Empirical Issues in Value at Risk Estimation: Time varying Volatility, Fat Tails and Parameter Uncertainty*

Dennis Bams[1]
and
Jacco L. Wielhouwer[2]

2 November 1999

Abstract

This paper describes alternative approaches to estimate the Value at Risk (VaR) of a position. Four methods are compared: the unconditional case, the model with time varying drift (modeled as an AR(1) process), the model with time varying drift and time varying volatility (modeled as a GARCH(1,1) process) with error terms that are normally distributed, and the model with time varying drift and time varying volatility with error terms that are Student-t distributed. Two issues are important. First, different specifications for mean, variance and fat tailness lead to different point estimates for the associated distribution function and hence to other VaR measures. Second, uncertainty in parameter estimates implies that the VaR also is uncertain. The model with error terms that are t-distributed is the preferred model, since: (1) the time varying volatility incorporates that recent volatility is a better predictor for the future, (2) the time varying volatility makes it possible to use a longer time series which implies less uncertain VaR estimates and (3) the fat tail of the distribution is taken care of by the t-distributed error terms. An important contribution of the paper lies in the fact that we explicitly take account for parameter uncertainty and propose ways to deal with it.

Key Words: Value at Risk, time varying volatility, fat tails

Correspondence: Jacco Wielhouwer, ING, Market Risk Management, Location code : HC 01.04, PO BOX 1800, 1000 BV Amsterdam, The Netherlands, phone: + 31 20 563 4187, fax: +31 20 563 5063, email: Jacco.Wielhouwer@mail.ing.nl

* We thank Johan van der Ende for useful comments on an earlier draft. The usual disclaimer applies.
[1] Market Risk Management ING Group and Maastricht University
[2] Market Risk Management ING Group and Center, Tilburg University

1. Introduction

An important and popular risk management tool for financial institutions nowadays is Value-at-Risk (VaR) (see e.g. Jorion (1997) for a thorough overview). It is not only an internal management tool to check whether traders are within their limits, but it is also a by the Basle Committee prescribed risk measure for the (international) supervisor. Formally VaR is the maximum value that is lost over a certain period within a given confidence interval. The confidence level reflects 'extreme market conditions' with a probability of, for example, 2.5% or 1%.

The measure is based upon the probability distribution of the underlying return series. It deals with the maximum loss under extreme market conditions, which occur in the left tail of the distribution function of the future market value. More specifically, given the market value for which the probability mass of even lower market values is, for example, one percent, the change of the current market value to this extreme market value is the VaR.

Crucial for the determination of the extreme future market value, and hence for the VaR, is the distribution function of the return on market value. Usually, as allowed by the Basle Committee, a normal or lognormal distribution is assumed for the market return. Nowadays, an important issue is the desirability of a distribution function that has fat tails (see e.g. Embrechts et al. 1997), since it is claimed that the normal distribution underestimates the probability in the tail and hence the VaR. An appropriate alternative distribution would be the Student-t distribution, since it allows for fatter tails in the distribution than the normal distribution.

A second important issue is the way the parameters of the underlying distribution functions are determined. Usually these parameters are estimated using historical data. These parameter estimates are plugged into the distribution function and are assumed to be given fixed figures. A problem with using point estimates is that in fact these parameters still incorporate uncertainty, which is usually reflected in the form of standard errors. Uncertainty in the parameter estimates implies uncertainty about the underlying distribution function and hence about the VaR.

In the following we determine the VaR and provide both expected value for the VaR and a standard deviation which reflects the uncertainty about the actual VaR. The two methods that are mostly used for the estimation of the VaR are Risk-Metrics of JP Morgan (JP Morgan, 1994) and historical simulation (see e.g. Jorion 1997 pp. 193-195). The Risk-Metrics method describes the possibility of calculating an exponentially weighted covariance matrix. If this is not done, using few historical data leads to very

uncertain VaR-results (reflected by large confidence intervals around the point estimate). But using a long time series leads to an implausible point estimate since it neglects that recent data are more representative for tomorrow than older data. The problem with the exponential weighted volatility is that it must be kept constant in the prediction of volatilities into the future, or otherwise decreases to zero. Besides, a second parameter (the decay factor λ) has to be estimated. The fact that only recent data have an important influence on λ implies that this parameter estimate will be very uncertain.

Since parameter uncertainty is an important issue, we propose a method to arrive at a univariate time series of market values for the portfolio under consideration, in stead of modeling the multivariate behavior of different economic variables (like interest rates and exchange rates). In the latter case many parameters have to be estimated leading to higher uncertainty of the ultimate VaR. Moreover only the linear co-movements between the different economic variables are incorporated. By evaluating the current portfolio composition against the yield curve and exchange rates at historical times, we arrive at a time series of portfolio market values that also incorporates the effects and co-movements of economic variables that influence the value of the portfolio. The advantage of this method is that we arrive at a univariate time series, which requires less parameters to be modeled and hence leads to increased efficiency.[3]

Four models are compared. First, the simple model in which both the expected value and the variance of the underlying return are assumed to be constant over time. Usually, the parameters of this model are estimated over a short historical period, since this best reflects the current situation in the market. Return observation of periods further away are less likely to represent the current situation. Second, we consider an AR(1) model for the market value return, to take account of the fact that there is some persistence in the level of returns. Third, we consider a GARCH(1,1)[4] model that takes account of the fact that volatility is time-varying, where error terms are assumed to be normally distributed. Fourth, we consider a GARCH(1,1) model where we assume that the error term is Student-t distributed.

