse voorspellingen af te leiden. Zowel de prestaties van het model binnen de sample als buiten de sample voor VaR berekeningen zijn zeer belovend. Er is een legio aan praktische voordelen van dit model, die in toekomstig onderzoek bekeken kunnen worden.

In het vierde hoofdstuk richten wij ons op de VaR-modelselectie voor wisselkoersposities. We behaalden een aantal interessante resultaten: De wisselkoersfluctuaties die we bekeken vertoonden dikstaartig gedrag, doch het aantal extreme waarnemingen is beperkt. Modellen die dit staartgedrag in acht nemen zijn noodzakelijk, aangezien de VaR anders tot een onderschatting van het neerwaarts risico van de wisselkoersportefeuille leidt. Van de andere kant veronderstellen de meer geraffineerde aanpakken, die gebruik maken van dikstaartige verdelingen met slow power-law decay in de staarten, dat meer kans massa in de staart van de verdeling zit dan dat werkelijk het geval is. Omdat er zeer weinig extreme observaties zijn, wordt het werkelijke staartgedrag met een grote onzekerheid gemeten. In het algemeen geldt dat modellen met geraffineerdere staarteigenschappen leiden tot meer parameter onzekerheid, hetgeen leidt tot een grotere onzekerheid in de weergegeven VaRs. Als we onze VaRmaatstaf corrigeren voor deze parameter onzekerheid, krijgen we bijv. voor een GARCH(1,1)-Student-t model

verwerpingsratio's die dicht bij de betrouwbaarheidsintervallen van de weergegeven VaR liggen. Onder de modellen die geanalyseerd zijn, lijkt de GARCH(1,1)-Student-t specificatie met correctie het meest geschikt om de extreme verliezen voor wisselkoersposities te bepalen.

In het vijfde hoofdstuk combineren wij de informatie in historische rendementen en optieprijsen om te onderzoeken welke bron de meest waardevolle informatie verschaft voor het voorspellen van de volatiliteit van de DAX index. We vergelijken de voorspelling van een tijdreeks GARCH model met de voorspelling van een GARCH model wiens parameters zijn afgeleid van de gelijktijdige optieprijsen. Onze resultaten wijzen op het volgende. Ten eerste, de impliciete volatiliteitsvoorspelling Granger-causes de voorspelling van de tijdreeksvolatiliteit. Ten tweede, de optiemethode presteert beter dan de tijdreeks methode in out-of-sample voorspellingen, vergeleken met een realized volatiliteitsmaatstaf op basis van 5-minuuts intraday rendementen. Ten derde, beide maatstaven lijken enigszins foutieve voorspellers van realized volatiliteit. De grootste bijdrage in dit hoofdstuk is dat we een afwijkende doch theoretisch correcte methodiek introduceren, om de informatie die in optieprijsen zit te extraheren. Wij laten zien dat optieprijsen over meer informatie beschikken dan toekomstige volatiliteit, met in het bijzonder meer informatie over het toekomstige prijsproces. Daardoor kunnen optieprijsen als een maatstaf voor markt specifiek risico beschouwd worden, aangezien alle praktische maatstaven, zoals VaR, vooruitziende maatstaven zijn en optieprijsen zouden informatiever kunnen zijn. Voorgaande studies proberen informatie te benaderen vanuit een impliciete volatiliteitsindex gebaseerd op Black-Scholes impliciete volatiliteit van verhandelde opties, voor het schatten en voorspellen van toekomstige volatiliteit zonder het onderliggende rendementsproces expliciet te modelleren. Los van de theoretische kritieken, geloven wij dat het negeren van smile en termijnstructuur informatie bij Black-Scholes methodes hun voorspellingskracht schaden, vooral als we naar een langere tijdshorizon kijken.

Het laatste hoofdstuk legt een alternatieve oplossing tot optiewaardering voor. De hypothese van een scheve, dikstaartige stochastische component in het rendementsproces wordt in een conditioneel optiewaarderingsframework getoetst. In een in-sample analyse laten we zien dat conditionele scheefheid en dikstaartigheid meer flexibiliteit geeft aan het GARCH optiewaarderingsmodel. Empirische analyse van DAX index opties laat een verbetering in out-of-sample waarderingen zien vergeleken met alternatieve methodes. Er bestaan reeds bewijzen dat optiewaarderingsmodellen de geobserveerde scheefheid in de volatility smile

niet kunnen voorspellen, waardoor de grootste waarderingsfouten worden gemaakt voor opties die (ver) out-of-the-money zijn. De specifieke vorm van de verdeling van innovaties die wij beschouwen, resulteert in substantiële waarderingsverbeteringen voor korte termijn out-of-the-money opties. Vanuit onze analyse kunnen wij concluderen dat conditionele scheefheid en dikstaartigheid voor de aandelen rendementsverbeteringen, en de correlatie van volatiliteit met spotrendementen tezamen leiden tot een betere verklaring voor de vorm van de volatility smile. Dit komt overeen met de resultaten van hoofdstuk 2 en 3.

1. The first part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

2. The second part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

3. The third part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

4. The fourth part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

5. The fifth part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

6. The sixth part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

7. The seventh part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

8. The eighth part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

9. The ninth part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

10. The tenth part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

11. The eleventh part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

12. The twelfth part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

13. The thirteenth part of the text discusses the importance of maintaining accurate records of all transactions, including sales, purchases, and expenses. It emphasizes that these records are essential for determining the correct amount of tax liability.

Curriculum Vitae

Thorsten Lehnert was born on June 23rd, 1970 in Neuwied, Germany. He studied Economics at the Rheinische Friedrich Wilhelms University in Bonn, Germany. In 1998 he finished his Master's thesis, "The role of return distributions of financial price fluctuations in the modern theory of option pricing", and joined the Ph.D. program in Economics at the Universitat Carlos III in Madrid, Spain. In 1999, Thorsten moved to Maastricht University to work for the Limburg Institute of Financial Economics (LIFE). During the period January 1999 until August 2002 he completed his work on his Ph.D. dissertation on extreme events in financial risk management. He has presented his work at various international conferences. Since January 2002, he is employed as an Assistant Professor in Finance at Maastricht University. Thorsten is currently involved in further research, as well as undergraduate teaching.

