Value at Risk & vega risk

- using the implied volatility surface to compute VaR for option positions.

Mathias Barkhagen

24th January 2002

Acknowledgements

The work has been done at Bankaktiebolaget JP Nordiska (former Nordiska Fondkommission AB). I would like to thank Pär Stjernberg, Risk Manager at JP Nordiska, for being supportive and helpful during the process, and for giving me the opportunity to do my master thesis at JP Nordiska.

I would also like to thank Senior Lecturer Gunnar Englund, at the Department of Mathematics at the Royal Institute of Technology, for his support and help with the statistical issues in this thesis.

Stockholm, January 2002

Mathias Barkhagen

Abstract

Option positions are exposed to a range of market risks. Delta and gamma risks are the exposures of an option position to changes in the prices of the underlying assets. Vega risk is the exposure of an option position to changes in the implied volatility of the option. This risk can be a large part of the risk in a portfolio containing options. This emphasises the need for a risk model, which correctly incorporates vega risk.

The treatment of vega risk is complicated by the prevalence of volatility smiles and term structures in option markets. The main object of this thesis is to develop a method, which incorporates these aspects of vega risk into a Value at Risk model. In the model proportional changes in implied volatility and the underlying asset price are assumed to be jointly log-normally distributed.

The model is then implemented on option data from the Stockholm Exchange. A general backtesting method is developed in order to be able to verify the accuracy of the model and to identify its possible shortcomings.

Contents

1 Introduction

Financial risk is financial exposure to uncertainty. The goal of risk management is to improve a firm's financial result and to ensure that the firm's losses do not become unacceptably large. Risk management therefore involves being aware of and understanding the risks that face the firm, and being able to measure and control these risks. Financial risks which faces a firm can be broadly divided into the following categories:

• *Business risks* are those risks specific to the industry and market in which the firm operates.
• *Market risks* are the risks of losses arising from adverse movements in market prices (e.g. equity or option prices) or market rates (e.g. interest or exchange rates).
• *Credit risks* are the risks of loss arising from the failure of a counterparty to make a promised payment.
• *Liquidity risks* are the risks arising from the cost or inconvenience of unwinding a position.
• *Operational risks* are risks arising from the failure of internal systems or the people who operate them.
• *Legal risks* are risks arising from the prospect that contracts may not be enforced.

In this master thesis only market risks will be considered.

One approach to market risk management, known as Value at Risk (VaR), has gained increasing acceptance in the last five or so years. VaR originated on derivatives trading desks and then spread to other trading operations. It is a measure of risk based on a probability of loss and a specific time horizon in which this loss can be expected to occur. VaR has become an accepted standard in the banking industry and it forms the basis of bank capital requirements for market risk.

For assets which does not depend on the price of an underlying asset, such as stocks or exchange rates, the VaR methodology is fairly straightforward. However, when a portfolio includes assets which depends non-linearly on the price of an underlying asset, such as options, accurate VaR-measures become more difficult to perform. An option position does not only depend non-linearly on the price of the underlying asset, but also on the so called implied volatility of the option. An option's exposure to the change in implied volatility is called vega risk. For a portfolio containing options, vega risk can be a significant component of the portfolio's risk. This is especially true for portfolios whose delta exposures have been hedged. The main object of this thesis is to develop[1] and implement a method which includes the different aspects of vega risk in a standard VaR model.

[1] The method which is developed in this thesis is based on Malz [16].

2 Value at Risk

As an introduction to VaR an overview is given of some of the standard methods, described in the literature, used for computing VaR.

2.1 Value at Risk basics

Suppose that the object is to measure the risk of a portfolio containing different financial assets. One approach to measure the risk in the portfolio is through Value at Risk, which is defined as:

Definition:
Value at Risk, VaR, is defined as the maximum expected loss over some specified time-horizon with a given confidence level.

Typically the timehorizon is chosen to be one day and the confidence level to be 95%. Then VaR is defined as the loss that is expected not to be exceeded in more than 5% of the outcomes. That is, VaR is taken to be the appropriate quantile from the profit/loss-distribution of the portfolio.

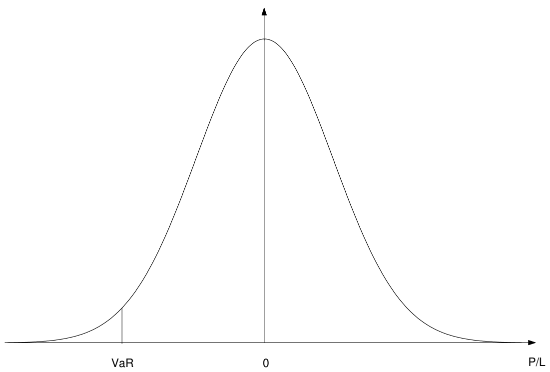

Figure 1: VaR from the P/L-distribution.

To compute the VaR of a portfolio, define W_0 as the initial investment, and R as its rate of return. Then the portfolio value at the end of the target horizon is $W = (1+R)W_0$. Now define the lowest portfolio value at the given confidence level α as $W^* = (1+R^*)W_0$. The absolute and relative VaR is now defined to be:

$$VaR_{absolute} = -R^* W_0. \tag{1}$$

That is, the absolute VaR is defined to be the maximum likely loss, in monetary units, to occur in (1-α) % of cases. The negative sign arises because we can assume that the extreme tail returns are losses, and the convention is to give VaR in positive numbers. The relative VaR is similarly defined to be the maximum likely loss relative to the mean return of the portfolio:

$$VaR_{relative} = -R^*W_0 + \mu W_0, \tag{2}$$

where μ is the mean return of the portfolio. Now suppose that there exists a stochastic process that describes the behaviour of the portfolios daily return. Suppose that the portfolio return at time t has a density function $f_t(R_t)$. The probability of a return less than R^* at time t and with the confidence level α is:

$$P(R_t < R^*) = \int_{-\infty}^{R^*} f_t(R_t)dR_t = 1 - \alpha. \tag{3}$$

In order to calculate VaR, some assumption must be made about $f_t(R_t)$. In practise, it is often assumed that $f_t(R_t)$ represents a normal distribution. Naturally, the choice is free to make other assumptions about the density function. However, the assumption of normality has the advantage of making calculations simple.

2.2 Methods for computing VaR for linear contracts

There are three standard types of methods used for computing VaR: Parametric methods, historical simulation methods and Monte Carlo methods. In this section, only VaR for linear contracts, e.g. stocks, will be considered.

2.2.1 Parametric VaR

Assume that the portfolio return at time t, R_t, is normally distributed with mean μ_t and volatility σ_t. The probability of a return less than R^* is given by:

$$P(R_t < R^*) = P\Big(Z_t < \frac{R^* - \mu_t}{\sigma_t}\Big) = 1 - \alpha \quad ; \quad Z_t \in N(0,1). \tag{4}$$

This implies that $(R^* - \mu_t)/\sigma_t = -\lambda_\alpha$, where λ_α is the α-quantile of the $N(0,1)$-distribution. This gives:

$$R^* = \mu_t - \lambda_\alpha \sigma_t. \tag{5}$$

Now plug (5) into (1) and (2) to arrive at parametric expressions for VaR:

$$VaR_{absolute} = \lambda_\alpha \sigma_t W - \mu_t W. \tag{6}$$

$$VaR_{relative} = \lambda_\alpha \sigma_t W. \tag{7}$$

In practise, instead of working with two different types of VaR, one would only work with one, the relative VaR. When using a parametric approach to VaR, the relative VaR is easier to handle since it doesn't require that the mean return μ_t is known. Also, if the time horizon is short, the difference between the absolute and relative VaRs will be small. From now on only the relative VaR will be considered.

Portfolio VaR

Suppose that the portfolio consists of n assets, S_i, $i = 1, 2, ..., n$, each with a relative amount w_i, $i = 1, 2, ..., n$, where:

$$\sum_{i=1}^{n} w_i = 1. \tag{8}$$

Then the portfolio return R is given by:

$$R = \sum_{i=1}^{n} w_i r_i, \tag{9}$$

where r_i is the relative return of asset i. Now suppose that the variance of the return of asset i is $V(r_i) = \sigma_i^2$, and that the correlation between the returns of assets i and j is ρ_{ij}, where $\rho_{ij} = \rho_{ji}$. Then the variance of the portfolio return, σ_p^2, is given by:

$$V(R) = \sigma_p^2 = V(\sum_{i=1}^{n} w_i r_i) = \sum_{i=1}^{n} w_i^2 \sigma_i^2 + 2 \sum_{i=1}^{n} \sum_{j<i} w_i w_j \rho_{ij} \sigma_i \sigma_j. \tag{10}$$

Assume that the returns of the individual assets are jointly normally distributed, then also the weighted sum of the returns is normally distributed, and the VaR of the portfolio is:

$$VaR_p = \lambda_\alpha \sigma_p W = \lambda_\alpha W \sqrt{\sum_{i=1}^{n} w_i^2 \sigma_i^2 + 2 \sum_{i=1}^{n} \sum_{j<i} w_i w_j \rho_{ij} \sigma_i \sigma_j}. \tag{11}$$

This can be expressed more compactly using matrix notation. Let $\boldsymbol{C}$ be the $n \times n$ covariance matrix of the asset returns, $\boldsymbol{w}$ the $n \times 1$ weight vector. Then the variance of the portfolio return is:

$$\sigma_p^2 = \boldsymbol{w^T C w}, \tag{12}$$

and the portfolio VaR is:

$$VaR_p = \lambda_\alpha W \sqrt{\boldsymbol{w^T C w}}. \tag{13}$$

2.2.2 Monte Carlo simulation

The Monte Carlo method estimates VaR on the basis of simulation results. The idea is to simulate repeatedly the stochastic processes governing the financial instruments of the portfolio. In order to do this, one must take a stand on the modelling of these stochastic processes. Each simulation will yield a possible return for the portfolio at the end of the target horizon. In repeating this procedure, a distribution of the portfolio return is created. Given that the chosen model describes the true evolution of the portfolio, this simulated distribution will converge to the distribution of the portfolio return.

The first step in Monte Carlo simulation consists of choosing a stochastic model for the evolution of stock prices. A commonly used model is the geometric Brownian motion. The model assumes that innovations in the stock price are uncorrelated over time and is described by:

$$dS_t = \mu_t S_t dt + \sigma_t S_t dW_t, \tag{14}$$

where W_t is a standard Wiener process with $E(dW_t) = 0$ and $V(dW_t) = dt$. The parameters μ_t and σ_t represent the instantaneous drift and volatility at time t. For simplicity, assume that these parameters are constant over the chosen time horizon. Integrating equation (14) over $[t_0\ ,\ t]$ gives the stock price at time t:

$$S_t = S_0 \exp\big((\mu - \frac{1}{2}\sigma^2)(t - t_0) + \sigma(W_t - W_{t_0})\big), \tag{15}$$

where S_0 is the known stockprice at time t_0.

The next step is to discretize (15) in order to simulate the evolution of the stockprice over the target time horizon. Divide the target time interval τ into n equal steps by creating the discrete times $[t_0, t_1, ..., t_n]$, where $\tau = t_n - t_0$. Then the stock price at time t_{i+1} given the stock price at time t_i is expressed as:

$$S_{t_{i+1}} = S_{t_i} \exp\big((\mu - \frac{1}{2}\sigma^2)\Delta t_i + \sigma \Delta W_i\big), \tag{16}$$

where:

$$\Delta t_i = t_{i+1} - t_i, \tag{17}$$

and

$$\Delta W_i = W_{t_{i+1}} - W_{t_i}. \tag{18}$$

Replace ΔW_i with $\epsilon_i \sqrt{\Delta t_i}$, where $\epsilon_i \in iidN(0,1)$. Then substitute for ΔW_i in (16) and arrive at the discretization of the time evolution of the stock price:

$$S_{t_{i+1}} = S_{t_i} \exp\big((\mu - \frac{1}{2}\sigma^2)\Delta t_i + \sigma \epsilon_i \sqrt{\Delta t_i}\big). \tag{19}$$

By creating n pseudo random numbers ϵ_i and using (19) the stock price at time t_n, S_{t_n}, is simulated. This procedure is then repeated N times until a stable distribution of S_{t_n} is reached.

VaR at confidence level α is then calculated as the $(1 - \alpha)N$ lowest value from this distribution.

Monte Carlo simulation for portfolios

The Monte Carlo method can be easily extended to the case of a portfolio containing several stocks. Since the stocks in the portfolio are correlated, pseudo random numbers that are correlated must be created for the simulation.

Suppose that the portfolio P consists of M stocks, $S_j, j = 1, ..., M$, with returns $r_j, j = 1, ..., M$. Let $Corr(r_i, r_j) = \rho_{ij}$ and $Cov(r_i, r_j) = \sigma_{ij}\Delta t$. The dynamics of the stock prices is given by:

$$dS_{t,j} = \mu_j S_{t,j} dt + \sigma_j S_{t,j} dW_{t,j}, \tag{20}$$

and in discrete form

$$\Delta S_{t_i,j} = \mu_j S_{t_i,j} \Delta t_i + \sigma_j S_{t_i,j} \epsilon_{t_i,j} \sqrt{\Delta t_i}, \tag{21}$$

where $\epsilon_{t_i,j} \in N(0,1)$. The covariance of the returns $r_{t,j}$ and $r_{t,k}$ is:

$$\begin{aligned} Cov(r_{t_i,j}, r_{t_i,k}) &= Cov\left(\frac{\Delta S_{t_i,j}}{S_{t_i,j}}, \frac{\Delta S_{t_i,k}}{S_{t_i,k}}\right) = Cov(\sigma_j \epsilon_{t_i,j}\sqrt{\Delta t_i}, \sigma_k \epsilon_{t_i,k}\sqrt{\Delta t_i}) \\ &= \Delta t_i \sigma_j \sigma_k Cov(\epsilon_{t_i,j}, \epsilon_{t_i,k}), \end{aligned} \tag{22}$$

and consequently $Cov(\epsilon_{t_i,j}, \epsilon_{t_i,k}) = Corr(r_{t_i,j}, r_{t_i,k}) = \rho_{t_i,jk}$. Let $\boldsymbol{R}$ be the correlation matrix, at time t_i, of the stock returns, and let $\boldsymbol{\epsilon}$=[$\epsilon_{t_i,1}$, $\epsilon_{t_i,2}$, ...,$\epsilon_{t_i,M}$]. Then $V(\boldsymbol{\epsilon})$=$E(\boldsymbol{\epsilon\epsilon^T})$=$\boldsymbol{R}$, and $E(\boldsymbol{\epsilon}) = \mathbf{0}$.

In order to create pseudo random numbers with these properties, start with the $1 \times M$ vector $\boldsymbol{\eta}$, which is composed of $iidN(0,1)$ variables. Then $V(\boldsymbol{\eta}) = \boldsymbol{I}$, where $\boldsymbol{I}$ is the $M \times M$ identity matrix, and $E(\boldsymbol{\eta}) = \mathbf{0}$. As the matrix $\boldsymbol{R}$ is symmetric, real, and positive semi-definite, it can be decomposed into its Cholesky factors $\boldsymbol{R} = \boldsymbol{TT^T}$, where $\boldsymbol{T}$ is a lower triangular matrix with zeros on the upper right corners. Now, construct the variable $\boldsymbol{\epsilon} = \boldsymbol{T\eta}$. Its covariance matrix is $V(\boldsymbol{\epsilon}) = E(\boldsymbol{\epsilon\epsilon^T}) = E(\boldsymbol{T\eta\eta^T T^T}) = \boldsymbol{T}E(\boldsymbol{\eta\eta^T})\boldsymbol{T^T} = \boldsymbol{TIT^T} = \boldsymbol{TT^T} = \boldsymbol{R}$. And thus, the pseudo random numbers $\boldsymbol{\epsilon}$ have the desired correlations.