We show that the latter model is preferred over the other three models. First, the variation in volatility takes account for the most recent volatility regime. Second, the Student-t distribution takes account for distributions with more probability mass in the left tail. This provides a more accurate picture of the VaR since it better models the probability of extreme returns. Third, it leads to lower variability in the VaR, i.e. it is more precisely estimated than methods in which an unconditional distribution is

[3] The fact that the series is univariate also simplifies the econometric treatment of more complex models.
[4] For a review of (G)ARCH models we refer to Bollerslev, Engle and Nelson (1994).

estimated on a recent sub-sample of the data. In general a lot of historical observations are required to arrive at precise parameter estimates, and hence at precise VaRs.

In the next section we set up the econometric framework for a VaR analysis, section 3 describes the return series that is of interest and in section 4 we provide the estimation results for the four models that are under consideration. We perform both in-sample and out-of-sample analyses. Also, the expected VaRs are given for the alternative models together with a standard deviation to reflect uncertainty in the given figure. We finish with some concluding remarks in section 5.

2. The econometric framework for the VaR analysis

Exposure to downside risk can be summarized in a single number by an estimate of the VaR. Jorion (1996) formally defines VaR as "the worst expected loss over a target horizon within a given confidence interval". The extreme future investment value is set equal to the quantile that is associated with a probability mass of, for example, $\alpha = 0.01$ in the left tail of its probability distribution. We define W_t as the investment value at time t and $R_{t,t+N}$ as the return over the interval $[t, t+N]$. Daily returns are defined with continuously compounding, and are denoted by

$$r_t = \ln\left[\frac{W_t}{W_{t-1}}\right] \tag{1}$$

and $R_{t,t+N}$ simply follows as the sum of the daily returns in the interval $[t, t+N]$. Crucial for the determination of the value at risk is a probability distribution for the future investment value. This distribution is fully determined by the distributions of future periodical returns. In the following we deal with normally distributed returns and with returns that follow a Student-t distribution. Normal distributed returns are fully determined by the first and second moment of the returns. The expected value of $R_{t,t+N}$ is defined by $\mu_{t,t+N}$ and its volatility by $\sigma^2_{t,t+N}$. Under the assumption of normally distributed returns the VaR is easy calculated as

$$VaR = W_t\{1 - \exp[\mu_{t,t+N} + sign(W_t)\sigma_{t,t+N}\ z_\alpha]\} \tag{2}$$

where $sign(W_t)$ is +1 if W_t is positive and −1 if W_t is negative, and z_α denotes the critical value of the normal distribution with a probability mass in the left tail of α. The only unknown parameters in equation (2) are the expected return and its volatility over the interval $[t, t+N]$. In case of Student-t distributed periodical returns, the resulting

future distribution of the position of the bank is determined by means of simulation. The critical one percent value follows immediately.

We propose and compare four methods to arrive at a distribution function for future returns: (1) the unconditional model, that only uses unconditional estimates for both expected value and volatility of normally distributed returns, (2) the AR(1) model, that assumes an auto regressive relation for daily normally distributed returns with constant volatility, expanding the first model with an auto-regressive part, (3) the N-GARCH(1,1) model for daily returns with normally distributed error terms and a GARCH(1,1) structure on the volatility, and (4) the t-GARCH(1,1) model for daily returns with Student-t distributed error terms and a GARCH structure on the volatility. All models are nested into the fourth case, which we will describe in detail.

In a GARCH(1,1) model the volatility of the return is given by a time varying process, where volatility at time t, σ_t^2, depends upon the volatility of the day before and upon the shock in the return in the previous period. For the return itself we assume that it can be represented by an AR(1) model. The complete GARCH(1,1) model with t-distributed error terms reads

$$r_t = \mu + \rho r_{t-1} + \varepsilon_t \tag{3}$$

$$\varepsilon_t \sim t\left(0, \sigma_t^2, \theta\right) \tag{4}$$

$$\sigma_t^2 = \gamma_0 + \gamma_1 \sigma_{t-1}^2 + \gamma_2 \varepsilon_{t-1}^2 \tag{5}$$

The N-GARCH(1,1) model follows after we let the degrees of freedom, θ, go to infinity $(\theta \to \infty)$. The AR(1) model with constant volatility follows after $\theta \to \infty$, $\gamma_1 = 0$ and $\gamma_2 = 0$. The Unconditional model follows after the restrictions $\theta \to \infty, \gamma_1 = 0$, $\gamma_2 = 0$ and $\rho = 0$ are imposed. Intuitively, the model without restrictions is most flexible since it allows for time varying forecasts of return levels, time varying volatility and for fat tails in the return distribution.

All models are estimated with maximum likelihood, where the loglikelihood is given by

$$\begin{aligned} \ln L \;=\; & T \ln \Gamma\left(\frac{\theta+1}{2}\right) - T \ln \Gamma\left(\frac{\theta}{2}\right) - \tfrac{1}{2} T \ln\left[\pi(\theta-2)\right] + \\ & - \tfrac{1}{2} \sum_{t=1}^{T} \ln\left(\sigma_t^2\right) - \left(\frac{\theta+1}{2}\right) \sum_{t=1}^{T} \ln\left[1 + \frac{\varepsilon_t^2}{(\theta-2)\sigma_t^2}\right] \end{aligned} \tag{7}$$

in the most general case where the error terms follow a Student-t distribution. In the special case of normally distributed error terms, the loglikelihood function reduces to

$$\ln L = -\tfrac{1}{2}T\ln(2\pi) - \tfrac{1}{2}\sum_{t=1}^{T}\ln(\sigma_t^2) - \tfrac{1}{2}\sum_{t=1}^{T}\left[\frac{\varepsilon_t}{\sigma_t}\right]^2 \tag{8}$$

where T denotes the number of historical observations. Parameter estimates follow after application of the Newton-Raphson algorithm. The covariance matrix that is associated with the parameter estimates follows from the negative of the inverse of the information matrix, I, where the (i,j)-th element is given by

$$I_{ij} = \left[\frac{\partial^2 \ln L}{\partial p_i \partial p_j}\right] \tag{9}$$

and p is the vector that consists of the parameters that have to be estimated.

