

A Comparison of Value at Risk Approaches and Their Implications for Regulators¹⁾

Gabriela de Raaji,
Burkhard Raunig

1 Introduction

Over the last decades many financial institutions have significantly stepped up their trading activities, especially in the field of derivatives. Jorion (1997) identifies increased volatility, technologically enhanced physical equipment, advances in finance theory and political developments, such as more market-oriented policies and deregulation of financial markets, as the driving force behind this process. In addition, many new financial products have been developed, some in response to regulation. Although such products may offer certain advantages in hedging financial risks or provide speculative opportunities, under certain circumstances they may also generate huge losses. The last decade has, indeed, seen spectacular financial disasters related to derivatives trading. Britain's 233-year-old Barings bank went bankrupt on February 26, 1995, when Nick Leeson lost USD 1.3 billion from derivatives positions. Other well-known cases are the USD 1.3 billion loss the German Metallgesellschaft firm suffered and Daiwa Bank's USD 1.1 billion loss. Most of those (and other) financial disasters could probably have been avoided, had properly functioning internal controls and adequate risk management systems been in place.

The financial industry and regulatory authorities have clearly recognized that, in order to ensure financial stability, it is imperative to accurately measure financial risks and implement sound risk management. The concept of Value at Risk (VaR), in particular, has received much attention and is now widely accepted as a useful measure of financial risk. In short, VaR is the expected maximum loss over a target horizon for a given confidence interval. To be more precise, let P be the price of a portfolio that contains m contracts C_j ($j = 1 \dots m$) where the changes in value of the contracts ΔC_j depend on n risk factors r_i ($i = 1 \dots n$). These risk factors are stochastic and may be specific equity prices, interest rates, exchange rates, etc. The portfolio profit/loss ΔP over a given horizon is a function of the changes in the value of the contracts. Thus, the change in the value of a portfolio $\Delta P(r_1, \dots, r_n) = \sum_m \Delta C_m(r_1, \dots, r_n)$ may be expressed as a function of the underlying risk factors.²⁾ Let $F(\Delta P)$ be the cumulative probability distribution of the changes in the market value of a portfolio. So, VaR can formally be defined as

$$\text{VaR} = F^{-1}(p), \quad (1)$$

where p is a specified probability, for example 0.05 or 0.01, and $F^{-1}(p)$ denotes the inverse of $F(\cdot)$. Thus, losses greater than the estimated VaR should only occur with the probability p . For example, if a VaR calculated at the 95% confidence level is accurate, then losses greater than the VaR measure, so-called "tail events", should on average only occur five times in every 100 trading days.

The VaR approach is attractive because it is easy to understand (VaR is measured in monetary units) and it provides an estimate of the amount of capital that is needed to support a certain level of risk. Another advantage of this measure is the ability to incorporate the effects of portfolio diversification. Many banks (and other financial institutions) now base their

assessment of financial risk and risk management practices on VaR or plan to do so in the future.

According to the Basle Committee proposals, in most countries banks have the option to use VaR models, upon their regulatory authorities' approval, for determining their capital requirements for market risk. For capital requirement purposes the model parameters are standardized and require banks to use a one-sided confidence interval of 99%, an assumed holding period of 10 days and at least one year of historical data for the market risk factors underlying their trading books. Although the model parameters are standardized, banks do not have to employ any one particular approach to estimate VaR. In other words, banks may choose their individual approach towards VaR. This liberal view makes sense because there is no single "best" VaR approach. Besides, ongoing research in this area is far from being completed.

Financial institutions use their VaR models on a day-to-day basis, and reported VaR numbers may also provide regulators with useful information. From a regulator's point of view it would be valuable if reported VaR numbers could be utilized to compare risk-taking across different banks at a given point in time and to track market risk exposures over time. For example Hendricks and Hirtle (1997) argue that³

"...the information generated by the models will allow supervisors and financial market participants to compare risk exposures over time and across institutions."