The Monte Carlo simulation in the multidimensional case is now carried out in the same way as in the one dimensional case, using the correlated vector $\boldsymbol{\epsilon}$ in the simulation process.

2.2.3 Variance reduction techniques

For most Monte Carlo-simulation procedures, the accuracy of any estimated parameter will be proportional to $1/\sqrt{N}$, where N is the number of iterations. The Monte Carlo approach is thus very computer-intensive. In addition, the number of calculations required also depends on the number of random drawings in each iteration. A simple linear position in n imperfectly correlated assets will typically depend on the realisations of n random variables, and non-linear or more complex positions will often depend on even more. Hence, the value of a position of 50 assets would generally depend on the realisations of at least 50 random variables. If, for example, 10000 iterations is needed in order to get results of acceptable accuracy, one would therefore have to construct at least 500000 hypothetical price paths. If one wishes to use Monte Carlo-simulation when operating under real-time constraints, accuracy may have to be sacrificed for speed.

However, it is possible to obtain improvements in accuracy (and/or speed) by using one of a number of variance-reduction techniques. The standard variance reduction techniques include: the antithetic variable technique, the control variate technique, importance sampling, and moment matching. The Monte

Carlo-methods used in this master thesis was performed without variance reduction. For details about the different techniques, see e.g. Hull [1] or Moix [17].

2.2.4 Historical simulation

The historical simulation approach is based on the historical distribution of the portfolio returns. Suppose that VaR over the time horizon τ is to be calculated and that the current time is T. The historical returns of the stocks in the portfolio over the specified time horizon is collected. Then it is assumed that this historical distribution of returns is also a good approximation for the distribution of returns at time $T + \tau$. The relevant percentile from the distribution of historical returns is then the VaR for the current portfolio.

Suppose that observations are running from period 0 to period T. Let $R_{i,t}$ be the return over τ of asset i at time t. That is:

$$R_{i,t} = \frac{S_{i,t} - S_{i,t-\tau}}{S_{i,t-\tau}}. \tag{23}$$

Let $w_i, i = 1, 2, ..., n$, be the relative weight of asset i in the portfolio at time T. Then the portfolio return R_t^p at time t is:

$$R_t^p = \sum_{i=1}^{n} w_i R_{i,t}, \qquad t = \tau, ..., T. \tag{24}$$

Each observation will give a particular portfolio return R_t^p. The sample of historical observations will then yield an empirical distribution of portfolio returns. Then the portfolio returns are translated to portfolio profits and losses, and the VaR is read off from the histogram of profits and losses.

The main attractions of the historical simulation approach is that it is easy to understand and implement. Also the correlations between the assets are captured within the method, and no assumption of normality has to be made and therefore fat tails are included. However the method relies on the assumption that the past represents the immediate future. The choice of the sample period reflects a trade-off between accuracy and the dynamics of the method. Longer sample periods increase the accuracy of estimates but could use irrelevant data, and thereby missing important changes in the underlying process. However, if the sample size is too small the VaR may be subject to a large estimation error. Also, risk contains significant and predictable time variation. The historical simulation method will miss situations with temporarily elevated volatility.

2.2.5 Filtered Historical simulation

The idea behind filtered historical simulation (FHS) is to adopt the historical simulation method to the current volatility situation. Historical simulation depends on the assumption that stock returns are independent and stationary. The independence implies that the size of price movement in one period will not influence the movement of any successive prices. Stationarity implies that the probability of occurrence of a specified loss is the same for each day. Using a constant volatility model to calculate VaR when the distribution of returns is not stationary, as it is in the case with most daily financial time series, could be misleading. In fact, the probability of having a large loss is then not equal across different days. During days with higher volatility one would expect larger than usual losses. Empirical studies have found that there is a tendency, known as volatility clustering, for large price changes to be followed by more large changes.

The above limitations of historical simulation, a limited set of outcomes and unresponsiveness to changes in market volatility, can be overcome with the use of filtered historical simulation. In filtered historical simulation the stationarity assumption is relaxed. Historical returns are first standardised by the volatility estimated on that particular day:

$$\eta_t = \frac{r_t}{\sigma_t}, \tag{25}$$

where r_t is the return at day t and σ_t is the estimated volatility for day t. This filtering process yields approximately i.i.d. returns[2] suited for historical simulation. Before filtered returns are used as innovations they are scaled by the current conditional forecast of volatility, and thus they reflect current market conditions. Suppose that the current time is T, and that VaR with time horizon τ is to be calculated. Then historical simulation is applied to:

$$\epsilon_T = \eta_T \sigma_{T+\tau}, \tag{26}$$

where $\sigma_{T+\tau}$ is the forecast, made at time T, of the volatility at time $T + \tau$.

A major advantage of FHS over HS is that the filtering process increases the range of outcomes beyond the historical record through a change of scale. In other words, FHS provides a systematic approach to generate extreme events not present in the historical record, completing the tails of the distribution. FHS requires therefore a shorter historical record than HS to simulate the tails of the distribution of returns.

2.2.6 Bootstrap simulation

The purpose of Bootstrap simulation is to create a distribution for the VaR value. Let S be the sample of historical returns, r_i, $i = 1, 2, ..., N$. The idea is now to create a number of samples of the same size as S and then estimate the VaR value from each of these simulated samples. As a result an estimated distribution of the VaR value is created.

The empirical distribution of the return is created by assigning the same probability to each of the historical returns. Then a new sample of size N is created by drawing with replacement from the empirical distribution, and VaR is calculated from the sample. The procedure is repeated until a stable distribution for the VaR value is reached. The major advantage of this method is that it will provide information about the uncertainty in the estimated VaR.

2.3 Volatility estimation

When calculating VaR it is crucial to have good estimations of the volatilities of the risk factors and correlations between the risk factors. All the above described

[2] See Barone-Adesi & Giannopoulos [20].

methods for calculating VaR, except for historical simulation and bootstrap simulation, depends on accurate estimations of volatility and correlation.

2.3.1 Equally weighted estimation

The simplest measure of volatility is the sample variance. If r_t is the return in period t, and there are N observations, the variance is taken to be:

$$\hat{\sigma}^2 = \frac{1}{N-1}\sum_{i=1}^{N}(r_i - \overline{r})^2, \qquad \overline{r} = \frac{1}{N}\sum_{i=1}^{N} r_i. \tag{27}$$

The correlation between asset i and asset j is measured in a similar way. Let $r_{t,i}$ and $r_{t,j}$ be the returns of assets i and j in period t. Then the covariance of asset i and j is taken to be:

$$\hat{c}_{ij} = \frac{1}{N-1}\sum_{t=1}^{N}(r_{t,i} - \overline{r}_i)(r_{t,j} - \overline{r}_j). \tag{28}$$

The sample variance and covariance has the attraction of being easy to estimate, but has a number of drawbacks. The biggest problem is that it assumes that the underlying 'true' variance is constant, and so cannot accommodate any changes in the variance over time. In particular, it fails to incorporate the phenomenon of volatility clustering. In practise the volatility estimation will be based on, for example, the last year's historic returns. The volatility estimation for two following days will be based on the same historical data set, except that one of the components in the data set will be replaced. In this sense, the estimated volatility can be regarded as constant over a shorter time period, and will change only slowly over time.

Also, it ignores the temporal order of observations which implies that it ignores the dynamic information in the return series and gives earlier and more recent observations equal weight.

2.3.2 Exponentially weighted estimation

Equation (27) gives equal weight to all of the historic returns. Given that the object is to estimate the current level of volatility, it is appropriate to give more weight to recent data. A model that does this is:

$$\hat{\sigma}^2_{t+1} = \sum_{i=0}^{n}\alpha_i u^2_{t-i}; \qquad u_{t-i} = r_{t-i} - \overline{r}. \tag{29}$$

The variable α_i represents the weight given to the observation i days ago. Since less weight is given to older observations, $\alpha_i < \alpha_j$ when $i > j$. The weights sum to unity:

$$\sum_{i=0}^{n}\alpha_i = 1; \qquad \alpha_i > 0. \tag{30}$$

The exponentially weighted moving average (EWMA) model is a particular case of the model in equation (29), where the weights decrease exponentially with older observations. That is: $\alpha_{i+1} = \lambda\alpha_i$, where $\lambda \in (0,1)$ is a constant.

An attractive feature of the exponentially weighted estimator is that it can be written in recursive form. In order to derive the recursive form it is assumed that an infinite amount of data are available. Using $\alpha_{i+1} = \lambda\alpha_i$, and the fact that the weights sum to unity gives:

$$\sum_{i=0}^{\infty} \alpha_i = \sum_{i=0}^{\infty} \alpha_0 \lambda^i = \alpha_0 \sum_{i=0}^{\infty} \lambda^i = \alpha_0 \frac{1}{1-\lambda} = 1 \Rightarrow \alpha_0 = 1 - \lambda. \tag{31}$$

Substitute the α_i's with $(1-\lambda)\lambda^i$ in equation (29):

$$\begin{aligned} \hat{\sigma}_{t+1}^2 &= \sum_{i=0}^{\infty} (1-\lambda)\lambda^i u_{t-i}^2 = (1-\lambda)u_t^2 + \sum_{i=1}^{\infty} (1-\lambda)\lambda^i u_{t-i}^2 \\ &= \{j = i-1 \Rightarrow i = j+1\} = (1-\lambda)u_t^2 + (1-\lambda)\sum_{j=0}^{\infty} \lambda^{j+1} u_{t-(j+1)}^2 \\ &= (1-\lambda)u_t^2 + \lambda(1-\lambda)\sum_{j=0}^{\infty} \lambda^j u_{t-(j+1)}^2 = (1-\lambda)u_t^2 + \lambda\hat{\sigma}_t^2. \end{aligned} \tag{32}$$

Thus, the variance forecast of tomorrow is the weighted sum of today's variance and today's squared return. The covariances are estimated using the same weighting scheme:

$$\hat{c}_{t+1,ij} = (1-\lambda)u_{t,i}u_{t,j} + \lambda\hat{c}_{t,ij}. \tag{33}$$

The weight λ is chosen to minimise the error between the forecast and the measured volatilities over some period. J.P. Morgan uses the EWMA model with $\lambda = 0.94$ for updating daily volatility estimates in its RiskMetrics database. The company found that, across a range of different market variables, this value of λ gives forecasts of the variance rate that come closest to the realised variance rate.[3]

Compared with the equally weighted volatility estimation, the EWMA estimation will be much more responsive to changes in the market risk. Below is a plot of the estimated volatilities of Ericsson B and NOKIA over the period October 26^{th} 1995 until August 12^{th} 2001. The equally weighted estimation has been estimated with a rolling window of 100 data. That is, each volatility estimation is based on the 100 most recent returns. The EWMA estimates was calculated using a decay factor λ of 0.94.

[3] See Kim & Mina [5].

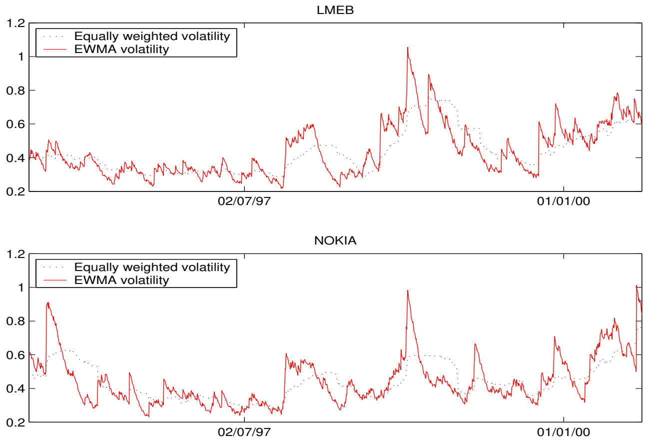

Figure 2: Volatility estimation

2.3.3 GARCH estimation

A more general class of volatility estimators is the GARCH (Generalised AutoRegressive Conditional Heteroskedastic) estimator. The generic GARCH model defines a volatility estimator that depends on both past values of squared returns and past volatility estimates. Let as before r_t be the return at time t. Then the GARCH(p,q) estimator is defined as:

$$\hat{\sigma}_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i r_{t-i}^2 + \sum_{i=1}^{q} \beta_i \hat{\sigma}_{t-i}^2, \tag{34}$$

where $\alpha_0 > 0$, and $\alpha_j, \beta_j \geq 0, j = 1, 2, \ldots$.

The GARCH framework allows for both time-varying volatility and volatility clustering. It has proven to be a flexible model that performs well empirically. The framework also encompasses a variety of specific models. These include the asymmetric GARCH (or AGARCH) model, which allows for skewness and kurtosis in the return distribution; and the exponential GARCH (or EGARCH) which formulates the conditional variance equation in logarithmic terms. The GARCH model also includes the EWMA as a special-case of the GARCH(1,1) model:

$$\hat{\sigma}_t^2 = \alpha_0 + \alpha_1 r_{t-1}^2 + \beta_1 \hat{\sigma}_{t-1}^2. \tag{35}$$

If $\beta_1 = \lambda, \alpha_1 = (1 - \lambda)$, and $\alpha_0 = 0$ then the GARCH(1,1) estimation and the exponentially weighted estimation are identical.

However, the GARCH model also has its drawbacks. It can be difficult to implement. In the GARCH(1,1) model 3 parameters must be estimated compared with one for the EWMA model. Also, it needs a large number of observations to get reliable eatimates, and these might not be available. Furthermore GARCH models can be unstable, and therefore produce unreliable forecasts[4].

One can also use multivariate GARCH procedures to forecast correlations, but the number of parameters involved grows exponentially with the number of correlations. The number of parameters makes estimates potentially unreliable. Multivariate GARCH procedures are therefore ill suited to estimates of high-dimension covariance matrices.

2.4 Model verification

2.4.1 Stresstesting

While VaR models have proven themselves to be very useful risk management tools, it is also important to realise that the models have limitations - in particular their excessive dependency on history or unrealistic statistical assumptions. The natural response to these limitations is to resort to stress tests to complement the results of the VaR analyses. Stress tests are exercises to determine the losses that might occur under unlikely but plausible circumstances. This method examines the effect of simulated large movements in key financial variables on the portfolio.

Stress tests can provide useful information about the risk exposure that VaR methods can easily miss, particularly if VaR models focus on 'normal' market risks rather than risks associated with rare or extreme events.

The essence of stress testing is the creation of user-defined scenarios, fed into a calculation procedure to produce estimates of the profits or losses that can be expected under these scenarios. Stress tests fall into three main types, which differ in how the scenarios are constructed.

• The first type uses scenarios from recent history, such as the 1987 equity crash, the ERM crises of 1992-1993, and the 1997 east Asian crisis and its aftermath. In these scenarios, the effect on profit or loss is simulated based on repeats of past extreme historical events.

• The second type uses predefined or set-piece scenarios that have proven to be useful in practise. For example one could simulate the impact on profit and loss, of a fall in the stock index by a predefined number of standard deviations.

• The third type are mechanical-search stress tests. These are automated routines to cover prospective changes in risk factors, evaluate profit and loss under each set of risk-factor changes, and report the worst-case results.

However, stress tests have a number of problems. Perhaps the most obvious is that stress tests are inevitably subjective because they depend on scenarios chosen by the stress tester. As a result, the value of stress testing depends critically on the choice of scenarios and therefore on the skill of the modeller. This subjectivity also makes it difficult to asses the value of a stress test objectively.

[4] See Dowd [2].

To sum up, stress tests should be used to complement the results of the specific VaR model used. It should test the effect on portfolio result of changes in the underlying risk factors, which are unlikely to be picked up by the VaR model.

2.4.2 Backtesting

Whatever the method used for calculating VaR, an important test of the validity of the method is backtesting. It involves testing how well the VaR estimates would have performed in the past. Suppose that VaR at confidence level α and time horizon one day is calculated. Backtesting involves comparing the calculated VaR values with the actual outcome of the portfolio for each day. If the method for calculating VaR is reliable one would then expect that the VaR value would be exceeded by the actual portfolio outcome in close to (1-α)% of the days. A more systematic approach to test the validity of the method is the likelihood ratio test proposed by Kupiec, which is explained in the next section.