We compare the *in-sample* performance of the four models by testing the restrictions that are imposed by the unconditional model, the AR(1) model, the N-GARCH(1,1) model and the t-GARCH(1,1) model. Formally, this results in a likelihood ratio test, where the critical value, X, is defined by twice the difference in loglikelihood, which is distributed as a chi-squared distribution where the degree of freedom is determined as the difference in number of parameters.

$$X = 2[\ln L_U - \ln L_R] \sim \chi^2(n_U - n_R) \tag{6}$$

where $\ln L_U$ denotes the value of the loglikelihood function in the optimum for the unconstrained model, and $\ln L_R$ denotes the value of the loglikelihood in the optimum for the constrained model. With n_U and n_R we denote the number of parameters in the unconstrained and the constrained model, respectively.

Also, we setup an *out-of-sample* test in which one part of the data is used to estimate a model for the return distribution. This model is applied to calculate the VaR for 20 periods ahead (which resembles one month forecast). The VaR then is compared with the realized change in the position of the bank over the same period. This procedure is repeated for many different periods. In the end we calculate the percentage of the cases in which the actual change in position of the bank exceeded the reported VaR. Since we calculate the VaR under 99 percent certainty, the percentage of times that the VaR is exceeded should not be greater than 1 percent if the model adequately describes the behavior of the daily returns.

In our view, an important observation is that there exists statistical uncertainty with respect to the VaR, because the parameter estimates for the underlying return distribution are uncertain. This uncertainty is most easily reflected by a confidence interval for the reported VaR. The covariance matrix of the parameter estimates reflects the uncertainty in the parameters. We take a close look at the special cases of the unconditional model and the AR(1) model, since in these cases the parameters estimates

and standard errors can easily be expressed analytically and therefore visualize the econometric properties of the parameters. In case of the unconditional model the parameter estimates leads to the following easy closed forms for the expected value

$$\hat{\mu} = \frac{1}{T}\sum_{t=1}^{T} r_t \tag{10}$$

and for the volatility

$$\hat{\sigma}_t^2 = \frac{1}{T}\sum_{t=1}^{T}(r_t - \hat{\mu})^2 \tag{11}$$

To reflect the uncertainty in these parameter estimates, the covariance matrix of the parameter estimates is usually calculated. In the unconditional model the covariance matrix of the parameter estimates is given by

$$V(\mu,\sigma^2) = \begin{bmatrix} \frac{\hat{\sigma}^2}{T} & 0 \\ 0 & \frac{2\hat{\sigma}^4}{T-1} \end{bmatrix} \tag{12}$$

In the special case of the AR(1) model, a vector Y of dependent variables is defined as

$$Y = \begin{bmatrix} r_2 \\ \vdots \\ r_T \end{bmatrix} \tag{13}$$

and a matrix X of explanatory variables is given by

$$X = \begin{bmatrix} 1 & r_1 \\ \vdots & \vdots \\ 1 & r_{T-1} \end{bmatrix} \tag{14}$$

The estimator for μ and ρ follow as the OLS estimator

$$\begin{pmatrix} \hat{\mu} \\ \hat{\rho} \end{pmatrix} = (X'X)^{-1}X'Y \tag{15}$$

and for σ^2 the estimator is given by

$$\hat{\sigma}^2 = \frac{1}{T}\left[Y - X\begin{pmatrix} \hat{\mu} \\ \hat{\rho} \end{pmatrix}\right]'\left[Y - X\begin{pmatrix} \hat{\mu} \\ \hat{\rho} \end{pmatrix}\right] \tag{16}$$

The associated covariance matrix of the parameter estimates is given by

$$V(\mu,\rho,\sigma^2)=\begin{bmatrix}\hat{\sigma}^2(X'X)^{-1} & 0\\ 0 & \dfrac{2\hat{\sigma}^4}{T-2}\end{bmatrix} \tag{17}$$

An important observation from the covariance matrix of the parameter estimates in equations (12) and (17) is that the variances of μ and σ^2 decrease with the number of observations T, and hence that a large sample of observations is required to arrive at efficient estimates. This is in conflict with approaches in which only a recent sub-sample of the data is incorporated to arrive at parameter estimates that are conform the most recent market developments. This motivates models that allow for time varying expected returns and time varying volatility.

3. Data

The data over the observed period represent the market value changes of a specific book of ING-Bank.[5] The book consists of interest rate risks in both NLG and DM. Furthermore, on average there is a long position in the book, i.e. the downside risk was related to increases in interest rates. The position ultimo April 1998 is valued against historical yield curves. The curves used are the daily swap-curves from the period of January 24, 1991 until April 30, 1998. This results in market value changes on a day to day basis. These data not only represent the changes in the interest rates and the slope of the yield curve, but also the changes in the spread between German and Dutch interest rates, since the part of the (investment) portfolio that is in German instruments is valued against the German interest rates. So, we consider historical changes of the factors that drive the market value changes (so German and Dutch yield curves). Furthermore, we focus on the effect these historical changes can have on the current portfolio. An important advantage is that the final market value changes are a one-dimensional time series. This simplifies estimation of GARCH models and increase reliability of the results.

