

EMPIRICAL ISSUES IN VALUE-AT-RISK*

BY

DENNIS BAMS¹ AND JACCO L. WIELHOUWER²

ABSTRACT

For the purpose of Value-at-Risk (VaR) analysis, a model for the return distribution is important because it describes the potential behavior of a financial security in the future. What is primarily, is the behavior in the tail of the distribution since VaR analysis deals with extreme market situations. We analyze the extension of the normal distribution function to allow for fatter tails and for time-varying volatility. Equally important to the distribution function are the associated parameter values. We argue that parameter uncertainty leads to uncertainty in the reported VaR estimates. There is a trade-off between more complex tail-behavior and this uncertainty. The “best estimate”-VaR should be adjusted to take account of the uncertainty in the VaR. Finally, we consider the VaR forecast for a portfolio of securities. We propose a method to treat the modeling in a univariate, rather than a multivariate, framework. Such a choice allows us to reduce parameter uncertainty and to model directly the relevant variable.

KEY WORDS

Value-at-Risk, Parameter Uncertainty, Time-varying volatility, Fat tails

1. INTRODUCTION

An important and popular risk-management tool for financial institutions nowadays is Value-at-Risk (VaR). (See, e.g., Jorion (2000) for a thorough overview). VaR analysis is not only a useful 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. VaR is an estimate for the maximum value that can be lost over a certain period

* We thank Johan van der Ende, two anonymous referees and participants of the AFIR-2000 colloquium for useful comments on an earlier draft. The usual disclaimer applies.

¹ ING Re Amsterdam and LIFE, Maastricht University.

² ING Bank Amsterdam and CentER, Tilburg University.

within a given confidence interval. The confidence level reflects extreme market conditions with a certain probability of, for example, 2.5% or 1%.

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. As allowed by the Basle Committee, a normal or lognormal distribution has usually been assumed for the market return. Recently, alternative distributions have been proposed that focus more on the tail behavior of the returns. See, for example, Embrechts, Kluppelberg and Mikosch (1997), McNeil and Frey (1999) and Lucas and Klaassen (1998) for a discussion. A normal distribution supposedly underestimates the probability in the tail and hence the VaR result. Popular alternatives in the financial literature include GARCH-type models which allow for time-varying volatility, and the Student-t distribution, which allows for more probability mass in the tail than the normal distribution. For a review of (G)ARCH models, see Bollerslev, Engle and Nelson (1994).

Other papers have focused on different risk measures and different VaR methods. See, for example, Drudi et al. (1997), Van Goorbergh and Vlaar (1999) and Jorion (1996). We focus here on portfolio treatment and the effect of parameter uncertainty on the reported Value-at-Risk estimates.

The parameter values of a distribution function are unknown and are normally estimated using historical data. The usual approach is to plug the point estimates for the parameter values into the distribution function and treat them as given fixed figures. In fact, however, the parameter estimates incorporate uncertainty, which may be quantified by the standard errors of the parameter estimates. Uncertainty in the parameter estimates implies uncertainty about the underlying distribution function, and hence about the VaR estimate. We claim that parameter uncertainty is an important issue that should be taken into account as a VaR-model selection criterion. Parameter uncertainty affects the model choice in at least three ways.

First, parameter uncertainty decreases with the number of historical observations that have been used in order to arrive at parameter estimates. This implies that models with constant drift and volatility specifications are less attractive, since – implicitly – the assumption is made that these variables have remained constant over a long period of time. The alternative of limiting the number of historical observations leads to enormous parameter uncertainty and hence to much uncertainty in the reported VaR estimates. This leads to a preference for time-varying volatility specifications.

Next, distribution functions that allow for more complex tail-behavior also suffer from increased parameter uncertainty. Extreme historical observations determine the estimates for tail parameters. The explanation is that extreme events have a low frequency, which leads to more parameter uncertainty and hence to more uncertainty in the reported VaR estimates. There seems to be a trade-off between model complexity and parameter uncertainty. In the empirical part of the paper we make explicit the effect of parameter uncertainty by reporting not only a “best-estimate” VaR, but also a standard deviation of this VaR estimate.