2.4.3 Likelihood ratio test

The likelihood ratio test is the quantitative standard for model verification embodied in the 1996 amendment to the Basle Capital Accord.[5]

Suppose that the portfolio performance is monitored, then the day-to-day profits and losses determine the outcome of a binomial event. Either the portfolio loss is less than the ex ante VaR estimate (a success), or the loss exceeds the ex ante VaR estimate (a failure). Assuming that the VaR estimates are accurate, such failures can be modelled as independent draws from a binomial distribution with a probability of occurrence equal to $p = 1 - \alpha$, where α is the confidence level of the VaR estimate. The probability of observing x failures regardless of order in a sample of size n is:

$$P(X = x) = \binom{n}{x}(1-p)^{n-x}p^x; \qquad x = 0, 1, ..., n. \tag{36}$$

Let p_0 denote the modelled probability of failure of the VaR metod, and let p denote the underlying 'true' probability of failure of the VaR method. The null hypothesis to be tested is:

$$H_0 : p = p_0. \tag{37}$$

This is tested against the alternative hypothesis H_A:

$$H_A : p \neq p_0. \tag{38}$$

The maximum likelihood principal says to look for a state $\hat{\theta}$ that maximises the likelihood function $L(\theta) = f(\boldsymbol{X}, \theta)$, where $f(\cdot, \cdot)$ is the density function. It regards such state as providing the 'best' explanation of the observed $\boldsymbol{X}$. In choosing between H_0 and H_A, it would be natural to compare the best explanation within H_0 with the best explanation within H_A and reject H_0 if and only if:

[5] The Basle Capital Accord sets down the agreement among the G-10 central banks to apply common minimum capital standards to their banking industries.

$$\sup_{H_0} f(\boldsymbol{x}, \theta) < \sup_{H_A} f(\boldsymbol{x}, \theta), \tag{39}$$

where $\boldsymbol{X} = \boldsymbol{x}$ is observed. Generalising this idea to provide some flexibility in the level of the test, one would consider critical regions of the form:

$$\Lambda^* \equiv \frac{\sup_{H_0} f(\boldsymbol{x}, \theta)}{\sup_{H_A} f(\boldsymbol{x}, \theta)} < \text{ constant.} \tag{40}$$

Such a critical region defines the likelihood ratio test.

In this instance it is convenient to consider the modified likelihood ratio in which the best θ in H_0 is compared with the best θ overall, that is, with the best θ in $H_0 + H_A$:

$$\Lambda = \frac{\sup_{H_0} f(\boldsymbol{x}, \theta)}{\sup_{H_0+H_A} f(\boldsymbol{x}, \theta)} < \text{ constant.} \tag{41}$$

If a test is to be based on the statistic Λ, it is necessary to know the distribution of Λ. This is ordinarily very complicated, and it is useful to know that the asymptotic distribution of $-2\log(\Lambda)$ is chi-squared[6]. The parameter of this chi-square distribution is the difference in the dimension of $H_0 + H_A$ and the dimension of H_0.

In testing the VaR method, the best explanation under H_0 is $p = p_0$, and the best explanation under $H_0 + H_A$ is $p = \frac{x}{n} = p_*$. This gives the following test statistic:

$$-2\log(\Lambda) = 2\big(\log(p_*^x(1-p_*)^{n-x}) - \log(p_0^x(1-p_0)^{n-x})\big). \tag{42}$$

Under H_0 the statistic is χ^2-distributed with one degree of freedom.

Let $\alpha = P(\text{reject } H_0 \text{ when } H_0 \text{ true})$ and $\beta = P(\text{do not reject } H_0 \text{ when } H_0 \text{ false})$. Then α is called the TYPE I error of the test and β is called the TYPE II error. The non-rejection regions for the test with $\alpha = 0.05$ and alternative sample sizes are given by:

Null hypothesis probability	Non-rejection region for x, n=255 days	Non-rejection region for x, n=510 days	Non-rejection region for x, n=1000 days
0.010	$x < 7$	$1 < x < 11$	$4 < x < 17$
0.025	$2 < x < 12$	$6 < x < 21$	$15 < x < 36$
0.050	$6 < x < 21$	$16 < x < 36$	$37 < x < 65$
0.075	$11 < x < 28$	$27 < x < 51$	$59 < x < 92$
0.100	$16 < x < 36$	$38 < x < 65$	$81 < x < 120$

Table 1.

where x is the number of failures that could be observed in a sample size equal to the specified number of trading days without rejecting the indicated null hypothesis at the level $\alpha = 0.05$.

[6] See Kupiec [6].

The problem with the likelihood ratio test is that it has poor power characteristics in small samples. That is, in small samples the null hypothesis acceptance regions are large, so there is a significant probability that one will accept the null hypothesis when it is false, i.e. a TYPE II error. The conclusion, according to Kupiec, is that this test should be carried out on large samples of about 1000 data in order to be reliable.

2.4.4 Creating a confidence interval

Let α be the confidence level of the VaR-method to be tested. Then the probability of the actual portfolio loss to exceed VaR is $p = 1-\alpha$. An easier alternative to the likelihood ratio test is to create a confidence interval for the underlying 'true' probability, p, of the VaR method. Let I_p be the confidence interval for p at the level α_0. If $p_0 \in I_p$, then the hypothesis that $p = p_0$ can not be rejected at the level α_0.

Let x be the number of times that the actual portfolio loss exceeds VaR in a sample of size n. Then the stochastic variable $X \in Bin(n,p)$, where x is an observation of X. The 'true' probability p is estimated by $p^* = \frac{x}{n}$. For large n the following approximation holds:

$$\frac{X}{n} \in \text{approx.}N\left(p, \sqrt{\frac{p(1-p)}{n}}\right). \tag{43}$$

Since p is unknown, instead p^* is used to create an approximate confidence interval for p. At the level α_0 the cofidence interval is:

$$I_p = \left(p^* - \lambda_{\alpha_0/2}\sqrt{\frac{p^*(1-p^*)}{n}}, p^* + \lambda_{\alpha_0/2}\sqrt{\frac{p^*(1-p^*)}{n}}\right), \tag{44}$$

where $\lambda_{\alpha_0/2}$ is the α-quantile for the $N(0,1)$-distribution.

2.5 Implementation

The methods described in section (2.2) was implemented on a portfolio consisting of the following stocks: Ericsson B, Nokia, H&M B and Astra Zeneca. The parametric and the Monte Carlo VaR-calculations was performed both using the equally weighted volatility estimation and the EWMA volatility estimation. The equally weighted volatility estimation was based on a rolling window of 100 historical data, and the EWMA estimation was based on a rolling window of 75 historical data[7]. The historical VaR calculation and the filtered historic VaR calculation was based on 252 historical data.

The VaR calculations was performed on three different portfolios consisting of the four stocks. Let $\boldsymbol{w} = [w_1, w_2, w_3, w_4]$ denote the relative portfolio weights for a specific portfolio. That is, w_i, $i = 1, 2, 3, 4$, is the value of stock i divided by the total portfolio value. The VaR methods was then tested on the following portfolios:

[7] This is the number of historical data recommended by RiskMetrics for EWMA volatility estimation.

Portfolio	rel. weight Ericsson B	rel. weight Nokia	rel. weight H&M B	rel. weight Astra Zeneca
A	0.25	0.25	0.25	0.25
B	0.40	0.30	0.20	0.10
C	0.10	0.20	0.30	0.40

Table 2.

The VaR methods was backtested on the three portfolios based on 1000 backtest days. Below is a table of the failure rates for the different methods when tested on the three portfolios. By the failure rate is meant the rate by which the actual portfolio outcome exceeds the estimated VaR value. Method I is the parametric method with equally weighted volatility estimation, method II is the parametric method with EWMA volatility estimation, method III is the historical estimation, method IV is the Monte Carlo method with equally weighted volatility estimation, method V is the Monte Carlo method with EWMA volatility estimation, and method VI is the filtered historical estimation. The Monte Carlo methods were performed by using 1000 simulations for each backtest date.

Portfolio	I	II	III	IV	V	VI
A	0.0501	0.0529	0.0629	0.0486	0.0572	0.0429
B	0.0629	0.0501	0.0715	0.0644	0.0572	0.0415
C	0.0415	0.0458	0.0501	0.0372	0.0443	0.0501

Table 3.

The VaR methods tested in the implementation all used a 95% confidence level in the calculations. The likelihood-ratio test provides a method for validation of the confidence level of the VaR methods. For a sample size of 1000 backtest data the non-rejection region for the number of failures is: (37 , 65). This corresponds to the following non-rejection region for the failure rate: (0.037 , 0.065).

The null hypothesis to be tested is: $H_0 : p = 0.05$, where p is the underlying 'true' probability of failure of the VaR method. If the failure rate is outside the non-rejection region given above, then the method is considered to be inaccurate. If the failure rate for a given method is close to 0.05, then the method would pass the likelihood ratio test also for a larger TYPE I error, which in turn implies a smaller TYPE II error. A small TYPE II error implies that the probability of not rejecting H_0 when H_0 is false is small, and thus the null hypothesis can be considered to be plausible in this case. Therefore, a VaR model is considered to be more accurate, the closer to 0.05 the failure rate is[8].

As can be seen from table (3), only one method is rejected: the historical simulation method when applied to portfolio B. For both the parametric- and the Monte Carlo-methods, the EWMA volatility estimation generally gives a failure rate that is closer to 0.05, than is the case for the equally weighted volatility estimation. Also, the filtered historic simulation produces failure rates which are closer to 0.05 than the corresponding failure rates of the historic simulation.

[8]This reasoning may be a bit awkward since there is an uncertainty in the estimated VaR-values, and thus it is difficult to distinguish between the models if H_0 is not rejected.

3 Value at Risk for options

3.1 Properties of European and American options

The object of this section is to establish some properties of European and American options. This will be needed in order to evaluate the possible difficulties arising when implementing the VaR-method developed in this thesis. Some relationships between option prices that do not require any assumptions about volatility and the probabilistic behaviour of stock prices are derived. It is assumed that there are some market participants, such as large investment banks, for which the following hold:

i) There are no transaction costs.
ii) All trading profits are subject to the same tax rate.
iii) Borrowing and lending at the risk-free rate is possible.

It is assumed that these market participants are prepared to take advantage of arbitrage opportunities as they arise. This means that any available arbitrage opportunity disappear very quickly. It is therefore reasonable to assume that there are no arbitrage opportunities.

3.1.1 Definitions and notation

There are two basic types of options. A *call option* gives the holder the right to buy the underlying asset by a certain date, *the expiration date* , for a certain price, *the exercise price.* A *put option* gives the holder the right to sell the underlying asset by a certain date for a certain price. *American options* can be exercised at any time up to the expiration date. *European options* can be exercised only on the expiration date itself. Since an American option offers the opportunity of early exercise, its price should always be greater than or equal to the corresponding European option on the same asset. Most of the options that are traded on exchanges are American, and one contract is usually an agreement to buy or sell 100 shares of the underlying asset. Examples of underlying assets are stocks, stock indices, foreign currencies, debt instruments, commodities, and futures contracts.

The following notation will be used:

Let t be todays time.
T, the time of expiration.
S_τ, the stock price at time τ.
K, the strike price of the option.
r, the risk-free rate of interest at time τ for maturity at T (continuously compounded).
C_τ, the value at time τ of an American call option to buy one share.
P_τ, the value at time τ of an American put option to buy one share.
c_τ, the value at time τ of an European call option to buy one share.
p_τ, the value at time τ of an European put option to buy one share.

3.1.2 Value of European puts and calls at expiry

At the time of expiry the European put and call have the following values:

$$c_T = max(S_T - K, 0), \tag{45}$$

$$p_T = max(K - S_T, 0). \tag{46}$$

A call option is said to be *in the money* if $S_t > K$, *out of the money* if $S_t < K$ and *at the money* if $S_t = K$. Similarly, a put option is in the money if $S_t < K$, and out of the money if $S_t > K$.

3.1.3 Lower bound for European calls on non-dividend-paying stocks

Consider the following portfolios:

A: Buy the call option and a zero-coupon bond maturing at time T to K. Value: $c_t + Ke^{-r(T-t)}$.
B: Buy the stock. Value: S_t.

At the time of expiry portfolio A will be worth $max(S_T - K, 0) + K = max(S_T, K)$ and the value of the stock is S_T. Hence, portfolio A is worth at least as much as the stock at time T. Since the option is European it cannot be exercised prior to the expiration date. It follows that it must be worth at least as much as the stock today. Hence:

$$c_t + Ke^{-r(T-t)} \geq S_t, \tag{47}$$

or

$$c_t \geq S_t - Ke^{-r(T-t)}. \tag{48}$$

3.1.4 Put-call parity for European options

Put-call parity expresses a fundamental relationship between the value of a European put and a European call option of the same strike and expiration, and with the same stock as the underlying asset. The underlying stock is assumed to pay no dividend. Then put-call parity can be written as:

$$c_t - p_t = S_t - e^{-r(T-t)}K. \tag{49}$$

To see why this is true a basic arbitrage argument is used. Consider the following two portfolios

A: Buy a call option with strike K, and sell a put option with the same strike. Value: $c_t - p_t$.
B: Buy the stock and sell and sell a risk-less, zero-coupon bond maturing at time T to K. Value: $S_t - Ke^{-r(T-t)}$.

Neither of these investments incur any costs during their lifetime. Since the call and the put have the same strike, at expiration either the call will be in the money or the put will be in the money, but never both. If the call is in the money, the payoff is $S_T - K$, since the position is long. On the other hand, if the short put is in the money, then the payoff is $-(K - S_T) = S_T - K$. Consequently, the payoff on the put-call portfolio is $S_T - K$ at time T.

At time T, the bond will have matured to a value of K. Therefore the second investment, portfolio B, will have a value of $S_T - K$ at time T. Hence, the portfolios have the same value at the time of expiration. Since the options are European, they cannot be exercised prior to the expiration date. The portfolios must, therefore, have identical values today.

3.1.5 Early exercise of American options

This section demonstrates that it is never optimal to exercise an American call option on a non-dividend-paying stock before the expiration date. That is, the value of immediate exercise, the *intrinsic* value of the option, never exceeds the theoretical value of holding the option. This is demonstrated by using an arbitrage argument.

Start with an American call option on a non-dividend paying stock. Pick a random time t during the lifetime of the option and assume that the option is in the money at the time. Consider the two separate investments:

i) Exercise the option at time t.
ii) Hold the option until its expiration and exercise it.

First consider what each of these investments will be worth at time T. To finance the first investment, sell a bond maturing at time T and use the proceeds to purchase the stock for K. Thus, at expiration the total value of the first investment is: $S_T - e^{r(T-t)}K$.