Table 1 provides summary statistics on the daily returns. The empirical distribution is a bit skewed and furthermore it exhibits fat tails, given a kurtosis that exceeds 3, sometimes leading to extreme positive or extreme negative daily returns. In figure 1 the daily returns are represented as a time series. From the figure we observe that there are periods in which returns are more volatile than in other periods, which

[5] Because of confidentiality, we have scaled the data so that the results do not represent the actual market value changes. The implications for the VaR are the same as with the real data.

motivates a model that accounts for time varying volatility. In figure 2 the daily returns are presented in a histogram, which shows that there are quite some extreme observations in the left tail. This motivates a distribution function that has fatter tails than a normal distribution.

Table 1. Summary Statistics

Mean	-0.03%
Standard Deviation	0.82%
Minimum	-4.31%
Maximum	3.71%
Skewness	0.24
Kurtosis	5.32

Notes The table presents summary statistics of the returns sample, which consists of daily observations for the period June 25, 1991 until April 30, 1998.

Figure 1. Time Series Returns

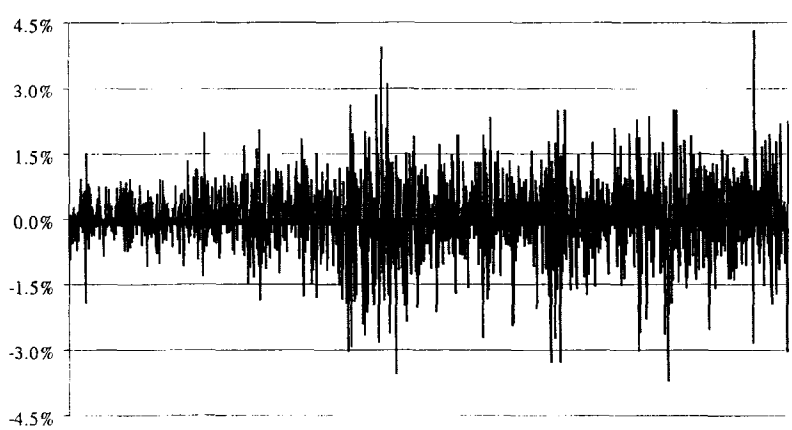

Notes: This figure presents the daily returns for the period from June 25, 1991 until April 30 1998 in the form of a time series.

Figure 2. Histogram Returns

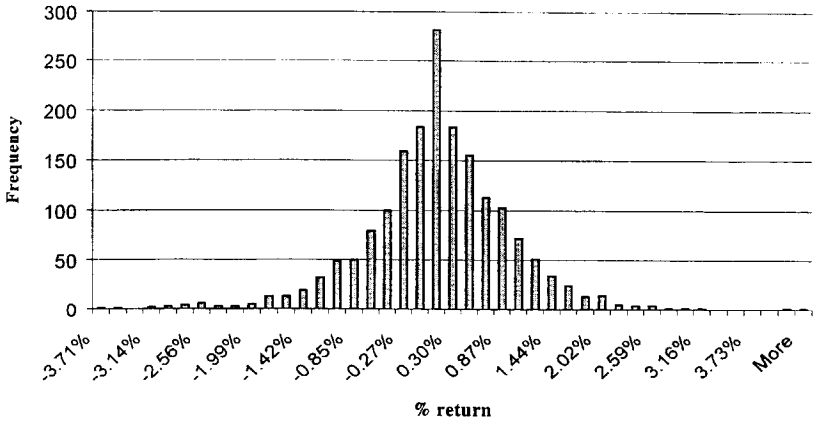

Notes: This figure presents the daily returns for the period from June 25, 1991 until April 30 1998 in the form of a histogram.

4. Results

In this section we describe the estimation results for the four models that are under consideration. We start with an in-sample description and formally test the alternative specifications against each other. Second, we present the implied VaR under the alternative specifications and provide a measure of reliability. Then we investigate the relation between the number of historical observations that are used to estimate the models and the associated reported VaR. We end with an out-of-sample test of the models, in which we compare the reported VaR with the realized changes in the position of the bank for different sub periods.

All models are estimated using the maximum likelihood principle. In case of the t-GARCH model the likelihood function is given in equation (7). The remaining three models are special cases of the likelihood function in equation (8). In table 2 the parameter estimates for the four models are given. The average value in the unconditional model is not significantly different from zero, the variance is estimated precisely when a long data sample is incorporated. The parameter results of the AR(1) model show that there is a significant relation between two consecutive daily returns, given the estimate for ρ. The parameter estimates for the N-GARCH(1,1) model show the same relation for consecutive returns as in the AR(1) model. Furthermore the representation of the

conditional volatility shows that volatility is time varying. Volatility is persistent on a day-to-day basis, given the high value for γ_1 and the impact of an unexpected shock in the previous period is significant, given the parameter estimate and the associated standard error for γ_2. The t-GARCH(1,1) model shows parameter results that are similar to the N-GARCH(1,1) model, but now we also find a significant estimate for the degree of freedom. The estimate for θ implies that the distribution of the returns have more probability mass in the tail than in case of a normal distribution.

A formal way to compare the in-sample performance of the three models is to perform a likelihood ratio test. To compare the Unconditional model with the AR(1) model, we test the restriction that $\rho = 0$ holds. Two times the difference in loglikelihood is compared with the critical value of a chi-squared distribution with one degree of freedom. Since 6 is greater than $\chi^2_{0.95}(1) = 3.84$, we reject the restriction $(\rho = 0)$ and prefer the AR(1) model over the Unconditional model. The data hence suggest that a time varying drift is preferred over a model that assumes that conditional means are constant.