And that⁴) "...a capital charge based on internal models may provide supervisors and the financial markets with a consistent framework for making comparisons across institutions."

This view is unduly optimistic since different approaches and assumptions may produce systematically different VaR estimates. This paper therefore attempts to provide an answer to the question whether it makes sense to compare VaR numbers generated by means of different models. To this end, variance-covariance methods and historical simulation approaches are used to estimate VaR numbers for one equally weighted portfolio and nineteen randomly chosen linear foreign exchange portfolios over a period of 1,000 trading days. In addition, a new method is applied. The method, which was recently proposed in Hull and White (1998), deals with fat-tailed distributions, which are typical of fx returns, but also of many other financial returns. In a next step the performance of the various models is compared over the simulation period with the help of a simple backtesting procedure to determine how accurately the models match the specified confidence intervals.

The paper is structured as follows: Section 2 briefly outlines the VaR approaches on which the calculations in this paper are based. Section 3 provides a description of the data used. Section 4 deals with the application of the various methods, section 5 presents and explains the results. Finally, section 6 contains a number of concluding remarks.

2 VaR Methods

The VaR estimates presented in this paper were derived by employing variants of the variance-covariance approach, historical simulations and Monte Carlo methods based on mixtures of normal distributions as proposed in Hull and White (1998).⁵⁾ The variance-covariance approach assumes that the risk factors that determine the value of the portfolio are multivariate normally distributed, which implies that the changes in the value of a linear portfolio are normally distributed. Since the normal distribution is fully described by its first two moments, the VaR of a portfolio is essentially a multiple of the standard deviation. For a portfolio, the VaR under the variance-covariance approach is given by

$$\text{VaR} = -\alpha\sqrt{w'\Sigma w}, \quad (2)$$

where w is a vector of absolute portfolio weights, w' is its transpose, Σ denotes a variance-covariance matrix and α is a scaling factor, which is 1.65 for a 95% confidence interval and 2.33 for a 99% confidence interval. Formula (2) implies that an estimate of the covariance matrix of the risk factors is needed. The variances and covariances are usually estimated from daily historical time series of the returns of the relevant risk factors using equally weighted moving averages such as

$$\sigma_{ijT}^2 = \sum_{t=T-n}^{T-1} \frac{r_{it}r_{jt}}{n}, \quad (3)$$

where the mean is often assumed to be zero.⁶⁾ In expression (3) σ_{ijT}^2 denotes a variance (or covariance) at time T , r_{it} and r_{jt} are returns and n is the number of observations, i.e. the window length, used to calculate the variances and covariances.

Another frequently used estimator is the exponentially weighted moving average (EWMA). In contrast to equally weighted moving averages, the exponentially weighted moving average weights current observations more than past observations in calculating conditional variances (covariances). The EWMA estimator in its recursive form is given by

$$\sigma_{ij/t}^2 = \lambda\sigma_{ij/t-1}^2 + (1-\lambda)r_{it-1}r_{jt-1}. \quad (4)$$

In equation (4) the parameter λ , which is sometimes termed the “decay factor”, determines the exponentially declining weighting scheme of the observations.⁷⁾ One difference between the two estimators is that the equally weighted moving average does not account for time-dependent variances, whereas the exponentially weighted moving average does.⁸⁾ From equation (4) it is evident that an EWMA model is equivalent to an IGARCH (1,1) model without intercept.⁹⁾

The second approach used is historical simulation. In contrast to variance-covariance methods, no specific distributional assumptions about the individual market risk factors, i.e. returns, are made, and no variances or covariances have to be estimated. Instead, it is only assumed that the distribution of the relevant market returns is constant over the sample

period. To calculate VaR numbers, the returns of the risk factors for each day within the historical sample period are viewed as a possible scenario for future returns. The portfolio is evaluated under each of the scenarios and the resulting profits/losses are ranked by size in ascending order. The resulting empirical distribution is viewed as the probability distribution of future profits and losses. The VaR is then determined as the quantile of the empirical profit/loss distribution that is implied by the chosen confidence level.