On the other hand, holding the option until expiration is worth $max(S_T - K, 0)$. If the option expires in the money, then it is worth $S_T - K$; otherwise, it is worth nothing. But, if interest rates are positive, then:

$$S_T - e^{r(T-t)}K < max(S_T - K, 0). \tag{50}$$

These arguments show that there are no advantages to exercising early if one plans to keep the stock for the remaining lifetime of the option. Another possible action is to exercise the option and sell the stock. In this case one is better off selling the option than exercising it. The price obtained for the option is greater than its intrinsic value $S_t - K$, which one would obtain when exercising the option. This can be seen from equation (48) and the fact that $C_t \geq c_t$. Thus, it is never optimal with early exercise of American calls on a non-dividend-paying stock. Hence, call options on a non-dividend-paying stock have the following property:

$$C_t = c_t. \tag{51}$$

It can be optimal to exercise an American put option on a non-dividend-paying stock early.[9] At any time during its life, a put option should always be exercised

[9] See Hull [1].

early if its sufficiently deep in the money. Because there are some circumstances when it is desirable to exercise an American put option early, it follows that an American put option is always worth more than the corresponding European put option. Hence:

$$P_t > p_t. \tag{52}$$

3.2 Dividends

This section describes how the stock price and the option value reacts to the payment of a dividend. Let the dates of dividend payments be denoted t_i and the amount of dividend paid on day i be D_i. The moment just before the stock goes ex-dividend will be denoted t_i^- and the moment just after is denoted t_i^+. The person who buys the stock on or before t_i^- will also get the rights to the dividend. The person who buys it at t_i^+ or later will not receive the dividend. Naturally, the advantage of receiving the dividend is balanced by a fall in the stock price as it goes ex-dividend. Across a dividend date the stock falls by the amount of the dividend. If it did not, then there would be arbitrage opportunities. Hence:

$$S_{t_i^+} = S_{t_i^-} - D_i. \tag{53}$$

This jump in the stock price will presumably have some effect on the value of an option. Jump conditions determine the value of a dependent variable, an option price, when there is a discontinuous change in one of the independent variables. In the present case, there is a discontinuous change in the asset price due to the payment of a dividend. The jump condition relates the values of the option across the jump. It is assumed that the amount and timing of the dividends during the life of an option can be predicted with certainty. As most exchange-traded stock options last for less than a year, this is not an unreasonable assumption. Since the dividend amount and date are known in advance there is no surprise in the fall in the stock price. The conclusion must be that the option does not change in value across the dividend date, its path is continuous. Let $\Pi(\cdot,\cdot)$ denote the value of the option. Then the continuity of the option value across a dividend date can be written as:

$$\Pi(S, t_i^-) = \Pi(S - D_i, t_i^+). \tag{54}$$

On the next page is a plot of the option value as a function of the stock price just before and just after the dividend date. The stock price before and after the dividend payment are represented by the intercepts between the horizontal line and the two curves. The before curve is identical to the after curve, but shifted by an amount D_i.

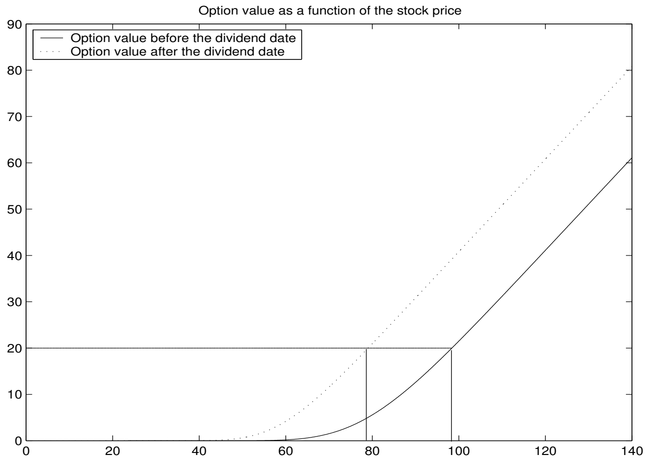

Figure 3: Effect of a dividend on the option price.

Sometimes it is optimal to exercise an American call immediately prior to an ex-dividend date. This is due to the jump in the stock price, making the option less attractive. It is never optimal to exercise a call at other times. For further discussion see Hull [1].

3.3 The Black-Scholes model

The fundamental assumption in the Black-Scholes analysis is that the asset price follows a Geometric Brownian Motion, which implies that asset returns will be normally distributed. That is, the dynamics of the asset price is described by a geometric Brownian motion:

$$dS_t = \mu S_t dt + \sigma S_t dW_t, \tag{55}$$

where μ is the expected yearly return of the asset and σ is the volatility of the asset. Suppose also that there is a risk-free asset, a bond, carrying a constant interest rate of r over $[t, t + \Delta t]$. The value B of this bond satisfies:

$$dB_t = rB_t dt. \tag{56}$$

These processes are the basis of the Black-Scholes theory. Black-Scholes propose that the option's price is determined by only two variables that are allowed to change over the time interval $[t, t + \Delta t]$: time and the underlying asset price. The other factors in the basic model - the volatility, the exercise price, and the risk-free rate - do affect the option's price, but they are not allowed to change.

Now, suppose that f is the price at time t of a derivative contingent on the asset S_t. The Black-Scholes assumptions says that f must be a function of S_t

and t, hence $f = f(t, S_t)$[10]. By the use of arbitrage arguments and Itô calculus, the Black-Scholes partial differential equation can be derived:

$$\frac{\partial f}{\partial t} + rS_t \frac{\partial f}{\partial S_t} + \frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 f}{\partial S_t^2} = rf. \tag{57}$$

In the special case when $f(t, S_t)$ is the price of a European call at time t, equation (57) has the following boundary value: $f(T, S_T) = max(S_T - K, 0)$. Then an analytical form of the option's price can be derived. This is the famous Black-Scholes formula:

$$c_t = S_t N(d_1) - Ke^{-r(T-t)} N(d_2), \tag{58}$$

where $N(\cdot)$ is the standard normal cumulative distribution function, that is $N(d) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{d} e^{-\frac{y^2}{2}} dy$,

and

$$d_1 = \frac{\ln(S_t/K) + (r + \sigma^2/2)(T - t)}{\sigma\sqrt{T - t}}, \tag{59}$$

$$d_2 = \frac{\ln(S_t/K) + (r - \sigma^2/2)(T - t)}{\sigma\sqrt{(T - t)}} = d_1 - \sigma\sqrt{T - t}. \tag{60}$$

The Black-Scholes price of an European put can be derived in a similar way, and is given by:

$$p_t = Ke^{-r(T-t)} N(-d_2) - S_t N(-d_1). \tag{61}$$

3.4 Methods for computing VaR for options

This section gives an overview of some of the methods used for computing VaR for portfolios including options. Most of the standard methods described in the literature only take delta- and gamma[11]-risks into account, and consequently ignores the vega risk. Furthermore, the standard parametric delta-gamma method employs crude approximations, resulting in poor estimates of the risk. This motivates the need for a more sophisticated model.

Since portfolios containing options are highly non-linear in the underlying risk factors, the methods described in section (2) must be adopted to handle this non-linearity. If one assumes that the option value depends quadratically on the underlying asset and linearly on the other risk factors, movements in the

[10] This is true for European options, but not for e.g. path dependent options.

[11] Delta- and gamma-risks are the option's first and second order exposure to changes in the price of the underlying asset.

option value can be found by a Taylor expansion:

$$dc_t = \frac{\partial c_t}{\partial S_t} dS_t + \frac{1}{2}\frac{\partial^2 c_t}{\partial S_t^2}(dS_t)^2 + \frac{\partial c_t}{\partial \sigma} d\sigma + \frac{\partial c_t}{\partial r} dr + \frac{\partial c_t}{\partial \tau} d\tau, \tag{62}$$

where $\tau = T - t$ is the time to expiration for the option.

3.4.1 The delta-gamma method

This section presents a delta-gamma method as described in Jorion [3]. The major risk factor in most cases is the underlying asset price. The delta-gamma method only takes this risk into account, and so ignores the other risk factors. The method assumes that movements in the option value can be approximated by:

$$dc_t \approx \frac{\partial c_t}{\partial S_t} dS_t + \frac{1}{2}\frac{\partial^2 c_t}{\partial S_t^2}(dS_t)^2. \tag{63}$$

Taking the conditional variance, given the information up to time t, of both sides of the quadratic approximation one obtains:

$$V(dc_t) = \Big(\frac{\partial c_t}{\partial S_t}\Big)^2 V(dS_t) + \Big(\frac{1}{2}\frac{\partial^2 c_t}{\partial S_t^2}\Big)^2 V((dS_t)^2) + 2\frac{\partial c_t}{\partial S_t}\frac{1}{2}\frac{\partial^2 c_t}{\partial S_t^2} Cov(dS_t, (dS_t)^2). \tag{64}$$

Assuming that dS_t and $(dS_t)^2$ are approximately jointly normally distributed (which in fact cannot be true) with zero mean, all of the odd moments are zero, and the last term in equation (64) vanishes. Under this assumption, it can be shown that $V((dS_t)^2) = 2(V(dS_t))^2$,[12] and the variance simplifies to:

$$V(dc_t) = \Big(\frac{\partial c_t}{\partial S_t}\Big)^2 V(dS_t) + \frac{1}{2}\Big(\frac{\partial^2 c_t}{\partial S_t^2}(V(dS_t))\Big)^2, \tag{65}$$

where $\Delta = \frac{\partial c_t}{\partial S_t}$ is called the *delta* of the option, and $\Gamma = \frac{\partial^2 c_t}{\partial S_t^2}$ is called the *gamma* of the option.

Since the sum of jointly normally distributed variables is also normally distributed, the distribution of dc_t is known. Let $\sigma = \sqrt{V(dS_t/S_t)}$ be the volatility of returns in the asset price. Since the asset price is known at time t (the present time) one has that $V(dS_t) = S_t^2 V(dS_t/S_t) = \sigma^2 S_t^2$. Then the VaR of the option can be written as:

$$VaR_{\delta\gamma} = \lambda_\alpha \sqrt{\Delta^2 S_t^2 \sigma^2 + \frac{1}{2}(\Gamma S_t^2 \sigma^2)^2}, \tag{66}$$

[12] This can be derived by the use of characteristic functions.

where λ_α is the α-quantile corresponding to the chosen confidence level.

This method can be expanded in order to calculate VaR for a portfolio containing different stocks and options on the stocks.

The assumption of joint Gaussianity of dS_t and its square is, of course, only a very rough approximation. In fact, dS_t and $(dS_t)^2$ cannot both be normal. If dS_t is normal, then $(dS_t)^2$ is chi-squared. Also, the method assumes that the distribution of movements in the option value is normally distributed. Empirical studies has shown that this is far from true.

Monte Carlo simulation

Another possible way to implement a delta-gamma method is to simulate the evolution in the underlying asset price through Monte Carlo simulation. This method can be regarded as a delta-gamma method since the first and second order dependence of the underlying asset price is incorporated in the method. Of course, also higher order dependence will be incorporated in the method.

When implementing the method, one simulates the evolution in the underlying asset price as described in section (2.2.2). For each simulation the new option value is estimated with the Black-Scholes formula or other methods. This procedure is repeated until a stable distribution of the option return is created. The VaR is then taken to be the appropriate percentile from this simulated distribution.

3.4.2 Moment matching methods

To avoid the assumption of Gaussianity one can use moment matching methods when estimating VaR for options.

The following method is proposed in Hull [1]. Start with the second order expansion of the change in option value:

$$dc_t = \Delta dS_t + \frac{1}{2}\Gamma(dS_t)^2. \tag{67}$$

As pointed out in the previous section, the variable dc_t is not normally distributed. Assuming that dS_t is normal with mean zero and variance $S_t^2\sigma^2$, the first three moments of dc_t are:

$$E(dc_t) = \frac{1}{2}S_t^2\Gamma\sigma^2, \tag{68}$$

$$E[(dc_t)^2] = S_t^2\Delta^2\sigma^2 + \frac{3}{4}S_t^4\Gamma^2\sigma^4, \tag{69}$$

$$E[(dc_t)^3] = \frac{9}{2}S_t^4\Delta^2\Gamma\sigma^4 + \frac{15}{8}S_t^6\Gamma^3\sigma^6. \tag{70}$$

These three moments can then be used in conjunction with the Cornish-Fisher expansion in order to estimate the appropriate percentile of the distribution of dc_t. For further detail see Hull [1].

Another method is proposed by *El-Jahel, Perraudin and Sellin* in a working paper from the central bank of Sweden, *Sveriges Riksbank.* The starting point of their analysis is a specification of the underlying stochastic processes that will drive the asset returns in the model. Stochastic volatility is included since volatility risk is often an important factor behind the riskiness of a portfolio of non-linear claims. The technique then makes extensive use of characteristic functions in order to generate conditional moments of the log asset prices and their volatilities. The moments are then used to derive a distribution for the portfolio return and infer the VaR. For details see El-Jahel, Perraudin and Sellin [10]. The advantage of this method compared to Monte Carlo simulation is a smaller computational cost.

4 Implied volatility

One way to use the Black-Scholes and other models is to take parameter values estimated from historical data, substitute them into a formula, and so derive the value of a derivative product. However, there is a difficulty in measuring the value of the volatility of the underlying asset. Despite the assumption of the Black-Scholes theory, it does not appear to be the case that volatility is constant for long periods of time. Remember that the estimated volatilities described in section (2.3) was used to predict the volatility for the next day. In option valuation, one needs to predict the future volatility over the entire lifetime of the option.

A direct measurement of volatility used in option pricing is therefore difficult in practise. However, despite these difficulties it is a fact that option prices are quoted in the market. This suggests that the market 'knows' the volatility of the underlying asset. Take the Black-Scholes formula for a European call, for example, and substitute in the interest rate, the price of the underlying, the exercise price and the time to expiry. All of these are very simple to measure. All that remains is to specify the volatility and the option price follows. Since a call option increases monotonically with volatility (see figure (4)) there is a one-to-one correspondence between the volatility and the option price. Thus one could take the option price quoted in the market and, working backwards, deduce the market's opinion of the value for the volatility over the remaining lifetime of the option. This volatility, derived from the quoted price for a single option, is called the *implied volatility.*

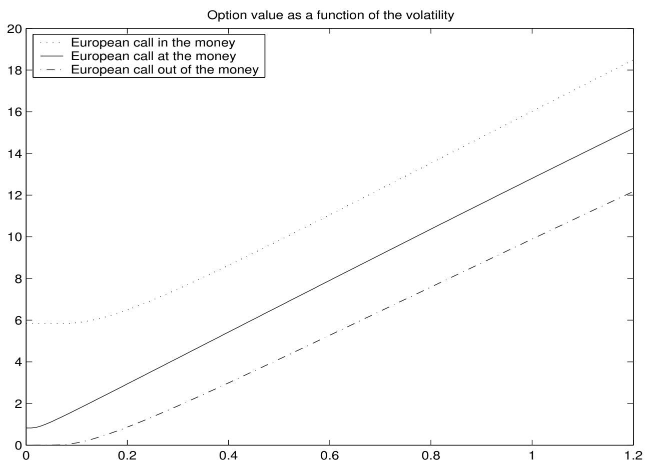

Figure 4: Call price as a function of volatility.

4.1 Computing the implied volatility

The options studied in this thesis are Ericsson stock options traded on the Stockholm Exchange. The option data was collected from the cd: *OM statistik;*

Slutkurser; Optioner & Terminer; 1994-1998. It contains, for each trading day between 1994 and 1998, the closing prices of all options and futures traded on the Stockholm Exchange. The closing price of each option is taken to be the average of the bid-ask spread.

All stock options traded on the Stockholm Exchange are American options. This complicates the computation of the implied volatility. In section (3.1.5) it was shown that if the underlying stock pays dividends then $C_t \geq c_t$, and if there are no dividends then $C_t = c_t$. The VaR-model of this thesis was tested on data from 1996 and 1997. The dividend per share paid by Ericsson was 0.31 SEK in 1996, and 0.44 SEK in 1997. This corresponds to about 1 percent of the share price at the time of the dividend payments. As was seen in section (3.2) this will shift the option-price curve as a function of the share price by an amount equal to the dividend. Also, assume that the option before the dividend is at the money. After the dividend is paid the option will be slightly out of the money. The effect of this is that the option-price curve as a function of volatility will be shifted downwards, which can be seen in figure (4). This will in turn effect the value of the implied volatility. However, this effect is small and will not change the general level of volatility. Furthermore, the implied volatility will be used in a consistent manner when applied to the VaR-method. This will lower the effect of the approximation that $C_t = c_t$, when computing the implied volatility.

The assumption of this thesis is that the dividends have negligible effect on the option price. This means that the price of an American call option is approximated by the Black-Scholes price of an European call option.