Comparing the N-GARCH(1,1) model with the AR(1) model boils down to comparing two times the difference in loglikelihood with the critical value of a chi-squared distribution with four degrees of freedom. Since 332 is greater than $\chi^2_{0.95}(4) = 9.49$, we reject the restriction of constant volatility and prefer the N-GARCH(1,1) model that allows volatility to vary over time. In figure 2 the in-sample volatility is depicted for the GARCH(1,1) model. Testing the t-GARCH(1,1) model against the N-GARCH(1,1) model results in a preference for the former model, since 84 is greater than the critical value of $\chi^2_{0.95}(1) = 3.84$. From the results in table 2 we conclude that with respect to the in-sample behavior of the four models the data suggest that the t-GARCH(1,1) model is the preferred model.

Table 2. Estimation Results

Parameters	t-GARCH(1,1)	N-GARCH(1,1)	AR(1)	Unconditional
μ $(\times 10^{+3})$	-0.440	-0.400	-0.320	-0.300
	(0.140)	(0.150)	(0.190)	(0.190)
ρ	-0.083	-0.059	-0.060	'0'
	(0.023)	(0.023)	(0.024)	(-)
γ_0 $(\times 10^{+5})$	0.025	0.033	6.57	6.59
	(0.011)	(0.014)	(0.220)	(0.221)
γ_1	0.950	0.946	'0'	'0'
	(0.011)	(0.010)	(-)	(-)
γ_2	0.046	0.049	'0'	'0'
	(0.010)	(0.009)	(-)	(-)
θ	6.309	'∞'	'∞'	'∞'
	(0.766)	(-)	(-)	(-)
σ_1 $(\times 10^{+2})$	0.390	0.408	0.810	0.812
	(0.106)	(0.096)	(0.014)	(0.014)
LnL	6238	6196	6035	6032

Notes This table gives the parameter estimates for the t-distribution model and special cases of this model. Standard errors are given within parentheses. With σ_1 we denote the standard deviation at the first time period and lnL denotes the value of the loglikelihood function. The models are estimated using daily observation for the period from June 25, 1991 until April 30, 1998.

Table 3. Residuals

	t-GARCH(1,1)	N-GARCH(1,1)	AR(1)	Unconditional
Skewness	0.178	0.169	0.254	0.225
Kurtosis	4.54	4.51	5.38	5.30

Notes This table reports the third and fourth moment of the standardized residuals for the t-GARCH(1,1) model, the N-GARCH(1,1) model, the AR(1) model and the Unconditional model. The associated standard errors are 0.058 for skewness and 0.12 for kurtosis.

In table 3 we present values for both the skewness and kurtosis of the standardized in-sample residuals. The standardized residual at time t is obtained by dividing the residual by the associated standard deviation at time t. All models result in residuals that show extreme kurtosis.

The four models are used to generate forecasts for future daily returns for up to 20 periods ahead (which corresponds with a period of one month). The 20 forecasted daily returns are added to arrive at the expected total return over the one month period. This enables us to calculate the VaR. Because the parameter values that serve as input for the forecast model are uncertain, we take account for this uncertainty by repeatedly drawing values for the parameters from a normal distribution with mean the parameter estimates and covariance the covariance matrix of the parameter estimates. This results in an expected VaR together with a standard error, that reflects the uncertainty in the VaR. In table 4 these values are given for all models. The expected VaR as given by the Unconditional model and the AR(1) model are almost the same. The N-GARCH(1,1) model assumes a significant higher value at risk, taking account for both the negative drift in expected returns and the higher volatility in the most recent period than on average in the historical sample (see also figure 3). This is in contrast with the previous two models in which it is assumed that volatility is constant. Finally, the t-GARCH(1,1) model reports the highest VaR, taking account for fat tails, which directly explains why the VaR is higher in this case. The more realistic representation of the t-GARCH(1,1) model comes with a price, since the associated standard error of the reported VaR is higher than in the other models. There is more uncertainty in the true value for θ since it is determined by extreme observations only.

Table 4. Value-at-Risk Results

t-GARCH(1,1)	N-GARCH(1,1)	AR(1)	Unconditional
309	245	230	231
(25)	(9)	(12)	(13)

Notes This table presents the expected VaR together with the associated standard error (within parentheses) for the position one month ahead. The underlying models have been estimated on daily returns for the period from June 25, 1991 until April 30, 1998.

From the previous remarks it becomes apparent that a desirable property of the underlying return distribution is the fact that it takes account for the most recent market circumstances. An interesting approach would be to incorporate only recent market

observations to take account for only the most recent market developments. In table 5 we show the implications for the expected VaR and the associated standard error for different number of observations in the estimation routine. The results are striking: the reported expected VaR differs a lot with the number of observations that is used to estimate the parameters of the underlying return distribution. This uncertainty is also reflected in the associated standard error, which is extremely high when only a limited number of observations is used. Only for a large data sample the reported VaR is more reliable. This suggests that the unconditional and the AR(1) model are less appropriate, since they would assume the same expected value and volatility over the entire sample.