Finally, the calculation of VaR estimates for a portfolio of securities usually requires a return specification for each security in the portfolio. Also the

interaction between the different securities has to be taken into account. This leads to multivariate models that include many parameters – all of which are estimated with uncertainty. This suggests that a VaR estimate for a portfolio of securities is more uncertain.

In this paper we propose a method to circumvent the adoption of multivariate models for the calculation of a portfolio's VaR. Instead of considering each asset individually, we determine the "constant maturity"-value of the current portfolio at all times in history. The "constant maturity"-value is determined by evaluating the current portfolio against historical yields, exchange rates, stock prices and volatility smiles. This constant maturity value is dealt with in detail in section 2.3. The resulting time series of constant maturity returns may be modeled in a univariate time series framework. The resulting model is used to provide a VaR estimate for the constant maturity value of the portfolio.

We claim that the constant maturity value of a portfolio over short forecasting horizons is (almost) equal to the actual value of the portfolio. The most pronounced examples of maturity-dependent securities are bonds and options. In case of bonds, the difference between the constant maturity value and the actual value arises because the discount yield with a different maturity has to be applied, and because the time-to-maturity is different. Since yields usually lie on a smooth curve, and the difference in maturity is small when the forecasting period is small, the difference will also be small. In case of options, the maturity of the contract affects the price through the time value of the option. Again, differences in time value are small when the difference in time-to-maturity is small. Moreover, the effect of the volatility of the underlying security is far more important than the time-value effect. The constant maturity assumption is in line with commonly used methods such as Risk Metrics.

An advantage of this approach is that it focuses on the portfolio return directly and not on individual security returns. In the latter approach it will in general be more difficult to adopt the complex (non-linear) relationships that are present between the individual security returns. The transformation to the univariate constant maturity portfolio automatically includes the effect of these features on the portfolio level. A second advantage is that we can restrict ourselves to univariate models, which have the advantage of less parameter uncertainty than multivariate models have, simply because univariate models have fewer parameters to estimate.

In the empirical analysis of this paper we consider an actual portfolio of securities that are affected by exchange rates and changes in the yield curve. We propose four different models for the constant maturity portfolio return. These models are the Student-t distribution with a GARCH(1,1) volatility specification and three special cases of this model. An in-sample comparison of the four models will be carried out to test which model best describes the historical data. We also analyse the out-of-sample implications for these models including the uncertainty in the VaR estimates.

The next section sets up the mathematical framework for the empirical analysis. Section 3 describes the data, section 4 reports on the empirical analysis, and section 5 concludes.

2. THE MATHEMATICAL FRAMEWORK

Consider a portfolio with value W_t at time t . The value-at-risk is a statistical estimate of a portfolio loss with the property that, with a given (small) probability α , we stand to incur that loss or more, over a given (typically short) holding period L . To reflect extreme market conditions, commonly selected values for α are 5% or 1%. Common time periods that are taken into consideration are $L = 1, 10, 20$ days. A formal definition for the VaR reads:

$$\Pr(W_t - W_{t+L} \geq \text{VaR}_{\alpha, L}) = \alpha \quad (1)$$

The probability that the decrease in value over a time period of L days exceeds the VaR estimate is α . Econometric models are not usually stated in terms of values, but rather in terms of returns. Define the continuously compounded return at time t as follows:

$$r_t = \ln \left[\frac{W_t}{W_{t-1}} \right] \quad (2)$$

We discuss three elements that are crucial for the determination of Value-at-Risk estimates in this framework. First, a probability distribution for the future investment value is required. Second, parameter uncertainty has to be taken into account. Third, we have to consider changes in *portfolio* values, which are usually not available in historical data.

2.1. Distribution Functions and Tests

We propose and compare four distribution functions for future returns: (1) the unconditional model, which uses only unconditional estimates for both expected value and the volatility of normally distributed returns, (2) the AR(1) model, which 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(1,1) structure on the volatility. All models are nested into the fourth case, which we will describe in some detail below.