4.1.1 Computing the continous interest rate

The deposit rates for different maturities was used to calculate the continuously compounded interest rate curve. The data was collected from Reuters. Interest rates with the following times to maturity was used: 1 month, 2 months, 3 months, 6 months, 9 months and 12 months. For each option, a linear interpolation was done between the two interest rates corresponding to the two numbers of months that are closest to the time to expiration of the considered option. If the time to expiry is greater than 1 year the 12 month interest rate is used. If the time to expiry is less than one month the one-month interest rate is used.

The deposit rates are annually compounded. Let r_d be the interpolated deposit rate. Then the continuously compounded interest rate is computed according to: $r = \ln(1 + r_d)$.

4.1.2 Volatility computation

Let c_{mkt} denote the market price of an European call option, and let c_{bs} denote the theoretical Black-Scholes price of the same option. Then the implied volatility at time t, $\sigma_{imp}(t)$, is computed by solving:

$$c_{mkt} = c_{bs}(S_t, K, T - t, r, \sigma_{imp}(t)). \tag{71}$$

This equation has no analytical solution and has to be solved numerically. In order to find the implied volatility for a specific call option, equation (71) is solved with the Newton Raphson method. The following algorithm is used:

$$\sigma_{i+1} = \sigma_i - \frac{c_{bs}(S_t, K, T-t, r, \sigma_i) - c_{mkt}}{\frac{\partial c_{bs}}{\partial \sigma_i}}, \tag{72}$$

until

$$|c_{mkt} - c_{bs}(\sigma_i)| \leq \epsilon. \tag{73}$$

The term $\frac{\partial c_{bs}}{\partial \sigma_i}$ is called the Black-Scholes *vega* of the option. It can be easily derived from the Black-Scholes formula and is given by:

$$\frac{\partial c_{bs}}{\partial \sigma_i} = S_t n(d_1)\sqrt{T-t}, \tag{74}$$

where $n(\cdot)$ is the standard normal density function. In some cases no solution exists to equation (71). This is mainly due to market imperfections, and will be discussed further in the next section.

5 The volatility surface

5.1 The smile

In the Black-Scholes model volatility is assumed to be constant. This means that when the Black-Scholes implied volatility is plotted against different exercise prices, the theory suggests a horizontal line. In practise, however, Black-Scholes implied volatilities tend to differ across exercise prices. In particular, the variation of implied volatility with exercise price for stock options exhibit a downward sloping, convex pattern.

5.1.1 Reasons for the smile

Since the market crash of October 1987, the volatility smile for most world equity markets has become more pronounced. In the article *Implied Binomial Trees (1994)* Mark Rubinstein presented the following theory explaining this:

One is tempted to hypothesise that the stock market crash of October 1987 changed the way market participants viewed options. Out of the money puts became valued much more highly, eventually leading to the situation where low striking price options had significantly higher implied volatilities than high striking options. The market's pricing of options since the crash seems to indicate an increasing 'crashophobia'.

Another explanation for the smile lies in the empirical fact that stock returns are not log-normally distributed as assumed by the Black-Scholes model. The skewness of the distribution of stock returns is more pronounced and the tails are fatter. This implies that the Black-Scholes implied volatilities are not constant across different exercise prices, but rather shows a 'smiley' pattern.

Yet another explanation for the smile is given in Hull [1]: As a company's equity declines in value, the company's leverage increases. This means that the equity becomes more risky and its volatility increases. As a company's equity increases in value , leverage decreases. The equity then becomes less risky and its volatility decreases. This argument supports the observation that the volatility of equity is a decreasing function of price.

5.2 Creating the smile

Let $M_t = S_t/K$ denote the *moneyness* at time t of an option with exercise price K. Let an *option series* be the collection of options with the same time to expiry, but with different exercise prices. Let $\tau = T - t$ be the time to expiry in days of a particular option series. In the option data studied there are, at each time t and with $\tau \in (1, 150)$, about 10 to 15 options in each option series. However, not all options in the option series are actively traded. The most liquid options are those closest to at-the-money, i.e. with a moneyness close to 1. Options far-out-of the money or deep-in-the money have much smaller trading volumes. For these options there may be a bid-ask spread quoted on the market but the actual trading volume may be null. This will have implications when creating the smile.

A *market maker* will provide liquidity on the market. The market maker in a certain contract guarantees that there always exists a bid-ask spread in that contract. The market makers bid-ask spread will be larger for options far from moneyness one. This implies that individual bid- or ask offers could have a large effect on the bid-ask spread for those options. The bid-ask spread in the option price corresponds to a bid-ask spread in implied volatility. Below is an example plot of the bid-ask spread in implied volatility as a function of moneyness.

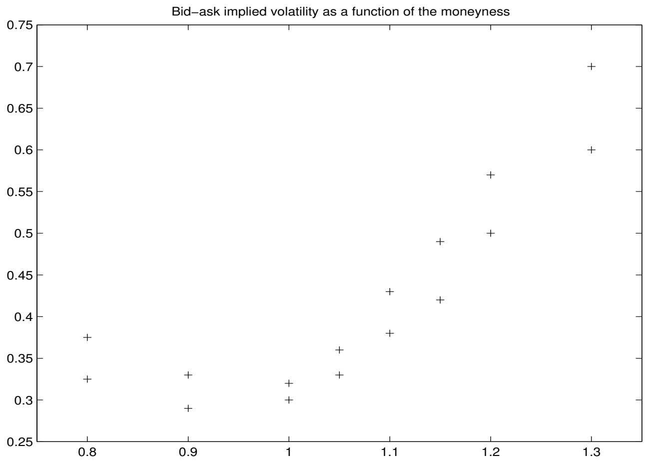

Figure 5: Bid-ask spread of the implied volatility.

The Black-Scholes implied volatility can be assumed to be approximately at the middle of the spread. Since the trading volume is small for options with moneyness far from 1, individual bid- or ask offers on those options could shift the bid-ask average by a large amount.

On the next page is a plot of the implied volatility as a function of the moneyness on February 20^{th} 1996 for the Ericsson option series expiring in April 1996. The time to expiry for the option series is 59 days. The implied volatility is calculated from the average of the bid-ask spread.

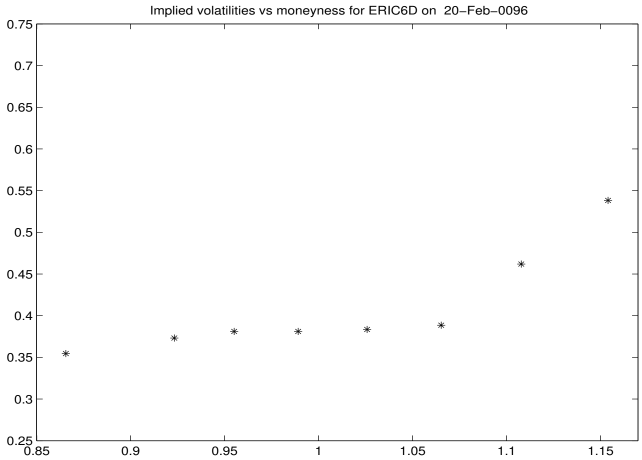

Figure 6: Implied volatility as a function of the moneyness.

The smile is created by fitting a polynomial of degree 2 in the least square sense. There is a large probability that the implied volatility for high moneyness does not correspond to the midpoint of the spread offered by the market maker. Since there is usually a higher concentration of options close to moneyness one than options far from moneyness one, more weight will be given to options which are likely to have stable implied volatilities. Another possibility is to give each option a weight corresponding to the relative trading volume of that option. Below is a plot of the smile curve fitted to the implied volatility points of February 20^{th} 1996.

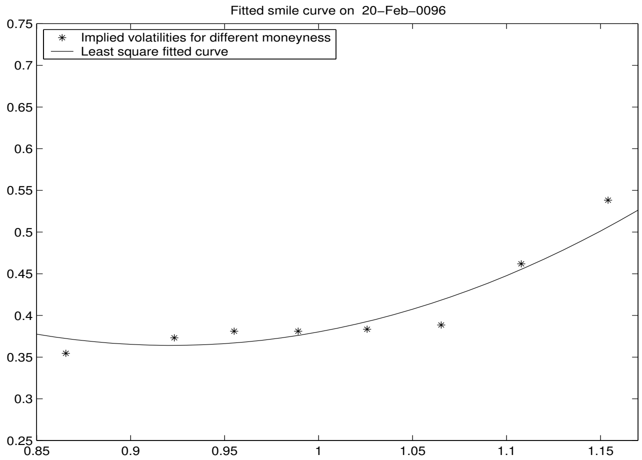

Figure 7: Creating the smile.

Note that the points close to moneyness one are closer to the fitted curve than the points further from moneyness one.

As was pointed out in section (4.1.2), there is not always a solution when computing the Black-Scholes implied volatility. Remember from section (3.1.3) that the theoretical lower bound for a call option is $S_t - Ke^{r(T-t)}$. For options far from moneyness one, single bid or ask offers could shift the bid-ask average by a large amount. This means that a single bid offer could shift the bid-ask average below the lower bound. The result is that no solution exists to the implied volatility equation. Also, when the time to expiry is less than ten days, option prices tend to fluctuate widely resulting in unstable prices. This implies a greater uncertainty in implied volatilities for options close to expiry.

5.3 The term structure

The Black-Scholes implied volatility does not only differ across different exercise prices, but also across different times to maturity. The dependence of implied volatility on the time to maturity is called the term structure of volatility. Implied volatilities for different times to expiry can be seen as the market's view of the long and short term volatility. This can be compared to the term structure curve for interest rates.

The shape of the term structure could be explained by the following: Volatilities tend to be mean-reverting. If the current volatility is historically low, then there is an expectation that it will increase and so the volatility used for an option should be an increasing function of its maturity. If the current volatility is historically high, there is an expectation that it will decrease and so the volatility used for an option should be a decreasing function of its maturity.

Below is a plot of the term structure of the at-the-money implied volatility for different times to maturity on May 15^{th} 1996. At this date Ericsson options with three different times to maturity are traded. The term structure curve is created by a linear interpolation between the three points.

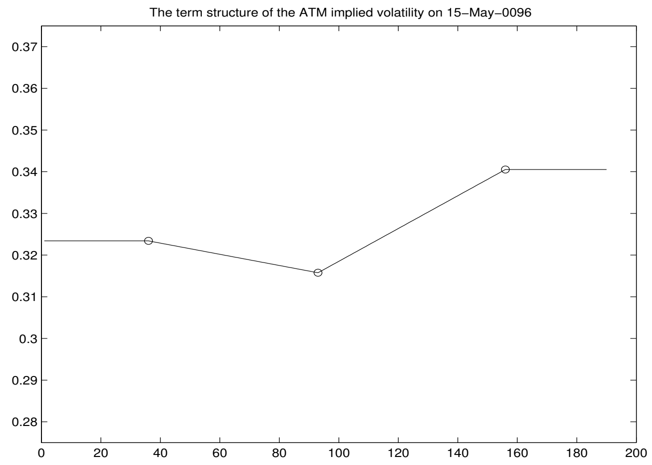

Figure 8: The term structure of implied volatility.

When the implied volatility is plotted against both the exercise price and the time to maturity one gets the implied volatility surface. This will consist of scattered points in a three-dimensional plot. To create the implied volatility surface different techniques can be used to interpolate between the points (see e.g. Wilmott [9]).

5.4 Extensions to Black-Scholes

The Black-Scholes model assumes that the probability distribution of the stock price at any given future time is lognormal. This assumption is not perfect. As was pointed out in section (5.1.1), empirical research has shown that the probability distribution of an equity price has a fatter left tail and thinner right tail than the lognormal distribution. In practise, traders use volatility smiles to allow for non-log-normality.

A number of alternatives to the Black-Scholes model have been suggested. These include models where the volatility of the underlying asset follows a stochastic process; and models where the asset price experiences occasional jumps rather than continuos changes.

As can be seen from figure (8) the implied volatility differs for options with the same underlying asset but with different times to maturity. At the same time the Black-Scholes model assumes a constant volatility. This paradox can be solved by allowing volatility to be time-dependent[13], as Merton did 1973. From the term structure of implied volatilities one can infer a time-dependent instantaneous volatility. Under the equivalent martingale measure[14] Q , the spot process is then governed by the following stochastic differential equation:

$$dS_t = r_t S_t dt + \sigma_t(T) S_t dW_t, \tag{75}$$

where r_t is the continuously compounded interest rate at time t.

However, the dependence of implied volatility on the strike, for a given maturity is trickier. Many researchers have attempted to enrich the Black-Scholes model to compute a theoretical 'smile'. Unfortunately they have to introduce a non-traded source of risk, thus losing the completeness of the model. Completeness is of the highest value; it allows for arbitrage pricing and hedging. In the article *Pricing and Hedging with smiles*, Bruno Dupire[15] builds a model which:

i) is compatible with the observed smiles at all maturities.
ii) keeps the model complete.

This is accomplished by allowing the instantaneous volatility to be a deterministic function of the spot and time. More precisely, given the prices of European calls of all strikes K and maturities T: $c_t(K, T)$, the task is to find a risk neutral *diffusion* process for the spot in the following form:

[13] That is, the volatility is modelled as a deterministic time-dependent function.
[14] See e.g. Hull [1] for definition.
[15] See Dempster & Pliska [13] pages 103-111.

$$dS_t = r_t S_t dt + \sigma_t(K,T) dW_t. \tag{76}$$

This approach extends the Black-Scholes model, taking full power of its diffusion setting, without increasing the dimension of the uncertainty. This model can then be used to elaborate sound pricing for path-dependent and American options. Moreover, it assesses the risk of such options by performing a risk analysis along both strikes and maturities.

The next section will describe a stochastic volatility model, which will incorporate the smile effect into the model.

5.4.1 A stochastic volatility model

This section describes the 'traditional' approach to modelling stochastic volatility that originates from, among others, Hull & White[16] (1987). The asset price is taken to be a diffusion process that is lognormal conditional on the path of a stochastic volatility process $\{\sigma_t\}$:

$$dS_t = \mu S_t dt + \sigma_t S_t dW_t. \tag{77}$$

Thus, S_t is driven by the process σ_t, and is subordinate to it. Let $V_t = \sigma_t^2$ be a continuos solution of:

$$dV_t = \alpha(V_t, t)dt + \beta(V_t, t)dZ_t, \tag{78}$$

where Z_t is a Brownian motion correlated with W_t. The correlation is constant with $Cov(dW_t, dZ_t) = \rho dt$, and $\rho \in$ [-1,1]. The price of an European call option is then modelled as a function of current price and volatility: $c_t = f(S_t, \sigma_t, t)$. This assumes that one can find the present volatility although it is not directly observable. Then, the *no-arbitrage* price of an European call option with strike price K and expiration date T is given by:

$$c_t = E_t^{Q(\lambda^\sigma)}[e^{-r(T-t)} max(S_T - K, 0)], \tag{79}$$

where $Q(\lambda^\sigma)$ is an equivalent martingale measure[17] depending on $\lambda^\sigma = \lambda(S_t, \sigma_t, t)$, the volatility risk premium, or the market price of risk. The notation E_t means the conditional expectation given the past up to time t: $E_t(\cdot) = E(\cdot|\mathcal{F}_t)$, with $\mathcal{F}_t$ the σ-algebra of events generated by $\{(S_u, \sigma_u) : u \le t\}$. The presence of a second non-traded factor that the derivative price depends upon, the stochastic

[16] See Sircar & Papanicolaou [11].

[17] Details of the theory of pricing by equivalent martingale measures can be found in *Arbitrage Theory in Continuos Time by Tomas Björk; Oxford University Press.*

volatility, gives rise to an incomplete market in which the risk inherent in holding an option cannot be completely hedged by trading in just the underlying asset and a risk-free bond. Thus, there must be a function λ^{σ} which represents how the market values the unhedged risk. This too must be estimated from data once the model is put into practise.