In table 6 we consider the out-of-sample performance of the different models. For different sub-samples the model is estimated over one period and applied to the following period to calculate a VaR. The reported VaR is then compared with the actual observed change in value over the out-of-sample period. We report the percentage of cases in which the actual change in the position of the bank exceeds the reported VaR. Because there is uncertainty in the VaR we also consider the upper and lower bounds of the 95 percent interval around the VaR. See figures 4, 5, 6 and 7 for the reported VaR, the associated 95 percent confidence interval for the VaR and the actual change in the position of the bank. A model that adequately takes account of the behavior of the return distribution would allow only a one percent violation of the VaR. The results show that the most restricted model performs worse and that the percentage of violations decreases as the model becomes more general. The t-GARCH(1,1) model still shows a violation of 2.29 percent for the reported VaR. If however we take account for the uncertainty in the VaR by considering the 95 percent upperbound on the reported VaR, we find that in that case the reported VaR is violated a little less than the allowed 1 percent.

Table 5. Relation betweenVaR and Number of Observations

#observations	Expected VaR	SE VaR
20	468	166
40	308	130
60	275	89
80	263	72
100	266	70
150	256	53
200	244	44
400	253	32
600	266	26
800	248	21
1000	256	19
1500	255	15
1787	231	13

Notes This table presents the expected VaR together with the associated standard error (within parentheses) for the position one month ahead for the unconditional model. The underlying model has been estimated on daily returns for different sub-samples in the period from June 25, 1991 until April 30, 1998. The first column gives the number of most recent observations that is used to estimate the parameters.

Table 6. Out-of-Sample Violations

	t-GARCH(1,1)	N-GARCH(1,1)	AR(1)	Unconditional
VaR (low)	5.42	6.82	10.92	11.06
VaR	2.29	5.42	9.18	9.18
VaR (high)	0.97	4.52	7.02	7.02

Notes. The table reports the actual percentage of violations of the predicted VaR under the alternative model specifications. An appropriate model should result in a violation of 1 percent. The uncertainty in the reported VaR is reflected with the 95 percent lower- and upperbounds.

5. Concluding Remarks

In this paper we have compared four alternative models to calculate the VaR for the value of the bank. Crucial for this calculation is the underlying return distribution, since it reflects the probability of extreme returns. A number of issues are important.

First, the underlying probability distribution should be able to reflect the behavior of extreme returns. Hence the tail of the distribution should be well-modeled. We have proposed to adopt a Student-t distribution since it allows for fatter tails than a normal distribution.

Second, the VaR is based on historical return observations. Recent market circumstances should be most informative on the implied future return distribution. This is either accomplished with a time varying return distribution based on a large historical data sample or with an unconditional distribution that is based on most recent observations only. The first method is preferred since a lot of observations are required to arrive at reliable estimates.

Third, since the parameters of the underlying return distributions are unknown, they have to be estimated. The associated standard errors of the parameter estimates reflect uncertainty in the underlying distribution, which implies that the reported VaR also incorporates uncertainty. We have reported the VaR together with a standard error. The empirical implications are that a relative long time series is required in order to arrive at relatively reliable VaR (i.e. with low associated standard errors). The preferred model is the t-GARCH(1,1) model since it allows for time varying drift and volatility to take account of most recent market circumstances, and for fat tails.

Fourth, in order to model the VaR of a set of positions, we valued the current position as if we had it in historical subsequent periods. This results in a univariate time series that increases reliability in the VaR estimation. The alternative, for example the estimation of a covariance matrix, so a multivariate time series, requires more parameters to be estimated, which decreases the reliability of the estimation results.

Fifth, the out-of-sample tests indicate that the t-GARCH(1,1) model indeed reports a VaR that shows the least number of violations. Taking into account the uncertainty of the reported VaR, we in fact cannot reject that the t-GARCH(1,1) model adequately describes the VaR.

Of course, there is always room for improvement. First, we find that the degrees of freedom for the Student-t distribution is an important parameter, because it takes account for fat tails in the return distribution. At the same time the uncertainty of this parameter is higher than for the other parameters. This increases the uncertainty in the

reported VaR. Second, from figures 6 and 7 we observe that the GARCH(1,1) specification indeed takers account for more volatile periods, but only after a first shocks, for which it does not take account. Only the degrees of freedom take account for that behavior. Finally, this model has been tested on a book that has positions in NLG and DM. When such a model is applied to for example emerging markets, distribution functions with even fatter tails might be required. In these cases we might think of, for example, a generalized Pareto distribution for the error terms. We leave these issues for further research. Notice however that also in these case it is worth reducing the input series to a univariate series and to take account for parameter uncertainty as in the proposed out-of-sample test.

References

BOLLERSLEV, T. , R.F. ENGLE and D. NELSON, (1994), ARCH Models , in R.F. ENGLE and D. MCFADDEN, eds, *Handbook of Econometrics*, *volume IV*, North Holland.

EMBRECHTS, P., C. KLUPPELBERG, and T. MIKOSCH, (1997): *Modelling Extreme Events*, Springer Verlag, Berlin.

JORION P., (1996), 'Risk: Measuring the risk in Value at Risk', *Financial Analyst Journal*, November/December.