The approaches described above offer certain advantages and disadvantages. For example, the variance-covariance approach is relatively easy to implement and VaR numbers can be calculated quickly. On the other hand, the method is problematic if the portfolio contains a significant amount of nonlinear financial instruments, such as options, because then the resulting profit/loss distribution is typically not normally distributed. Another problem arises if the distributions of the underlying risk factors are not normal. Then the joint distribution of the risk factors cannot be derived analytically in most cases. Finally, the resulting VaR is very much contingent on the method used to estimate the variance-covariance matrix. Historical simulation methods avoid many of the problems of the variance-covariance approach because the underlying risk factors need not be normally distributed and the method can deal with nonlinear portfolios. In addition, no variance-covariance matrices have to be estimated. On the other hand, the method is data-intensive and requires more computer power. What is more, the resulting VaR depends heavily on the chosen window length of historical data.

The main idea of the third approach is to transform the original data in such a way that the resulting data are normally distributed. The convenient properties of the normal distribution may then be exploited. Let e_{it} be the return of risk factor i on day t and let G_{it} be the assumed probability distribution for e_{it} . The goal is to transform e_{it} into a new variable f_{it} that is normally distributed using the transformation

$$f_{it} = N^{-1}[G_{it}(e_{it})], \quad (5)$$

where N is the cumulative probability function of a normal distribution and N^{-1} is its inverse. Thus, the original variables e_{it} are mapped into variables f_{it} that are standard normally distributed on a “fractile to fractile” basis. To make this method operational, the functional form of the G -distributions of the risk factors must be chosen and the parameters of these distributions have to be estimated using historical data. The choice of the G -functions obviously depends on the characteristics of the distributions of the risk factors that drive the value of the portfolio (the specific choice for this paper is presented in a later section). Given the parameters of the G -functions, the f_{it} variables can be mapped back into actual outcomes using the relationship

$$e_{it} = G_{it}^{-1}[N(f_{it})]. \quad (6)$$

This methodology has the advantage that it can deal with risk factors that are not normally distributed. This is important when the objective is to

calculate VaR numbers using financial returns which are typically fat-tailed. Fat-tailed distributions imply that extreme observations are more likely to occur in a normal distribution. In addition, the method can also easily deal with nonlinear portfolios.

For this paper, Monte Carlo simulations were run with a view to generating a large number of f_{it} variables from standard normal distributions. To simulate the joint distribution of market risk factors, the correlation between the risk factors is incorporated via Cholesky factorization. The generated f_{it} variables are mapped into actual outcomes by using relationship (6). Individual portfolios may then be evaluated under each simulation trial. From the resulting profit/loss distribution (under the mapped outcomes e_{it}) VaR numbers can be calculated by using the appropriate quantile of this distribution.¹⁰⁾

3 Data

The methods described above are applied to one equally weighted portfolio and nineteen randomly chosen fx portfolios. The assumption is that an investor holds a certain amount of dollars in foreign currencies. This is why changes in the value of these portfolios depend solely on changes in exchange rates. The calculations are based on the assumption that the amount invested in each portfolio is USD 100 billion. The reason for the choice of simple linear portfolio structures is that matters should not be complicated by issues concerning the valuation and mapping of complex financial instruments. Such complications would only add extra noise to the comparisons.

All of the portfolios contain the Australian dollar (AUD), Belgian franc (BEF), Swiss franc (CHF), Deutsche mark (DEM), Danish krone (DKK), Spanish peseta (ESP), French franc (FRF), British pound (GBP), Italian lira (ITL), Japanese yen (JPY), Dutch guilder (NGL), Swedish krone (SEK) and Austrian schilling (ATS). Daily exchange rates covering the period from June 16, 1986, to June 15, 1998, are used, which gives a total of 3,131 observations for each individual time series.¹¹⁾ All distributions of the returns of the individual currencies display excess kurtosis (see Table 1).¹²⁾

Table 1

Excess Kurtosis of Exchange Rate Distributions			
(Currencies Quoted against the USD)			
Currency	Excess Kurtosis	Currency	Excess Kurtosis
AUD	4.84	GBP	3.44
BEF	2.91	ITL	8.49
CHF	2.04	JPY	4.81
DEM	2.30	NLG	3.50
DKK	4.18	SEK	6.07
ESP	5.69	ATS	2.87
FRF	3.04		

Source: OeNB.