The stochastic volatility model is capable of explaining the smile behaviour documented in figure (7). That is, a smile that is skewed. The model of Hull & White (1987), for example, can explain this feature when the asset price and volatility are negatively correlated. The negative correlation is what produces the skew[18], not the stochastic feature itself.

However the model introduces another parameter λ^{σ} which cannot be directly measured, and therefore must be estimated. This parameter is difficult to estimate, and thus making the model subject to bias. Also, for reasonable choices of coefficients, the models tend to produce large smiles for long maturities and shallow smiles for short maturities[19]. In reality, one would observe the opposite.

5.4.2 A jump diffusion model

Merton has suggested a model where the asset price has jumps superimposed upon a geometric Brownian motion. Define μ to be the expected return from the asset, λ to be the rate at which jumps happen, and k to be the average jump size measured as a proportional increase in the asset price. The model then specifies the evolution in the asset price as:

$$dS_t = (\mu - \lambda k)S_t dt + \sigma S_t dW_t + S_t dq_t, \tag{80}$$

where dW_t is a geometric Brownian motion and dq_t is a Poisson process generating the jumps. The processes dW_t and dq_t are assumed to be independent.

The model gives rise to fatter left and right tails than the Black-Scholes model and is consistent with the implied volatilities observed for currency options. This model also gives rise to an incomplete market since adding another stochastic process to the model increases the dimension of uncertainty.

[18]See Dumas, Fleming & Whaley [14].

[19]See Mougin & Voinnesson [15].

6 Value at Risk & vega risk

Vega risk is the exposure of an option position to changes in the implied volatility of the option. This risk can be a large part of the risk of a portfolio containing options.

Vega risk can be thought of as the 'own' price risk of an option position. Since implied volatility can be used as a measure of option prices and has the interpretation as the market's perceived future volatility of returns, option markets can be viewed as markets for the 'commodity' asset price volatility. Exposures to volatility are traded on that market.

Vega risk is analytically easy to nest into a standard risk management framework. The treatment of vega risk in portfolios is, however, often impeded by the lack of availability of data on option implied volatilities. Vega risk is also complicated by the prevalence of volatility smiles and term structures in option markets. One way to deal with the smile effect is to use the stochastic volatility model described in section (5.4.1). One could then simulate the evolution in the option price, using equations (77 & 78), and thus creating a distribution for the change in the option price. However, this method introduces another parameter that must be estimated, the market price of risk. This will induce problems when implementing the model as was discussed in section (5.4.1).

This thesis will present an alternative and more direct way to incorporate the volatility smile and volatility term structure into a Value at Risk model. In the model proportional changes in implied volatility is assumed to follow a random walk which is correlated with the underlying asset price. The main ideas for the model is given in Malz [16]. The object here is to develop a method which is suitable for implementation and then implement it on option data from the Stockholm Exchange. Backtesting is then used to verify the accuracy of the model and to identify possible shortcomings of the model. The model is also tested against alternative methods.

7 The model

This section first presents a simple analytical model which incorporates vega risk. Then a more complicated model which captures the smile effect of implied volatility is developed. This model is then extended to take the term structure of implied volatility into account.

7.1 A simple model

A simple and tractable way of calculating vega risk is to treat implied volatility analogously to other market risk factors to which a portfolio is exposed. This approach assumes that proportional changes in implied volatility follow a random walk with mean zero. The method maintains the Black-Scholes hypothesis that implied volatility is constant across different exercise prices for a given asset and option maturity. This captures exposure to the general level of the implied volatilities of the options in the portfolio, including the correlations of implied volatility with returns on the underlying market factors.

To calculate VaR, one needs the standard deviation of the implied volatility over a sample period, which will be denoted ν_{vol} and expressed as a one-day rate. The term 'vol-of-vol' is often used to describe the volatility of implied volatility. An equally weighted estimation with a sample period of 100 days will be used to estimate the vol-of-vol. Under the assumptions $\log(\sigma_{t+\tau}/\sigma_t)$ is normally distributed with mean zero and a standard deviation $\nu_{vol}\sqrt{\tau}$, where τ is the VaR horizon in days.

Now, assume[20] that the option price depends linearly on the underlying asset price and linearly on the implied volatility. Then movements in the option price can be found by a Taylor expansion:

$$\Delta c_t = \frac{\partial c_t}{\partial S_t} S_t \frac{\Delta S_t}{S_t} + \frac{\partial c_t}{\partial \sigma_t} \sigma_t \frac{\Delta \sigma_t}{\sigma_t}. \tag{81}$$

Taking the conditional variance, given the information up to time t, of both sides of equation (81), and denoting $\delta = \frac{\partial c_t}{\partial S_t}$, and $\kappa = \frac{\partial c_t}{\partial \sigma_t}$ gives:

$$V(\Delta c_t) = \delta^2 S_t^2 V\Big(\frac{\Delta S_t}{S_t}\Big) + \kappa^2 \sigma_t^2 V\Big(\frac{\Delta \sigma_t}{\sigma_t}\Big) + 2\delta\kappa S_t \sigma_t Cov\Big(\frac{\Delta S_t}{S_t}, \frac{\Delta \sigma_t}{\sigma_t}\Big). \tag{82}$$

Now, using the series expansion: $\ln(1+x) = x - \frac{1}{2}x^2 + \frac{1}{3}x^3 - \cdots \quad ; |x| < 1$, and denoting $r_t = \frac{\Delta S_t}{S_t} = \frac{S_{t+1} - S_t}{S_t}$, one has that $\ln(S_{t+1}/S_t) = \ln(1 + r_t) \approx r_t$, when r_t is small. This implies that under the assumption of log-normality $\frac{\Delta S_t}{S_t}$ and $\frac{\Delta \sigma_t}{\sigma_t}$ are approximately jointly normally distributed. The variance in equation (82) can be expressed more compactly by using matrix notation. Let $\boldsymbol{C}$ denote the corresponding covariance matrix of the implied volatility and the asset price, and let $\boldsymbol{x} = [\delta S_t \ \ \kappa\sigma_t]^T$. Then the variance can be expressed as

[20] The assumption is, of course, only a crude approximation, but it makes it possible to construct a simple parametric VaR model that includes vega risk.

$V(\Delta c_t) = \boldsymbol{x^T C x}$, and the VaR at confidence level α is:

$$VaR_{\delta\kappa} = \lambda_\alpha \sqrt{\boldsymbol{x^T C x}}, \tag{83}$$

where λ_α is the corresponding α-quantile. The method could be expanded to work for portfolios containing stocks and options on different stocks. Then the covariance matrix containing the covariances between relative changes in the underlying stocks and relative changes in implied volatilities must be estimated. Since the at-the-money implied volatilities are used, the method will capture the portfolios exposure to the general level of implied volatilities. However, the method will not take the 'smile risk' or the 'term structure risk' into account. The following sections will present a method which addresses these risks.

7.2 Vega risk in the presence of a volatility smile and term structure

In fact, exposure to volatility is a non-linear risk, like the exposure to the underlying asset price. The return distribution to options is therefore skewed rather than symmetrical. This implies that to correctly measure the risks one ought to use the Monte Carlo approach with full repricing: applying correlated normally distributed shocks to the implied volatility and the underlying asset price, re-pricing the option, and computing the chosen quantile from the return distribution.

If attention is restricted to at-the-money options, vega risk is almost exactly linear, as seen in figure (4). Parametric computations of vega VaR are therefore a good approximation for at-the-money options, but may be a poor approximation for in- and out-of-the-money options. Even for at the money options, parametric VaR computations are not advisable, as the exposure to the underlying is highly non-linear, so the total VaR of a portfolio containing options will be inaccurate.

In some option portfolios, vega risk can be much greater than delta risk. This emphasises the need for a risk management model which correctly incorporates vega risk. Some examples of option portfolios which are high in vega risk but low in delta risk are

- a delta-hedged position.
- a long (short) *straddle* , which consists of a long (short) call and a long (short) put with the same exercise price.
- a *strangle* , which consists of a long (short) call and a long (short) put with different exercise prices.

As pointed out earlier portfolios containing options are exposed not only to changes in the general level of implied volatility, but also to

- changes in the curvature and skewness of the smile.
- changes in implied volatility *along* the smile.
- changes in the slope of the term structure of volatility.
- changes along the term structure of volatility as the option 'ages'.

The next section will present a model which takes changes in the general level of implied volatility, and changes in implied volatility along the smile into account.

7.3 The fixed smile approach

In the simpler approach to incorporating the smile, changes in the shape of the volatility smile are ignored and only parallel shifts in the volatility smile, and changes in implied volatility along the smile as the spot rate changes is taken into account. The changes in implied volatility and the underlying stock price are simulated through Monte Carlo simulation.

Implied volatility scenarios are determined not only by the normally-distributed shocks to the general level of implied volatility, but by a displacement along the smile. The displacement along the smile is determined by the shock to the stock price. The change in the stock price will change the moneyness of the option, and thus the implied volatility is changed along the smile.

In the fixed smile approach, it is assumed that the shock to implied volatility shifts the entire volatility smile up or down, but does not change the *shape* of the smile.

Below is a plot of one realisation of the change in implied volatility when the above described effects are taken into account.

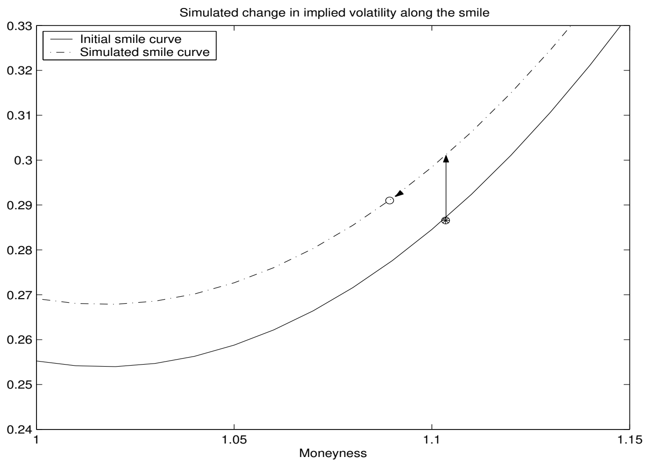

Figure 9: Simulation of the change in implied volatility.

In the example above, the smile curve is shifted upwards by a positive shock to the general level of implied volatility. The underlying stock price is affected by a negative shock and the moneyness of the option declines resulting in a downward shift for the implied volatility along the smile.

The fixed smile approach partially compensates for one of the shortcomings of conventional VaR, the fact that the lognormality assumption is only an approximate description of the actual behaviour of stock prices and implied volatilities. The volatility smile is one residue of this non-normality. The fixed smile approach makes it possible to take the non-normality into account using standard VaR computations.

In the Monte Carlo computation of VaR, correlated shocks to the stock price and the implied volatility are combined. For each realisation the option price is re-valued using the simulated stock price and the simulated implied volatility. This procedure is repeated until a stable distribution of the option return is created. Then the VaR is taken to be the relevant percentile of this distribution. To understand the net effect of the smile on VaR, it is helpful to analyse the effect on the option value of the four possible combinations of positive and negative shocks to these two factors. As a demonstrating example the net effect on a delta-hedged short at-the-money put is analysed.

• **Stock price up, volatility up:** This is the worst combination for the delta-hedged short at-the-money put. The rise in the stock price creates a 'gamma' trading loss (as will be the case for a drop in the stock price), since the gamma of the put is positive and the position is short. The appreciation in the stock price also drives the volatility of the put higher along the smile as the moneyness increases. Thus volatility rises both as a result of the shock to volatility and, in addition, as a result of the change in moneyness.

• **Stock price up, volatility down:** The rise in the stock price creates a gamma trading loss and the appreciation in the stock price also drives the volatility of the put higher along the smile. In this case, the rise in volatility along the smile is offset by its decline due to the volatility shock.

• **Stock price down, volatility up:** The fall in the stock price also creates a gamma trading loss. The volatility falls along the volatility smile as the put's moneyness decreases, offsetting the rise due to the volatility shock.

• **Stock price down, volatility down:** This is the best combination for the delta-hedged short at-the-money put. The volatility falls both along the smile and because of the volatility shock. The result is gain offset only by the gamma trading loss.

To understand the net effect of the smile on VaR, one also needs to consider the correlation between the at-the-money implied volatility and the stock price. Because of the pronounced negative correlation between the stock price and the implied volatility, there is a preponderance of scenarios in which the stock price and volatility move in opposite directions, which leads to offsetting effects of the smile and the volatility shock. As a result, the smile effect on VaR is lowered in the case above, although it does have an effect on the distribution of option returns. The distribution of gains and losses is less disperse, skewed and kurtotic than when the smile is ignored. In general, one can expect a bearish volatility smile, that is, one skewed to lower strikes, when there is a negative correlation between implied volatility and the underlying stock price.

However there are also cases when the smile effect on VaR is significant. Consider a delta-hedged *risk reversal* . A risk reversal consists of a long out-of-the money call and a short out-of-the money put. The vega of the risk reversal is very close to zero at initiation.

Below is a plot of the risk reversal payoff as a function of the stock price at maturity.

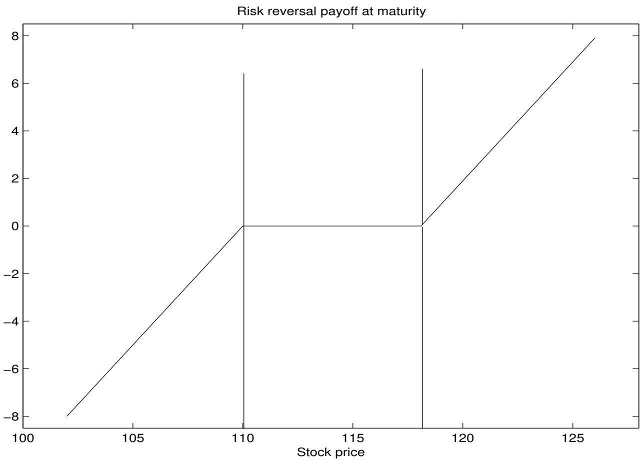

Figure 10: Value of risk reversal at maturity.

The vertical lines represents the strikes of the two options in the portfolio. At the time of initiation the stock price is in the middle between the strikes.

As in the case for any portfolio of options, the delta hedge of a risk reversal must be adjusted as the stock price changes. Similarly, as the stock price fluctuates, it moves closer to the exercise price of one of the options and further from the exercise price of the other, so the risk reversal is no longer vega neutral. In other words, both the delta and the vega must be actively managed as the stock price changes. However, at the time of initiation, in the absence of a volatility smile, the return profile of the risk reversal is symmetrical. In the presence of a smile, the return distribution is no longer symmetrical. In fact, the distribution of returns on the position can be heavily skewed[21].

Because a risk reversal is vega neutral and has low gamma risk, the VaR of a delta-hedged risk reversal is very low in the absence of a smile. In the presence of a smile, however, the VaR increases dramatically[22].

7.4 Random volatility smile

The fixed smile approach can be extended to take the change in the shape of the smile into account. In this approach the smile is not fixed, but rather implied volatilities for different moneyness is allowed to vary in correlated fashion. For example, one can permit the at-the-money implied volatility and the volatilities corresponding to one or two strikes on each side of the at-the-money volatility to vary together. To obtain the smile after the shocks to the implied volatilities of different moneyness, the same interpolation scheme as was described in section (5.2) is used.

[21] See Malz [16].

[22] For more details see Malz [16].

Generally, implied volatilities for options with moneyness greater than one has a markedly higher negative correlation with the stock price than the at-the-money volatility.

Below is a plot of the change in the general level of implied volatility and the curvature of the smile.