JORION P., (1997), '*Value at Risk: The New Benchmark for Controlling Market Risk*', McGraw-Hill

JP MORGAN BANK, (1995), *RiskMetrics Technical Document*, New York: JP Morgan Bank

Figure 3. GARCH(1,1) Volatility

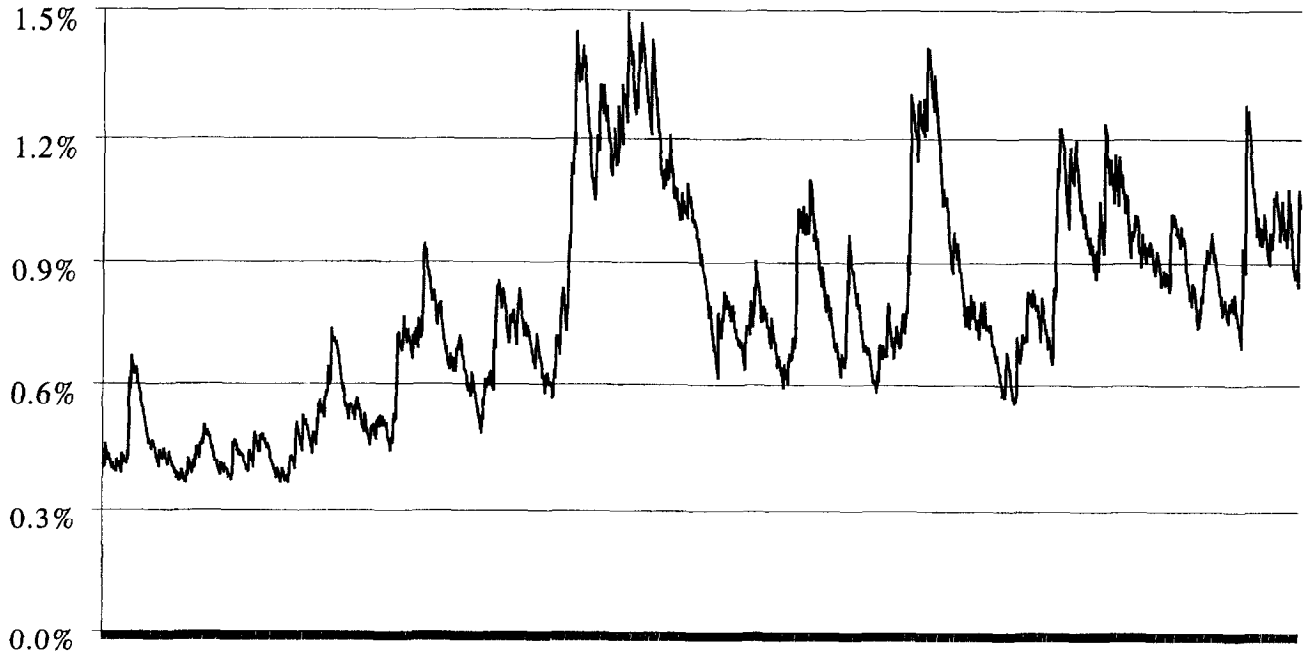

Notes: This figure presents the in-sample volatlity as implied by the N-GARCH(1,1) The underlying model has been estimated on daily returns for the period from June 1991 until April 30, 1998.

Figure 4. Unconditional model

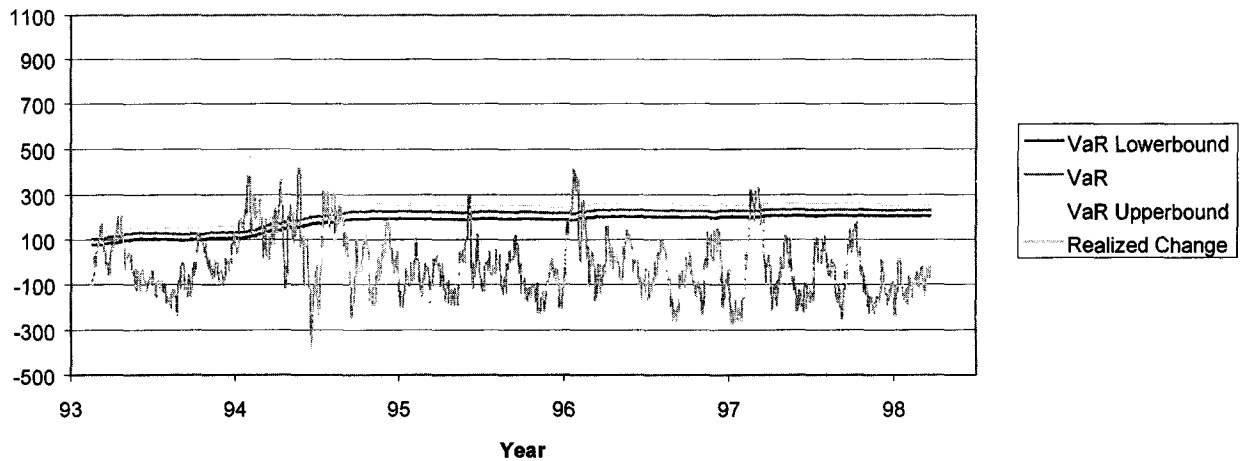

Notes. The figure displays the forecasted VaR values together with 95 percent upper- and lowerbounds. Also the realized change in position is displayed. The entire period preceding the time of forecast is used to estimate the model.