As mentioned above, the fat tails of the distributions imply that extreme market shocks are more frequently observed than in normal distributions.¹³⁾ For example, if we wanted to calculate the VaR for a position in a single currency at the 99% level of confidence, we would use 2.33 times the

standard deviation, assuming a normal distribution. If the underlying distribution has fat tails, we would underestimate the VaR because of the higher probability mass of the distribution on the left tail. Table 2 demonstrates this problem for the 1% and 5% quantiles of the empirical distributions.

For each currency the 1% quantile exceeds the 2.33 multiples implied by a normal distribution. On average, the 1% quantile is located 2.62 standard deviations below the mean. At the 5% quantile some multiples are above 1.65, as implied by the normal distribution, but the majority is below, indicating a tendency for slightly too conservative VaR estimates.

Table 2

Empirical 5% Quantiles and 1% Quantiles as Multiples of the Standard Deviation		
Currency	5% Quantile	1% Quantile
AUD	1.49	2.46
BEF	1.64	2.61
CHF	1.66	2.80
DEM	1.68	2.66
DKK	1.63	2.56
ESP	1.56	2.49
FRF	1.62	2.75
GBP	1.60	2.60
ITL	1.57	2.54
JPY	1.64	2.80
NLG	1.62	2.64
SEK	1.53	2.53
ATS	1.66	2.60
Average	1.61	2.62

Source: OeNB.

4 Application of the Various VaR Models

This section describes the specific applications and variants of the VaR models used in the comparisons. Each model is used to generate daily VaR estimates of the overnight risk, i.e. on the assumption of a one-day holding period, inherent in each of the twenty portfolios at the 99% and 95% confidence intervals for the last 1,000 trading days of the sample. All calculations rest on the assumption that the means of the daily return series are zero.

The first model is the variance-covariance approach. The first variant of this model is based on daily variances and covariances estimated by means of equally weighted moving averages with a window length of 250 actual trading days. The equal weighting scheme implies that the VaR numbers generated by this model do not account for time-dependent variances.¹⁴⁾ Since there is much empirical evidence that variances of financial returns may be predicted, equally weighted moving averages do not seem to be very attractive estimators.¹⁵⁾

The next model is the variance-covariance approach using exponentially weighted moving averages. As opposed to the first model, the resulting VaR estimates incorporate effects, e.g. the well-known volatility clustering, of time-dependent variances. It follows from equation (4) that the persistence of the estimated variances (and covariances) depends on the chosen lambda. In accordance with J. P. Morgan, lambda is set at 0.94 for estimating the daily variances and covariances.¹⁶⁾

Both the third and fourth model are based on historical simulation with a time window of 250 and 1,250 historical scenarios, respectively. To obtain the VaR numbers, the neighboring observations implied by the 1% and 5% quantiles of the ranked changes of portfolio values are interpolated linearly. Due to the equal weighting of each historical scenario these models do not discriminate between recent scenarios and scenarios further back in time. All scenarios (implicitly) carry the same probability of occurrence. Let us assume that markets are not very volatile at the moment, while the sample used for simulating VaR still contains a significant fraction of scenarios that stem from a highly volatile period. Such a case would result in an overestimation of VaR numbers. On the other hand, if we were in a highly volatile period, we would underestimate the VaR if the scenarios are based on a low volatility period. This issue will be discussed in greater detail when the results are presented.