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 Student-t distributed error terms reads

$$r_t = \mu + \rho(r_{t-1} - \mu) + \varepsilon_t \quad (3)$$

$$\varepsilon_t \sim t(0, \sigma_t^2, \theta) \quad (4)$$

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

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

All models are estimated with maximum likelihood. In the most general case (in which the error terms follow a Student-t distribution) the log-likelihood is given by

$$\begin{aligned} \ln L = & T \ln \Gamma\left(\frac{\theta+1}{2}\right) - T \ln \Gamma\left(\frac{\theta}{2}\right) - \frac{1}{2} T \ln[\pi(\theta-2)] + \\ & - \frac{1}{2} \sum_{t=1}^T \ln(\sigma_t^2) - \left(\frac{\theta+1}{2}\right) \sum_{t=1}^T \ln\left[1 + \frac{\varepsilon_t^2}{(\theta-2)\sigma_t^2}\right] \end{aligned} \quad (6)$$

where $\Gamma(\cdot)$ is the Gamma-function and T is the number of historical observations. In the special case of normally distributed error terms, the log-likelihood function reduces to

$$\ln L = -\frac{1}{2} T \ln(2\pi) - \frac{1}{2} \sum_{t=1}^T \ln(\sigma_t^2) - \frac{1}{2} \sum_{t=1}^T \left[\frac{\varepsilon_t}{\sigma_t}\right]^2 \quad (7)$$

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. The covariance matrix (C) is given by

$$C = - \left(E \frac{\partial^2 \ln L(p)}{\partial p \partial p'} \right)^{-1} \Bigg|_{p=\hat{p}} \quad (8)$$

where p is the vector that consists of the parameters that have to be estimated, and \hat{p} denotes the parameter values for which the log-likelihood is maximized.

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 in which the degree of freedom is determined as the difference in number of parameters. So

$$X = 2 [\ln L_U - \ln L_R] \sim \chi^2(n_U - n_R) \quad (9)$$

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.

We also set up 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 estimate the VaR for 20 days ahead (which resembles a one-month forecast)¹. The VaR estimate is then compared with the realized change in the position of the bank over the same period. This procedure is repeated using many different sub-periods, which are “moving windows” of T_e observations that are used to estimate the model. Let us denote T_t for the total number of comparisons between the actual change with the predicted VaR. Suppose that the number of violations is denoted by T_v . The null-hypothesis is that the percentage of violations of the VaR estimates is equal to a . Under the null-hypothesis, the number of violations follows a binomial distribution so a goodness-of-fit test exists. This χ^2 -test leads to a confidence level for the percentage of violations. We cannot reject the null-hypothesis at a confidence level of p that the percentage of violations of the VaR is equal to a if:

$$a - \sqrt{a(1-a) \frac{\chi_p^2(1)}{T_t}} < \frac{T_v}{T_t} < a + \sqrt{a(1-a) \frac{\chi_p^2(1)}{T_t}} \quad (10)$$

where $\frac{T_v}{T_t}$ denotes the percentage of violations and $\chi_p^2(1)$ is the critical value of the chi-squared distribution with one degree-of-freedom at a probability level p . When the percentage of violations is below the lower bound, then the VaR overestimates the risk; when the percentage of violations is above the upper bound, then the VaR underestimates the risk in the portfolio.

2.2. Parameter Uncertainty

We want to stress the existence of statistical uncertainty with respect to the VaR estimate, since the parameter estimates for the underlying return distribution are uncertain. This uncertainty is most easily reflected by a confidence interval for the reported VaR estimates. 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, since in this case the parameter estimates and standard errors can be expressed analytically and therefore illustrate the econometric properties of the parameters. For the

¹ We chose this holding period, since the book under consideration contains large parts of portfolios that can only be reported on a monthly basis. This implies, furthermore, that we must deal with time-aggregation.

unconditional model, the parameter estimates lead to the following easy closed forms for the expected value

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^T r_t \quad (11)$$

and for the volatility

$$\hat{\sigma}_t^2 = \frac{1}{T-1} \sum_{t=1}^T (r_t - \hat{\mu})^2 \quad (12)$$

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

$$C(\mu, \sigma^2) = \begin{bmatrix} \frac{\hat{\sigma}^2}{T} & 0 \\ 0 & \frac{2\hat{\sigma}^4}{T-1} \end{bmatrix} \quad (13)$$

An important observation from the covariance matrix of the parameter estimates in equation (13) 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 that incorporate only a recent sub-sample of the data to arrive at parameter estimates that conform to the most recent market developments. This motivates models that allow for time-varying expected returns and time-varying volatility.