Figure 5. AR(1) model

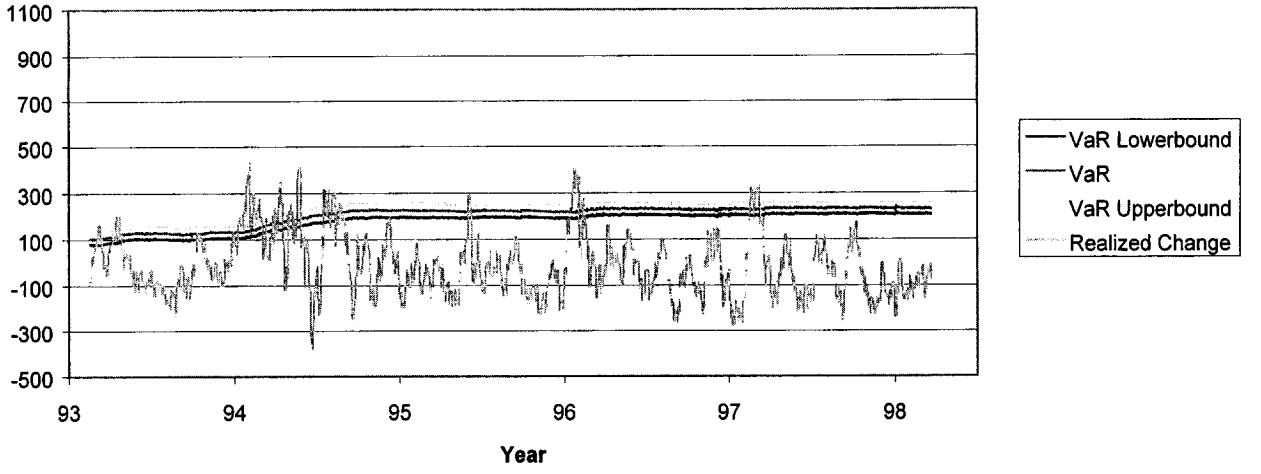

Notes. The figure displays the forecasted VaR values together with 95 percent upper- and lowerbounds. Also the realized change in position is displayed. The entire period preceding the time of forecast is used to estimate the model.

Figure 6. Normal Distribution GARCH(1,1)

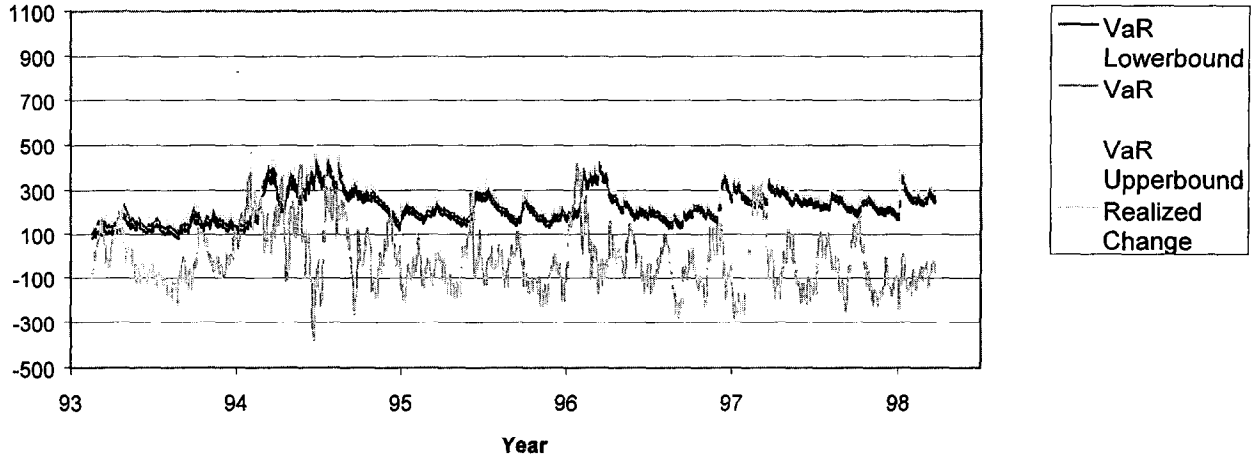

Notes. The figure displays the forecasted VaR values together with 95 percent upper- and lowerbounds. Also the realized change in position is displayed. The entire period preceding the time of forecast is used to estimate the model.

Figure 7. Student-t Distribution GARCH(1,1)

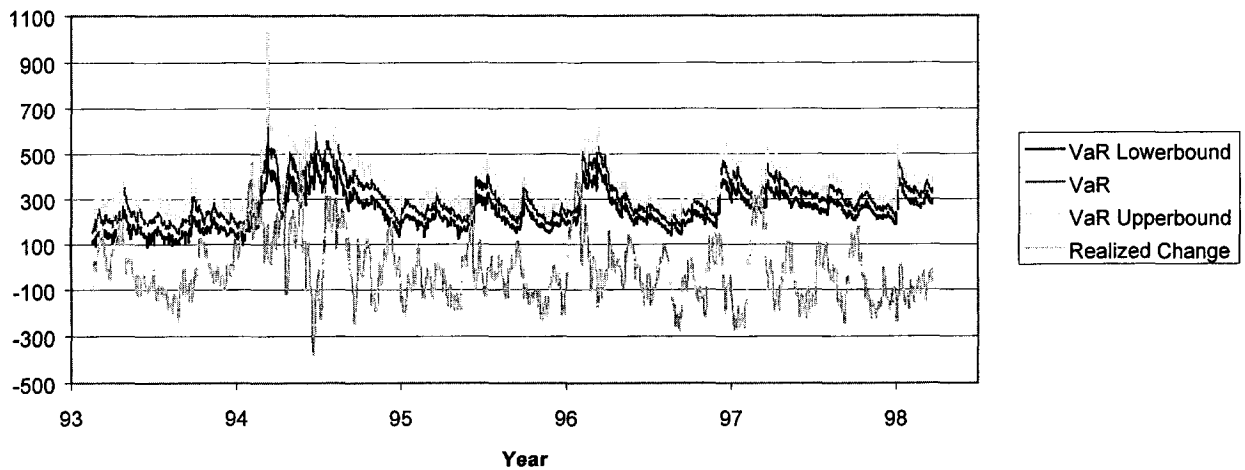

Notes. The figure displays the forecasted Var values together with 95 percent upper- and lowerbounds. Also the realized change in position is displayed. The entire period preceding the time of forecast is used to estimate the model.