To implement the Monte Carlo methods based on Hull and White, it is necessary to assume a particular form of the distributions of the risk factors that determine the values of the portfolios. Following Hull and White (1998), the assumption is that the empirical distributions are generated by a mixture of two normal distributions according to

$$G_{it}(e_t) = p_i N\left(\frac{e_{it}}{u_i \sigma_{it}}\right) + (1-p_i) N\left(\frac{e_{it}}{v_i \sigma_{it}}\right), \quad (7)$$

where $G_{it}(e_{it})$ denotes the value of the cumulative probability distribution for observation e_{it} , p_i and $(1-p_i)$ is a probability, N denotes the cumulative probability distribution of a normal distribution and u and v are parameters that scale the standard deviation σ_{it} .¹⁷⁾ The parameters of the distribution must satisfy the restriction

$$p_i u_i^2 + (1-p_i) v_i^2 = 1, \quad (8)$$

since the variance of the mixture distribution must be the same as the variance of the observed empirical distribution.¹⁸⁾

The p , u , v and σ parameters have to be estimated for each individual risk factor. For technical reasons, the implied likelihood functions are not maximized directly. Instead, the data are grouped by each risk factor i into four categories: less than one standard deviation ($|e_{it}| \leq \sigma_{it}$); one to two standard deviations ($\sigma_{it} < |e_{it}| \leq 2\sigma_{it}$); two to three standard deviations ($2\sigma_{it} < |e_{it}| \leq 3\sigma_{it}$); and more than three standard deviations ($|e_{it}| > 3\sigma_{it}$). The maximization of the log-likelihood function follows

$$\sum_{j=1}^4 \alpha_{ij} \log(\beta_{ij}), \quad (9)$$

which results from the comparison of the predicted fraction of data β_{ij} implied for particular values of p , u and v with the proportion α_{ij} of the data actually observed in each category, i.e. to determine the values of p , u and v that provide the best fit for the empirical distributions of the individual risk factors.

Two different versions of the model are estimated for each risk factor by using 1,880 historical observations. In the first version the data are categorized by means of equally weighted moving averages according to equation (3) with a window length of 250 trading days. In the second version the standard deviations are estimated by using exponentially weighted moving averages according to (4) with a weighting parameter of 0.94. The estimated parameters for both versions are summarized in Table 3.

Table 3

Parameter Estimates for Mixture of Normal Distributions						
Currency	Equal Weights			EWMA		
	u	p	v	u	p	v
AUD	0.68	0.71	1.52	0.64	0.36	1.15
BEF	0.71	0.68	1.43	0.45	0.21	1.10
CHF	0.74	0.63	1.33	0.45	0.15	1.07
DEM	0.73	0.74	1.53	0.44	0.19	1.09
DKK	0.77	0.81	1.65	0.45	0.18	1.08
ESP	0.70	0.72	1.52	0.49	0.25	1.12
FRF	0.74	0.77	1.59	0.47	0.14	1.06
GBP	0.64	0.68	1.50	0.45	0.24	1.12
ITL	0.69	0.71	1.51	0.48	0.22	1.10
JPY	0.71	0.73	1.53	0.67	0.49	1.24
NLG	0.72	0.73	1.52	0.45	0.18	1.08
SEK	0.78	0.81	1.63	0.49	0.20	1.09
ATS	0.69	0.71	1.51	0.47	0.23	1.11

Source: OeNB.

Under both sets of parameters 10,000 Monte Carlo trials are run for each of the 1,000 trading days to simulate the joint distributions of the market risk factors for each of the twenty portfolios. In the simulations the estimated correlation matrices are used according to equations (3) and (4) that correspond to each individual trading day. Thus, for day t the joint distribution of the risk factors is estimated by using Choleski factorization that is implied by the estimated variances and correlations for day t for both versions of the model.

Since there is no closed-form solution for the transformations of the simulated values into the “actual” outcomes implied by the mixture of normal distributions, these values are iterated by means of the Newton scheme. Once the transformed values for day t have been obtained, the portfolios under each of the 10,000 scenario vectors for day t are evaluated and the VaR numbers are calculated as the corresponding quantiles from the resulting profit and loss distributions.