Parameter uncertainty may be taken into account by the asymptotic distribution of the parameter estimates. In a Bayesian framework we sample the parameters p from the following parameter distribution:

$$p \sim N(\hat{p}, C) \quad (14)$$

Consider M samples of parameters, which are denoted by $p^{(1)}, \dots, p^{(M)}$. Different parameter values lead to different distribution functions of the future returns. For all these parameter values we calculate the VaR estimates following the procedure outlined earlier. This leads to M values for the VaR estimate, denoted by $VaR_{a,L}(p^{(1)}), \dots, VaR_{a,L}(p^{(M)})$. So, instead of arriving at one VaR estimate, we have come up with a sample of VaR estimates. The uncertainty in the VaR may be quantified by calculating the standard deviation of the estimates.

2.3. Portfolio Treatment

Let W_t denote the value of the portfolio at time t . The investment portfolio consists of N individual securities. The price of security i at time t is denoted

by $P_t^{(i)}$, and the number of shares of security i at time t in the portfolio is denoted by $w_t^{(i)}$. The value of the portfolio is thus

$$W_t = \sum_{i=1}^N w_t^{(i)} P_t^{(i)} \quad t = 1, \dots, T. \tag{15}$$

The price of a security depends upon security-specific properties and market variables. We assume that there are K security-specific properties. The security-specific properties of security i at time t will be denoted by $x_{1t}^{(i)}, \dots, x_{Kt}^{(i)}$. Examples of security-specific properties include the maturity of a bond and the exercise price and maturity of an option. We assume P market variables, which will be denoted by y_{1t}, \dots, y_{Pt} . Examples of market variables are stock prices, interest rates, exchange rates and volatility smiles. The price of security i at time t may be written as a function of the security-specific properties and the market variables:

$$P_t^{(i)} = f(x_{1t}^{(i)}, \dots, x_{Kt}^{(i)}; y_{1t}, \dots, y_{Pt}) \quad i = 1, \dots, N; t = 1, \dots, T \tag{16}$$

Let T be the current time, which will serve as the starting point for the VaR calculation. The number of shares in each security and the security-specific properties at time T will be kept fixed. We introduce the constant maturity prices that follow by changing only the market values and not the security specific properties:

$$\bar{P}_t^{(i)} = f(x_{1T}^{(i)}, \dots, x_{KT}^{(i)}; y_{1t}, \dots, y_{Pt}) \quad i = 1, \dots, N; t = 1, \dots, T \tag{17}$$

The constant maturity portfolio values follow as

$$\bar{W}_t = \sum_{i=1}^N w_T^{(i)} \bar{P}_t^{(i)} \quad t = 1, \dots, T. \tag{18}$$

It is possible to calculate a univariate time series of constant maturity returns for these constant maturity portfolio prices:

$$\tilde{r}_t = \ln \left[\frac{\bar{W}_t}{\bar{W}_{t-1}} \right] \tag{19}$$

These first stages have a lot in common with historical simulation, which is an alternative method to estimate the VaR. With historical simulation, these returns are seen as the “empirical distribution function” for the returns, from which the $\alpha\%$ worst case can be observed. This is the VaR estimate using historical simulation. In contrast to historical simulation, we interpret these returns as a time series instead of as a distribution function.

For these returns we adopt a time-series model that we apply afterwards to generate forecasts for future values of the constant maturity portfolio,

denoted by \tilde{W}_{T+L} . Although the forecast for the constant maturity portfolio is not equal to the forecast of the actual portfolio, for small values of L the difference will typically be negligible. For example, a bond or an option that has an actual maturity τ at time $T + L$ will have a maturity $\tau + L$ in the constant maturity portfolio. The effect on the price of a bond or option is small when the difference in maturity is small. Note that this is in line with the assumptions made in commonly used VaR estimation methods such as Risk Metrics. The classical approach adopts a multivariate time series model for all the market variables instead of creating the constant maturity portfolio. The advantage of the latter approach is that no parameters need to be estimated that are irrelevant in a generally implemented model. Moreover, a simple model for the constant maturity returns may be able to include the same information as a more complex model on the multivariate level. Also, note that some parameters that may not be finite for the separate underlying market factors may be finite for the constant maturity return as a whole. Finally, note that an important advantage of creating a constant maturity return model is that we focus exactly on what we are interested in.