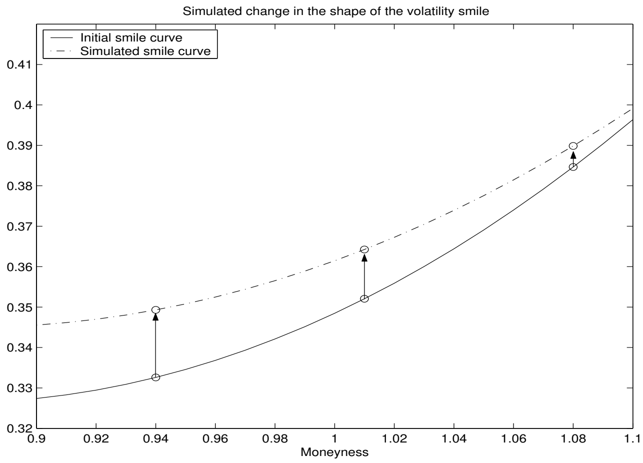

Figure 11: Random volatility smile.

In the example above the general level of implied volatility appreciates, and the curvature of the smile becomes less pronounced.

The shape of the volatility smile generally change slowly over time compared with the level of implied volatility[23]. However, when levels of the underlying price and implied volatility change abruptly, it is often accompanied by a sharp change in the curvature of the smile. These relationships are unlikely to be picked up with the lognormality assumption regarding implied volatilities. Malz [16] suggests that these effects should be incorporated in the model using stress tests[24]. A table of measures of smile curvature accompanying large changes in the underlying stock price and implied volatility can provide scenarios for such stress tests.

7.5 Incorporating the term structure of volatility

One effect of the term structure is that as an option ages, its implied volatility changes along the term structure. Except for very short-term options (less than one month), the VaR horizon of one day is typically short. The shortening of the option maturity by one day is therefore unlikely to change the implied volatility materially along the term structure.

The mechanics of the ageing of options in moving their implied volatilities along the prevailing term structure are analogous to the mechanics of stock price

[23] See Malz [16].
[24] See section (2.4.1).

shocks moving implied volatilities along the smile. The difference is that the effect of ageing is deterministic rather than random.

The term structure is created by a linear interpolation between the at-the-money implied volatilities for options with different maturities as was described in section (5.3). For portfolios containing options with different times to maturity, correlated shocks to the at-the-money implied volatilities corresponding to different maturities are simulated.

Below is a plot of how the term structure of implied volatility for Ericsson stock options changed between the 21 of May 1996 and the 22 of May 1996.

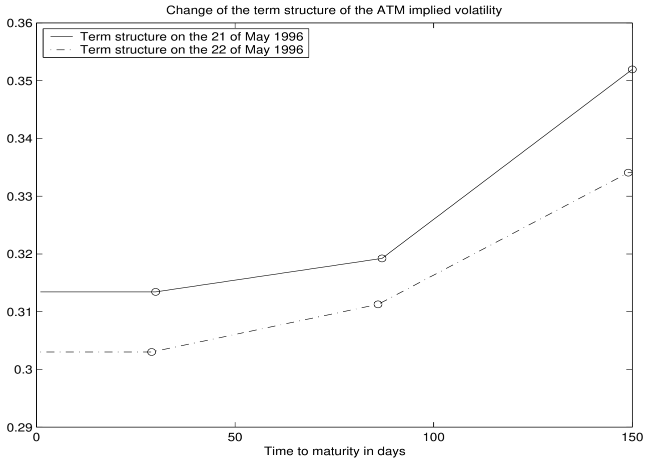

Figure 12: Daily change of the term structure.

Generally, the correlation of implied volatility with the underlying stock price decays rapidly as the time to maturity increases, while longer-term implied volatility is less variable than shorter term[25]. These important differences between shorter- and longer-term volatility are captured in the simulation method described above.

Like as for the smile, the term structure can change shape abruptly, generally at the same time that the stock price, the level of volatility and shape of the volatility smile experience sudden and abrupt changes. Stress testing is therefore a required supplement to standard VaR analysis.

[25] See Malz[16].

8 Implementing the model

The VaR model described in the previous section was implemented in three steps. A crucial part of the implementation is to be able to perform backtesting in order to test the validity of the model. The backtesting procedure, however, presents practical difficulties when implemented, as will be seen in the following sections.

First a method which computes the VaR for an option for every day during its lifetime[26] was developed. The computed VaR values are then compared to the actual outcome of the change in the option value.

The typical lifetime of an option is 100-150 days. As was pointed out in section (2.4.3), one needs to backtest on about 1000 historical trading days in order to reliably test the validity of the model. Also, it is desirable to be able to test how well the method performs with respect to the time to maturity and the moneyness of the option. Assume that the method is to be tested on an option with time to maturity τ^* and moneyness M^*. An option with these properties will not exist on the majority of historical trading days. The second step is therefore to develop a method which overcomes these practical difficulties.

The third step is to develop a method which computes VaR and performs backtesting for a portfolio which consists of options with different moneyness and different times to maturity, and the underlying stock. The options in the portfolio could be call or put options.

All implementations were done in MATLAB.

8.1 Computing VaR of an option during its lifetime

The first step is to create the data which is used in the implementation. Let, as before, an option series be the collection of options with the same maturity date but with different strikes. For each trading day during 1996 and 1997 the following data is created for each existing option series: the time to maturity in days, the moneyness of each option in the option series, the implied volatility of each option in the option series, and the date of the trading day. The data is stored in matrix form, with one matrix for each option series.

For each historical trading day also the term structure of volatility is created. The term structure curve is created by a linear interpolation between the at-the-money volatilities. The at-the-money volatility is taken to be the implied volatility of the option with moneyness closest to one. Options with time to maturity less than 10 days are not used in the creation of the term structure. This is because the price of such options tend to fluctuate widely, making the implied volatility of those options less reliable. The problem of non-existing implied volatilities, described in section (5.2), is typical for options close to maturity.

Let $D = \{d_1, d_2, ...\}$ be the collection of trading dates during the lifetime of a specific option. The method is then performed, for each trading day of the option, in the following steps:

- For day d_i, create that days' smile for the option as a function of the mon-

[26]Except when time to maturity is less than ten days.

eyness. The smile is created with the interpolation scheme described in section (5.2).

• Suppose that the time to maturity for the option on day d_i is τ^*. Then the assumption is that the at-the-money implied volatility with time to maturity τ^* at time t, $\sigma_{ATM,\tau^*}(t)$, and the underlying stock price S_t are jointly lognormal. The covariance matrix, $Cov\Big(\ln\left(S_{t+1}/S_t\right), \ln\left(\sigma_{ATM,\tau^*}(t+1), \sigma_{ATM,\tau^*}(t)\right)\Big)$, is estimated based on the 75 most recent historical data.

• In the next step the change in the stock price and the implied volatility for the option is simulated through Monte Carlo simulation. The correlated shocks to the stock price and the implied volatility is created by Cholesky factorisation of the correlation matrix, as was described in section (2.2.2). The fixed smile approach is used in order to simulate the implied volatility corresponding to the moneyness of the option. For each realisation of the new stock price and implied volatility, σ_{new} and S_{new}, the option is revalued. In the revaluation of the option also the shortening of the time to maturity is taken into account. This means that the method will include delta-, gamma-, vega- and theta[27]-risk.

• For each realisation the relative change of the option price is calculated:

$$r = \frac{\Pi_{new}}{\Pi_{old}} - 1, \tag{84}$$

where $\Pi_{new} = \Pi(\sigma_{new}, S_{new}, \tau^* - 1)$ is the new option price, and $\tau^* - 1$ is converted from days to the corresponding number of years. The simulation procedure is repeated until a stable distribution of the option return is created.

Below is a plot of the simulated distribution of the return for the option ERIC6L170 on the 2 of April 1996. ERIC means that the underlying stock is Ericsson B, 6L means that the expiration date is on the third Friday of December 1996, and 170 means that the strike price is 170 SEK.

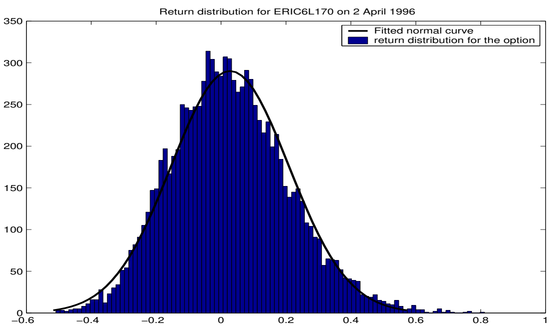

Figure 13: Simulated distribution of the option return.

[27]The sensitiveness of the option value with respect to the time to maturity.

The return distribution in this example is slightly skewed to the left and with a fatter right tail than the fitted normal curve. Monte Carlo simulation with 10000 simulations was used to create the distribution.

• The VaR is taken to be the appropriate percentile from this distribution. The VaR value is then compared to the actual outcome of the change in the option value.

• All the above steps are performed for each trading day of the option's 'life'.

Below is a plot of the estimated VaR-values and the actual outcome for ERIC6L170 during 1996.

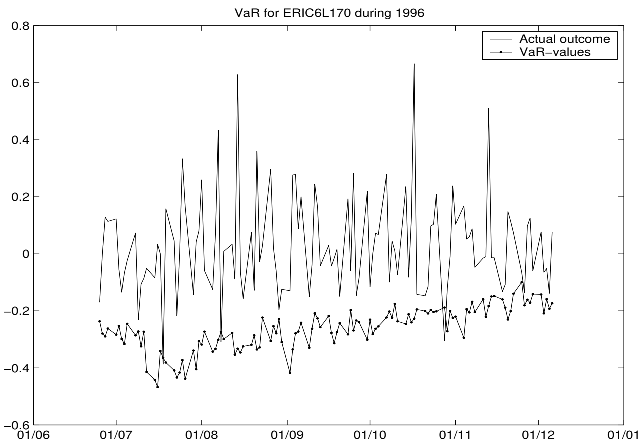

Figure 14: VaR values and the actual outcome.

The number of backtest days for the option above is 119. This is not enough in order to perform a reliable validity test for the method. In the next section a more general method is developed which will make it possible to increase the number of backtest days.

8.2 A general backtesting method

As was pointed out in section (2.4.3) the validity of the model should be tested on about 1000 backtest data in order to be reliable. At the same time the typical lifetime of an option is 100-150 days.

Also, one wants to test how well the method performs for options with specific characteristics. Here the characteristics of the option is determined by the moneyness and the time to maturity of the option. Suppose that the option to be tested has the following characteristics: moneyness M^*, and time to maturity τ^*. For most historical trading days, such an option will not exist.

In order to overcome these difficulties the idea is to indirectly use the implied volatility surface in the backtesting procedure. The method consists of the following steps:

• For each historical trading day: If there exists an option series with time to maturity τ^*, then the corresponding smile curve is used in the simulation process. If such an option series doesn't exist, find the two (or one) option series with time to maturity closest to τ^*.

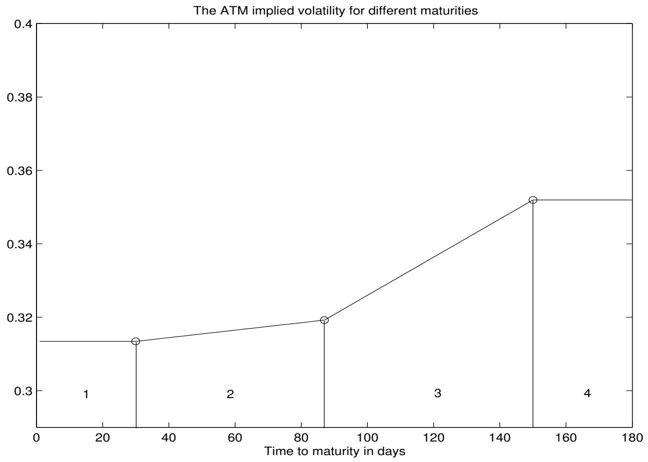

Figure 15: Interpolating the smile.

Consider the example above when option series with three different times to maturity are being traded on the backtest date in question. Let the existing times to maturity be denoted: τ_1, τ_2 and τ_3, with $\tau_1 < \tau_2 < \tau_3$. Let ξ_1, ξ_2 and ξ_3 be the smile curves corresponding to the three times to maturity. Let ξ^* be the unknown smile curve corresponding to τ^*.

• If $\tau^* < \tau_1$ (section 1 in figure 14), ξ^* is approximated by ξ_1.

• If $\tau^* > \tau_3$ (section 4 in figure 14), ξ^* is approximated by ξ_3.

• If $\tau_1 < \tau^* < \tau_2$ (section 2 in figure 14), ξ^* is estimated through a linear interpolation between the pairs of corresponding points on ξ_1 and ξ_2. By corresponding points is meant the two points on the respective smile curves, which have the same moneyness.

• If $\tau_2 < \tau^* < \tau_3$ (section 3 in figure 14), ξ^* is estimated in a similar fashion as above.

The VaR is then calculated in the same way as was described in section (7.3), using the estimated smile curve, ξ^*, in the simulation process.

Below is a plot of the relative return distribution for a call option with moneyness 0.8 and time to maturity 100 days on the 30 of January 1997.

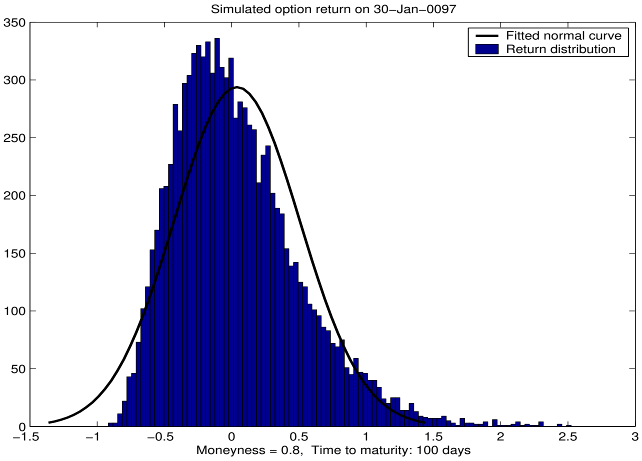

Figure 16: Return distribution 1.

The distribution is heavily skewed to the left and has a large fat tail to the right. It is interesting to compare this distribution with the return distribution of an option with moneyness greater than one. Below is a similar plot for a call option with moneyness 1.1 and time to maturity 100 days.

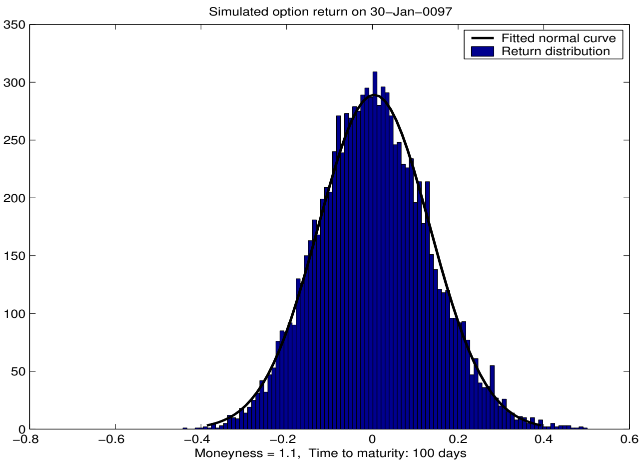

Figure 17: Return distribution 2.

The distribution in this case has a much less pronounced skew to the left and a thinner right tail. This is due to the offsetting effects of the smile which was described in section (7.3). The reason that the offsetting effects are small in the case when the moneyness is 0.8, is that the slope of the smile is much smaller when the moneyness is less than one than when the moneyness is greater than one. This can be seen in figure (7).

8.2.1 Performing the backtest

For each historical trading day the computed VaR value is then to be compared to the actual outcome for the option. However, for most days there is no observable outcome to compare with, that is, no option with moneyness M^*, and time to maturity τ^* exists on that particular day. Again this difficulty is overcome by using the implied volatility surface.