5 Results

This section states and discusses the results of the daily VaR estimates at the 99% and 95% confidence levels. All calculations are based on the assumption of a one-day holding period. In addition, the results derived from backtesting each method are presented. These results should provide information on how accurately the various methods perform.

The discussion starts out with an inspection of the plots of VaR numbers generated by the different models. Chart 1 shows the daily VaR at the 99%

confidence interval for the hypothetical portfolio with equal portfolio weights for each of the six methods.¹⁹⁾

It clearly turns out that, although all methods measure the VaR of the same portfolio, the patterns are quite different for the various methods. On average, the historical simulation with 1,250 days of historical data produces the highest VaR. It is also evident that the plots for both historical simulation methods look rather different from the plots for all other methods. The VaR computed by historical simulation often does not change for rather long periods of time, yet once it changes, it changes in an abrupt fashion. Such changes are more drastic when only 250 historical observations are used. These patterns are driven by extreme events that influence the VaR numbers over long periods of time.

The VaR numbers computed by means of variance-covariance approaches are driven by the methods for estimating the daily variance-covariance matrices. The equally weighted moving average estimator produces a much smoother VaR series than the EWMA estimator. The VaR obtained with EWMA reflects to some extent the kind of “volatility clustering” that is typical of most financial return series. When the VaR series computed with unweighted and weighted moving averages are compared more closely, it turns out that the unweighted VaR reacts more slowly to changes in market volatility. For example, although market volatility falls sharply over the period from the 300th to the 500th day of the simulations according to the EWMA-based VaR series, it stays at an approximately constant high level until around the 400th day and then falls only gradually for the equally weighted VaR series. Throughout this period the VaR always lies above the EWMA-based VaR.

This is followed by an examination of the differences between the various methods and the EWMA-based variance-covariance approach (benchmark) for the 99% confidence interval. This approach was chosen as the benchmark because it is frequently used and the variance-covariance matrices are available on the Internet (provided by J. P. Morgan/Reuters) free of charge. Table 4 illustrates the results of the comparisons. Obviously, the differences can be extremely large (276%), as observed for the historical simulation method with a window length of 250 trading days.

Table 4

Differences between VaR Methods

with a 99% Confidence Interval in Percent
(Benchmark: Variance-Covariance Model with EWMA)

Model	Minimum ¹⁾	Maximum ²⁾	Mean ³⁾
VCunw	0.0	133.7	25.7
HS250	0.0	243.1	59.1
HS1250	0.0	276.5	53.0
MNunw	0.0	167.6	31.0
MNewma	0.0	11.2	4.1

Source: OeNB.

VCunw: Variance-covariance approach with equally weighted moving averages.

HS250: Historical simulation with 250 days of historical data.

HS1250: Historical simulation with 1,250 days of historical data.

MNunw: Mixture of normal distributions approach with equally weighted moving averages.

MNewma: Mixture of normal distribution approach with exponentially weighted moving averages.

¹⁾ Minimum denotes the minimum difference observed across all twenty portfolios.

²⁾ Maximum denotes the maximum difference observed across all twenty portfolios.

³⁾ Mean denotes the average difference across all twenty portfolios.

Consider a regulator who compares the VaR numbers of two banks in that period, with bank A using the parametric approach with EWMA and bank B using a 250-day historical simulation. Relying solely on the reported numbers, the regulator would conclude that bank B's trading book was nearly three times as risky as that of bank A, although both banks hold identical portfolios. It is obvious from Table 4 that similar conclusions hold for all other comparisons with the exception of the EWMA-based mixture of normal distributions model.²⁰) The results presented in Table 4 clearly demonstrate that comparisons of risk exposures across financial institutions with VaR measures generated by different methods may lead to serious misinterpretations.