3. DATA

The data over the observed period represent the market-value changes of a specific portfolio of ING Bank². The book consists of interest rate risks in both NLG and DM. Furthermore, there is, on average, a long position in the book (i.e. the downside risk was related to increases in interest rates). We apply the algorithm outlined in section 2.3 to arrive at a constant maturity portfolio. 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 through April 30, 1998.

Table 1 provides summary statistics of the daily returns of the constant maturity portfolio. The empirical distribution is a bit skewed and exhibits fat tails, given a kurtosis that exceeds 3, which sometimes leads 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 some periods in which returns are more volatile than in others, which motivates a model that accounts for time-varying volatility. Figure 2 presents the daily returns in a histogram, which reveals quite a number of extreme observations in the left tail. This motivates a distribution function that has fatter tails than a normal distribution.

The last observed value for the portfolio is the starting point from the out-of-sample VaR analysis. This value is equal to $W_T = 3000$.

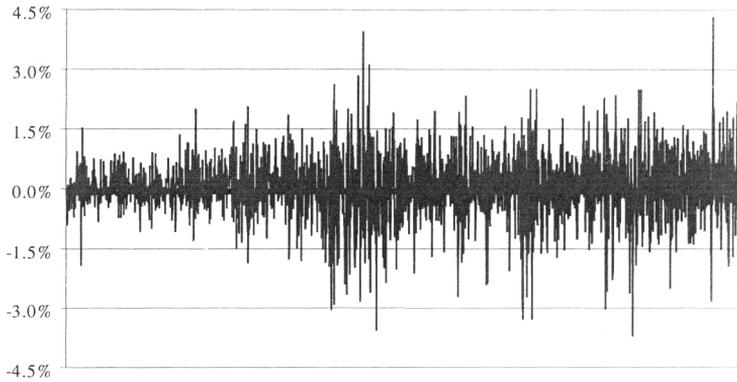
² Because of confidentiality, we have scaled the data so that the results do not represent the actual market value changes. The implications for the VaR estimates are the same as with the real data.

TABLE I
SUMMARY STATISTICS

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

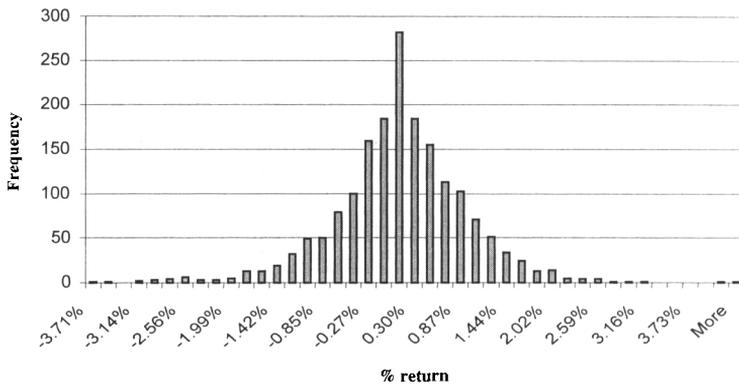
Notes: The table presents summary statistics of the returns sample, which consists of daily observations for the period of Jan. 25, 1991 through April 30, 1998.