Suppose that one wants to determine the actual outcome for time t. The VaR value for time t has been computed from the interpolated smile curve at time $t-1$, that is $\xi^*(\tau^*, t-1)$. Now the question is how to determine how the actual option value changes from time $t-1$ to t. This is done in the following steps:

- First the interpolated option price at time $t-1$, $\Pi(M^*, \tau^*, t-1)$, is determined.

- The next step is to determine how the option value evolves to time t. In this step it is crucial to realise that the option is likely to change its moneyness. This effect is taken into account by extracting the strike at time $t-1$: $K(t-1) = S(t-1)/M^*$. Remember that the VaR method performs backtest on options with a specific moneyness M^*, which implies that the strike is not constant but varies with the underlying stock price. The moneyness at time t is then taken to be: $M(t) = S(t)/K(t-1)$.

- The new option price is then determined from the interpolated smile curve at time t. The time to maturity will decrease by one day. This means that the interpolated smile curve corresponding to time to maturity $\tau^* - 1$ is used when determining the new option price. That is, the option price at time t is given by:

$$\Pi = \Pi\big(S(t), \sigma_{M(t),\tau^*-1}(t), \tau^* - 1\big). \tag{85}$$

- The actual outcome of the relative change in the option price is then given by:

$$\frac{\Pi\big(S(t), \sigma_{M(t),\tau^*-1}(t), \tau^* - 1\big)}{\Pi\big(S(t-1), \sigma_{M^*,\tau^*}(t-1), \tau^*\big)} - 1. \tag{86}$$

Below is a plot of estimated VaR values and the interpolated actual outcomes for a call option with moneyness one and time to maturity 100 days over the period: 3 of April 1996 to 30 December 1997.

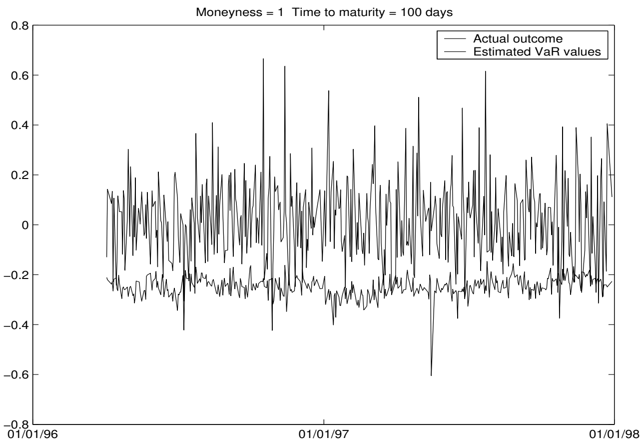

Figure 18: VaR and realised outcomes.

The number of backtest days in the above example is 434.

8.3 VaR for a portfolio of different options

The computation of VaR for a portfolio of options is carried out in a similar way as before. The new wrinkle here is that as the portfolio can include options with different maturities, the correlation between at-the-money implied volatilities of different maturities must be taken into account. The options in the portfolio could be both calls and puts.

The computation of VaR is done in the following steps:

• First the options in the portfolio are divided into different buckets, where each bucket represents options with the same time to maturity.

• For each maturity date in the portfolio, the corresponding smile curve is estimated in the same way as was described in the previous section.

• The correlation between the at-the-money implied volatilities for different maturities must be taken into account in the simulation process. The correlation matrix will also include the correlation between the underlying stock price and the different volatilities.

• The VaR is then calculated as before, using the estimated smile curves in the simulation.

The backtest procedure is also carried out in the same way as in the previous section, where the estimated smile curves are used to find the unobservable changes in the option values. The actual outcome of the relative option return is computed separately for each option in the portfolio. Let $r_i(t) = \frac{\Delta\Pi_i(t)}{\Pi_i(t)}$, where $\Delta\Pi_i(t) = \Pi_i(t) - \Pi_i(t-1)$, be the relative return for option i at time t and let w_i be the relative weight in the portfolio of option i. Then the actual outcome for the relative portfolio return at time t is:

$$R(t) = \sum_i w_i r_i(t). \tag{87}$$

This relation also holds when the underlying stock is included in the portfolio. Then one of the r_i's will be given by the relative return for the stock.

Below is a plot of the VaR values of the individual options of an example portfolio and the corresponding actual outcomes. Also the total VaR of the portfolio and the corresponding actual outcome is plotted.

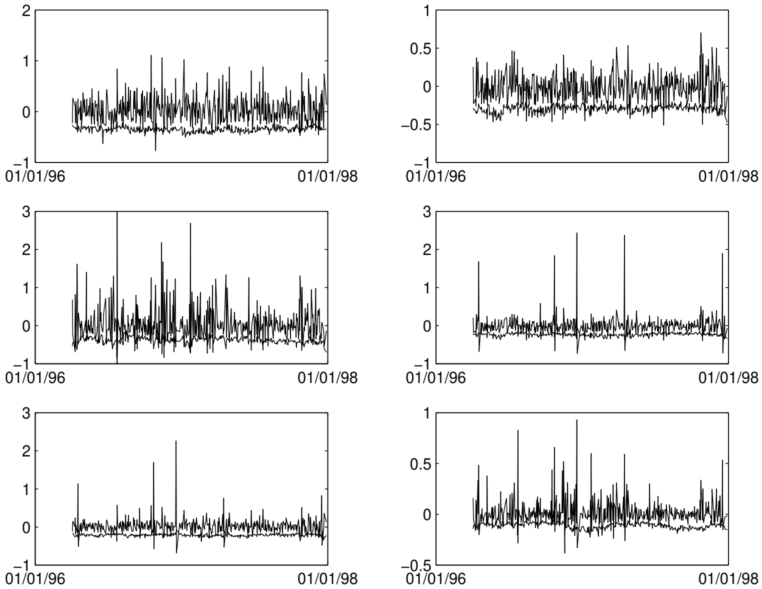

Figure 19: Portfolio VaR

The portfolio consists of, from top left to bottom right: a call with $M = 0.95$ and $\tau = 60$, a put with $M = 0.95$ and $\tau = 60$, a put with $M = 1.1$ and $\tau = 60$, a put with $M = 0.95$ and $\tau = 120$, a put with $M = 1.05$ and $\tau = 120$. The bottom right figure is a plot of the total outcome for the portfolio.

9 Results and conclusions

This section presents the results of the backtesting procedure for options with different characteristics, that is options with different moneyness and different times to maturity. This will make it possible to decide how well the method performs for options with different characteristics.

The VaR model developed in this thesis is also compared to alternative, less complicated, models. This is done in order to decide if the higher computational cost is justified by a better accuracy in the estimated VaR values.

The alternative models which was used in the comparison were:

• A model which takes delta-, gamma-[28], and theta-risk into account when computing VaR. This was done by simulating the evolution in the stock price using Monte Carlo simulation, but assuming that the implied volatility remains constant until the next trading day.

• A parametric VaR model which was presented in section (7.1). That is, VaR is computed according to:

$$VaR_{\delta\kappa} = \lambda_\alpha \sqrt{\boldsymbol{x}^{\boldsymbol{T}} \boldsymbol{C} \boldsymbol{x}}, \tag{88}$$

where $\boldsymbol{x} = [\delta S_t \quad \kappa\sigma_t]^T$, and $\boldsymbol{C}$ is the covariance matrix of the stock price and the implied volatility; κ and δ are calculated from the at-the-money call option.

For each backtest date 1000 simulations was used in the Monte Carlo simulation. The total number of backtest days were 434. Below is a table of the failure rate of the three models for **call** options with different moneyness and times to expiration. By failure rate is meant the rate by which the actual outcome exceeds the estimated VaR value. Since the VaR models used in the implementation have confidence level 95%, ideally the failure rate should be 0.05.

	Moneyness	0.8	0.9	1	1.1	1.2
Time to maturity 50 days	I	0.0642	0.0562	0.0481	0.0535	0.0963
	II	0.0802	0.0535	0.0455	0.0990	0.123
	III	0.0775	0.0508	0.0348	0.0668	0.0909
100 days	I	0.0531	0.0554	0.0462	0.0439	0.0554
	II	0.0624	0.0508	0.0323	0.0670	0.0924
	III	0.0393	0.0393	0.0346	0.0462	0.0716
150 days	I	0.0442	0.0482	0.0442	0.0643	0.0683
	II	0.0843	0.0602	0.0643	0.0763	0.0964
	III	0.0402	0.0361	0.0361	0.05232	0.0683

Table 4.

In the table, I represents the VaR model developed in this thesis, II represents the Monte Carlo model with constant volatility, and III represents the parametric model.

[28] As well as all higher order risks in the underlying asset.

In section (2.4.3) the likelihood ratio test proposed by Kupiec was presented. The test provides a method for validation of the confidence level of a VaR method. For a sample size of 510 backtest days and a 95% confidence level for the VaR method, the non-rejection region for the number of failures is :
(16 , 36). This corresponds to the following non-rejection region for the failure rate: (0.0314 , 0.0706).

The null hypothesis to be tested is: $H_0 : p = 0.05$, where p is the underlying 'true' probability of failure of the VaR method. The main problem with the test is that it only concludes whether H_0 is rejected or not. It doesn't tell if H_0 is accepted. Let $\alpha = P(\text{reject } H_0 \text{ when } H_0 \text{ true})$ and let $\beta = P(\text{do not reject } H_0 \text{ when } H_0 \text{ false})$. The non-rejection regions given above corresponds to the likelihood ratio test with $\alpha = 0.05$. This means that there is a small probability of rejecting the null hypothesis when it is true, but at the same time the probability of not rejecting H_0 when it is false is quite large. By increasing α in the test the probability of falsely not rejecting H_0 when it is false becomes smaller.

The following criterias will be used in evaluating the three VaR models:

• If the failure rate is outside the non-rejection region given above, then the method is considered to be inaccurate.

• The closer to 0.05 the failure rate is, the more accurate the VaR model is considered to be[29]. If the failure rate is very close to 0.05, then the method will pass the likelihood ratio test even for large α, and the probability that H_0 is true is large.

Summary of results from table 4:

The VaR model developed in this thesis (model I) is only rejected when the moneyness is 1.2 and the time to maturity is 50 days. The failure rate of model I is generally closer to 0.05 than both of the other models.

Model II which assumes that the implied volatility remains constant is rejected in 7 cases. The model is rejected when the moneyness is greater than one ($M = 1.1$ and $M = 1.2$), and when the moneyness is much smaller than one ($M = 0.8$). Generally the model underestimated VaR and consequently has a failure rate which is greater than 0.05 in most cases.

The parametric model (model III) is rejected in 3 cases, two times with moneyness 1.2 and one time with moneyness 0.8. The model generally underestimates the VaR of call options which are out-of-the-money and at-the-money. On the other hand, the model generally underestimates VaR of call options which are in-the-money.

All of the models had problems with accurate measures of VaR when the moneyness is 1.2, that is when the call option is deep in-the-money. This could be caused by the approximation that the price of the American call is approximated by the Black-Scholes price of an European call. In some cases it could be optimal with early exercise for the call. For model I another plausible reason for the inaccurate measure of VaR when $M = 1.2$ is the inaccuracy in the

[29]This reasoning may be a bit awkward since there is an uncertainty in the estimated VaR-values, and thus it is difficult to distinguish between the models if H_0 is not rejected.

corresponding implied volatility. When the moneyness is 1.2 there is a greater uncertainty in the calculated implied volatilities which in turn will effect the accuracy of the VaR measure.

Model II, which assumes that the implied volatility is constant until the next trading day, generally performs worst of the three models. This implies that the vega risk in most of the cases is significant and must be captured to successfully measure the VaR of an option position.

The model developed in this thesis generally performs well, and outperforms the two other models in most cases, which justifies the higher computational cost. The number of backtest data used in testing the models was 434. Ideally one would want to test the models on about 1000 backtest data, in order to robustly test the accuracy of the models. The backtest method described in section (8.2) makes it possible to test the models on 1000 backtest data. For this master thesis, however, time was too scarce in order to perform this extensive backtest procedure.

So far nothing has been said about backtesting results for put options with different characteristics. Ideally one would want to perform the same backtesting procedure for puts as was done for the calls. However, the assumption that the value of an American option can be approximated by the European Black-Scholes price of the option does not work well for puts. If the approximation is supposed to be valid, then the backtesting procedure would be straightforward. Remember the put-call parity for European options in section (3.1.4). Because put-call parity holds for the Black-Scholes model, the following is valid:

$$p_{bs} + S_t = c_{bs} + Ke^{-r(T-t)}. \tag{89}$$

Put-call parity also holds for market prices, and thus:

$$p_{mkt} + S_t = c_{mkt} + Ke^{-r(T-t)}. \tag{90}$$

Subtracting these two equations gives:

$$p_{bs} - p_{mkt} = c_{bs} - c_{mkt}. \tag{91}$$

From equation (91) the conclusion can be drawn that the implied volatility of the call option is equal to the implied volatility of the put. However, for American options put-call parity does not hold. Instead the following relation is valid for American options[30]:

$$C - P \leq S - Ke^{-r(T-t)}. \tag{92}$$

Hence, the implied volatilities used when computing VaR for the call options may be a poor approximation for the puts. Therefore backtesting was not performed for puts with different characteristics, and must be left for future investigation.

[30]When the stock price pays no dividend.

Bibliography

[1] John C. Hull (2000), *Options, futures & other derivatives, 4^{th} edition,* Prentice-Hall International.

[2] Kevin Dowd (1998), *Beyond Value at Risk,* John Wiley & sons.

[3] Philippe Jorion (1997), *Value at Risk: The new benchmark for controlling derivatives risk,* McGraw-Hill

[4] Bernard W. Lindgren (1968), *Statistical theory, 3^{d} edition,* Macmillan Publishing.

[5] Jongwoo Kim and Jorge Mina (2000), *Riskgrades Technical Document,* RiskMetrics Group.

[6] Paul H. Kupiec (1995), *Techniques for verifying the accuracy of risk measurement models,* Journal of Derivatives, 3, pages 73-84.

[7] Gunnar Blom (1989), *Sannolikhetsteori och statistikteori med tillämpningar,* Studentlitteratur.

[8] Johan Liwing and Katrin Näsgårde (1999), *Value at Risk - En kvantitativ undersökning av olika beräkningsmetoder för den svenska marknaden med avseende på enkla linjära tillgångar,* Bachelor thesis at Stockholm University.

[9] Paul Wilmott (2000), *Paul Wilmott on quantitative finance* , John Wiley & sons.

[10] Lina El-Jahel, William Perraudin and Peter Sellin (1997), *Value at Risk for Derivatives,* Sveriges Riksbank Working Paper Series.

[11] K. Ronnie Sircar and George C. Papanicolaou (1997), *Stochastic Volatility, Smile & Asymptotics,* Department of Mathematics, Stanford University.

[12] Kim Sundkvist (2001), *Essays on option pricing with smiles and non-constant volatility,* Doctoral Thesis at Swedish School of Economics and Business Administration, Vaasa, Finland.

[13] Dempster and Pliska (1997), *Mathematics of Derivative Securities,* Cambridge University Press.

[14] Dumas, Fleming and Whaley (1998), *Implied Volatility Functions: Empirical Tests,* The Journal of Finance, 6, pages 2059-2106.

[15] Mougin and Voinesson (2000), *On Volatility Surfaces for American Equity Options,* Master Thesis at the Royal Institute of Technology, Stockholm, Sweden.

[16] Allan M. Malz (2000), *Vega risk and the smile* Working Paper from the RiskMetrics Group

[17] Pierre-Yves Moix (2001), *The measurement of market risk,* Springer-Verlag

[18] Jonas Lundberg (2001), *Option pricing with a 'smile' - derivation of probability distribution via relative entropy minimization,* Master Thesis at the Royal Institute of Technology, Stockholm, Sweden.

[19] Jose A. Lopez (1998), *Reglatory evaluation of Value at Risk models,* Working paper from the Federal Reserve Bank of New York.

[20] Barone-Adesi & Giannopoulos (2000), *Non-parametric VaR techniques: myths and realities,* Working paper.