Figure 1. Time Series Returns



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

Figure 2. Histogram Returns



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

4. RESULTS

This section describes the estimation results for the four models under consideration. We start with an in-sample description and formally test the alternative specifications against each other. We then present the implied VaR estimates under the alternative specifications and provide a measure of reliability. We end with an out-of-sample test of the models, in which we compare the reported VaR estimates 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 (6). The remaining three models are special cases of the likelihood function in equation (7). Table 2 gives the parameter estimates for the four models. 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 relationship between two consecutive daily returns, given the estimate for ρ . The parameter estimates for the N-GARCH(1,1) model show the same relationship for consecutive returns as that found 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 those of the N-GARCH(1,1) model, but now we also find a significant estimate for the degrees of freedom. The estimate for θ implies that the distribution of the returns has more probability mass in the tail than in the 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. Twice 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 $\chi_{0.95}^2(1) = 3.84$, we reject the restriction ($\rho = 0$) and prefer the AR(1) model over the unconditional model. The data thus 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 $\chi_{0.95}^2(4) = 9.49$, we reject the restriction of constant volatility, and prefer the N-GARCH(1,1) model, which allows volatility to vary over time. Testing the t-GARCH(1,1) model against the N-GARCH(1,1) model results in a preference for the former, since 84 is greater than the critical value of $\chi_{0.95}^2(1) = 3.84$. From the results in table 2 we conclude that, with respect to the in-sample behavior of the four models, the data favor the t-GARCH (1,1) model.

TABLE 2
ESTIMATION RESULTS

Parameters	t-GARCH(1,1)	N-GARCH(1,1)	AR(1)	Unconditional
$\mu (\times 10^{+3})$	-0.440 (0.140)	-0.400 (0.150)	-0.320 (0.190)	-0.300 (0.190)
ρ	-0.083 (0.023)	-0.059 (0.023)	-0.060 (0.024)	"0" (-)
$\gamma_0 (\times 10^{+5})$	0.025 (0.011)	0.033 (0.014)	6.57 (0.220)	6.59 (0.221)
$\gamma_1 (*)$	0.950 (0.011)	0.946 (0.010)	"0" (-)	"0" (-)
$\gamma_2 (*)$	0.046 (0.010)	0.049 (0.009)	"0" (-)	"0" (-)
θ	6.309 (0.766)	" ∞ " (-)	" ∞ " (-)	" ∞ " (-)
$\sigma_1 (\times 10^{+2})$	0.390 (0.106)	0.408 (0.096)	0.810 (0.014)	0.812 (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. σ_1 stands for the standard deviation at the first time period, and lnL denotes the value of the loglikelihood function. The models are estimated using daily observations for the period of June 25, 1991 through 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

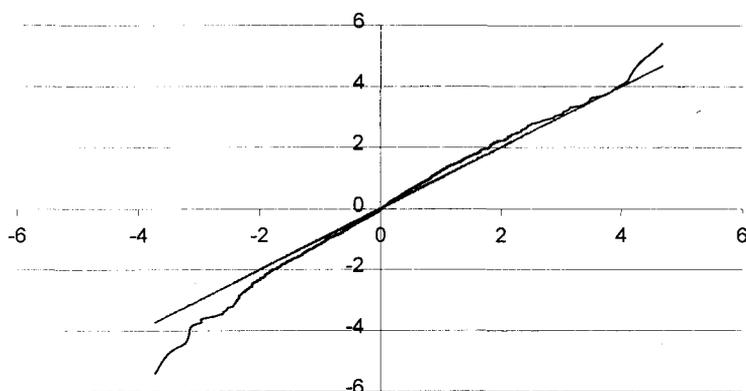
Notes: This table reports the third and fourth moments of the standardized residuals for the t-GARCH(1,1) model, the N-GARCH(1,1) model, the AR(1) model and the Unconditional model. The associated standard errors are 0.058 for skewness and 0.12 for kurtosis.

Table 3 presents 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

(*) The condition for a bounded fourth moment is just violated. In order to conclude that the reported standard errors for the volatility do have a meaning, we also determined the estimates using bootstrapping (so by simulation). The results obtained by bootstrapping are nearly identical to the results above. Therefore we conclude that, when the condition for a bounded fourth moment is only just violated, the reported standard errors of the volatility as calculated in traditional ways still have a meaning.

result in residuals that show extreme kurtosis. Furthermore, to see whether the t-GARCH(1,1) model is a good model, we performed some more residual analysis which is summarized in figure 3: the QQ-plot. The QQ-plot shows that only in the extreme part of the tail there is a deviation between the observed residuals and the modelled residuals, which could indicate even fatter tails than modelled with the t-distribution. Note that this may well be the result of parameter uncertainty. We explicitly take parameter uncertainty into account in this paper. The out-of-sample results to be discussed shortly show that this is exactly what needs to be done in order for the observed violations to be consistent with the implied violations.

Figure 3. QQ-plot for residuals t-GARCH(1,1)



Notes: The figure displays a QQ-plot of the observed and the modelled residuals for the t-GARCH (1,1) model.

The four models are used to generate scenarios for future daily returns for up to 20 periods ahead (which corresponds to a period of one month). The 20 generated daily returns are added to arrive at the total return over the one-month period. Generating many different scenario-paths provides a distribution of possible future portfolio values. Given all these possible future portfolio values, we are able to calculate the VaR estimate.

Since the parameter values that serve as input for the scenario model are uncertain, we take account of this uncertainty by repeatedly drawing values for the parameters as given in equation (14). This results in an expected VaR estimate, together with a standard error that reflects the uncertainty in this estimate. In table 4 these values are given for all models. The value of the initial portfolio is equal to $W_T = 3000$, so the different models roughly imply a VaR of 8% to 10% for a one-month period. The expected VaRs that are given by the unconditional model and the AR(1) model are almost the same. The N-GARCH(1,1) model assumes a significantly higher value at risk, since it accounts for both the negative drift in expected returns and the higher volatility in the most recent period than on average in the historical sample. This contrasts 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 of fat tails, which directly explains why the VaR estimate is higher in this case. The more realistic representation of the t-GARCH(1,1) model comes at a price, since the associated standard error of the reported VaR estimate is higher than in the other models.

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 of Jan. 25, 1991 through April 30, 1998.

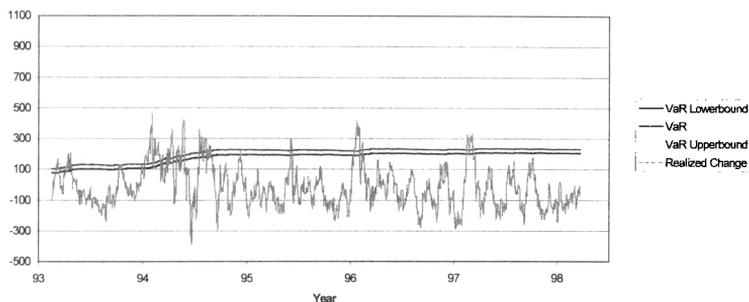
Table 5 considers the out-of-sample performance of the different models described in section 2. We consider $T_t = 1700$ sub-samples, each consisting of $T_e = 200$ observations. The sub-samples are related to each other, since they are constructed as moving windows. To arrive at a new sub-sample, we delete the first observation from the old sub-sample and add the next observation after the old sub-sample. From the sub-sample the parameters of the four models are estimated. The estimated models are used to generate VaR forecasts for an out-of-sample period of $L = 20$ days. 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. Also the upper and lower bounds of the 95 percent interval around the VaR are determined. Figures 4 and 5 present the reported VaR estimate, the associated 95 percent confidence interval for the VaR estimate and the actual change in the position of the bank for the unconditional and the t-GARCH(1,1) model, respectively.

TABLE 5
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

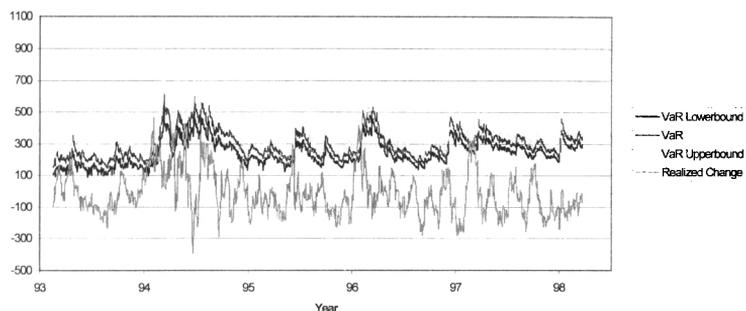
Notes: The table reports the actual percentage of violations of the predicted VaR under alternative model specifications. An appropriate model should result in a violation of 1 percent. The uncertainty in the reported VaR is reflected by the 95 percent lower- and upperbounds.

Figure 4. Unconditional model



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

Figure 5. Student-t Distribution GARCH(1,1)



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

For each of the different models, we test whether the models that forecast 1% violations of the VaR estimates, indeed result in 1% violations (within a certain confidence interval). According to equation (10), under the null-hypothesis of $\alpha = 1\%$ at a confidence level of $p = 5\%$, the number of violations

T_v should be in the confidence interval: $0.53\% < \frac{T_v}{T_t} < 1.47\%$. A model that

adequately takes account of the behavior of the return distribution would show a percentage of violations of the VaR that is within this interval. The results show that the most restricted model performs the worst, and that the percentage of violations decrease as the model becomes more general. The t-GARCH(1,1) model still shows a violation of 2.29 percent for the reported VaR estimate. If, however, we take account of the uncertainty in the VaR by considering the 95 percent upperbound on the reported VaR estimate, we then find that the reported VaR estimate implies a number of violations that is inside the proposed interval.

5. CONCLUDING REMARKS

In this paper we have compared four alternative models to calculate VaR estimates for the value of a certain portfolio 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 proposed adopting a Student-t distribution, since it allows for fatter tails than a normal distribution.

Second, the VaR estimate is based on historical return observations. Recent market circumstances should be most informative on the implied future return distribution. This is accomplished with either a time-varying return distribution (based on a large historical data sample) or with an unconditional distribution (based on most recent observations only). The first method is preferable 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 estimates also incorporate uncertainty. We have reported the VaR estimates together with a standard error. The empirical implications are that a relatively long time series is required in order to arrive at a relatively reliable VaR estimate (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 the most recent market circumstances, and since it allows for fat tails.

Fourth, in order to model the VaR of a portfolio of securities, we transform the data into a constant maturity portfolio. This results in a univariate time series that increases the reliability of the VaR estimation. The alternative is to consider a multivariate time-series model that requires estimation of more parameters. This decreases the reliability of the estimation results.

Fifth, the out-of-sample tests show that, when comparing the realized change in value of the constant maturity portfolio with the VaR estimates, all models result in significantly more than 1% violations of the 1%-VaR. The t-GARCH(1,1) model reports a VaR that shows the least number of violations. Taking into account the uncertainty of the reported VaR estimate, we cannot, in fact, reject the result that the t-GARCH(1,1) model adequately describes the VaR of a portfolio.

REFERENCES

- BOLLERSLEV, T., ENGLE R.F. and NELSON, D. (1994) ARCH Models, in R.F. Engle and D. McFadden, eds, *Handbook of Econometrics, volume IV*, North Holland.
- DRUDI, F., GENERALE, A. and MAJNONI, G. (1997) "Sensitivity of VaR measures to different risk methods", *Temi di Discussione*, 317, Banca D'Italia. Available upon request from Banca D'Italia, www.bancaditalia.it

- EMBRECHTS, P., KLUPPELBERG, C. and MIKOSCH, T. (1997) *Modelling Extreme Events for Insurance and Finance*, Springer Verlag, Berlin.
- JORION P. (1996) "Risk: Measuring the risk in Value at Risk", *Financial Analysts Journal*, **52**, 6, 47-56.
- JORION P. (2000) *Value at Risk: The New Benchmark for Managing Financial Risk*, McGraw-Hill, New York.
- JP MORGAN BANK (1996) *RiskMetrics Technical Document*, New York: JP Morgan Bank. Available from: www.riskmetrics.co.jp
- LUCAS, A. and KLAASSEN, P. (1998) "Extreme returns, downside risk and optimal asset allocation", *Journal of Portfolio Management*, **25**, 71-79.
- MCNEIL, A.J. and FREY, R. (2000) "Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach", *Journal of Empirical Finance*, **7**, 3-4, 271-300.
- VAN DEN GOORBERGH, R.W.J. and VLAAR, P.J.G. (1999) "Value at Risk Analysis of stock returns: Historical simulation, variance technique or tail index estimation?", *research memorandum WO&E*, 579, De Nederlandsche Bank. Available via: <http://www.gloriamundi.org/var/wps.html> from www.dnb.nl/publicaties/pdf/staff40.pdf

JACCO WIELHOUSER
ING Bank
PO BOX 1800
1000 BV Amsterdam – The Netherlands
Telephone: + 31 20 576 1442
Fax: +31 20 576 1473
Email: Jacco.Wielhouwer@mail.ing.nl
J.L.Wielhouwer@kub.nl