Although the results bear testimony to the fact that the differences between the methods may be huge, one should not conclude that the Value at Risk concept itself is flawed. First, the differences are, on average, in the range from 25 to 59%, which is not negligible but far below the observed maximum differences. Second, comparisons of risk exposures within institutions, e.g. among trading desks or different risk categories, etc., are useful if the calculations are based on the same methodology and VaR numbers are interpreted not only in an absolute sense but also in a relative context.

It is interesting to compare the VaR estimates from the mixture of normal distributions methods and the variance-covariance approaches. It is obvious from Chart 2 that at the 99% confidence interval the mixture of normal distributions model with variances based on equally weighted estimators always produces higher VaR numbers than the corresponding variance-covariance approach. This reflects the fact that the mixture of normal distributions method incorporates the excess kurtosis of the underlying market risk factors. On the other hand, the differences are small for both methods if the variances are estimated with exponentially weighted moving averages, although in this case the mixture of normal distributions VaRs provide a kind of upper boundary (see Chart 3). This finding indicates that the EWMA-based variances reduce, but do not eliminate, the effects of excess kurtosis of the distributions of the risk factors.

Chart 4 shows the VaR numbers at the 95% confidence level for the mixture of normal distributions model and the variance-covariance approaches. In the case of equally weighted estimators the mixture of normal distributions method generates VaR numbers that most of the time are slightly below the numbers derived from the corresponding variance-covariance approach. The reverse pattern occurs in the case of EWMA-based VaR calculations.

Having discussed the VaR patterns of the various approaches, it is interesting to compare the methods via backtesting to evaluate their accuracy with respect to the specified confidence intervals. To this effect, the models are tested by comparing the estimated VaR of each portfolio for day t with the profits/losses realized by the portfolios on day t . The cases in which the realized losses exceed the estimated VaR are then counted for each portfolio and method. Table 5 presents the percentages of observed "outliers" or "tail events" for each method averaged over the twenty portfolios. The minimum

and maximum numbers (in percent) of tail events are stated as well. The average percentage of tail events provides information about how accurate a method matches a specified confidence interval, i.e. the implied quantile of the profit/loss distribution.

Table 5

Backtesting of VaR Estimates				
Method	Minimum	Maximum	Mean	Standard deviation
99% Confidence Interval				
VCunw	1.3	2.1	1.790	0.20494
VCewma	0.9	1.6	1.305	0.17614
HS250	1.3	2.1	1.790	0.20494
HS1250	0.9	1.6	1.305	0.17614
MNunw	0.8	1.6	1.170	0.23864
MNewma	0.7	1.7	1.010	0.22455
95% Confidence Interval				
VCunw	4.3	5.3	4.780	0.28023
VCewma	3.8	4.8	4.250	0.23508
HS250	4.3	5.3	4.250	0.28023
HS1250	3.8	4.8	4.250	0.23508
MNunw	4.6	5.6	5.030	0.28488
MNewma	3.7	4.6	4.160	0.25215

Source: OeNB.

VCunw: Variance-covariance approach with equally weighted moving averages.

VCewma: Variance-covariance approach with exponentially weighted moving averages.

HS250: Historical simulation with 250 days of historical data.

HS1250: Historical simulation with 1,250 days of historical data.

MNunw: Mixture of normal distributions approach with equally weighted moving averages.

MNewma: Mixture of normal distributions approach with EWMA.

A perfect model would, for instance, produce 1% and 5% tail events at the 99% and 95% confidence interval, respectively. The variance-covariance method with equally weighted moving averages and the historical simulation with 250 historical scenarios show the weakest performance at the 99% level. Both methods tend to produce the greatest fraction of outliers on average. Note that even the minima are above 1% in both cases. If we interpret the average percentages of tail events as tail probabilities, we see that for both methods the probability of losses exceeding the estimated VaR is 1.8% and not 1% as implied by a 99% confidence interval. The historical simulation with 1,250 days of data and the EWMA-based variance-covariance approach produce somewhat better results. Not unexpectedly, the Monte Carlo methods based on mixtures of normal distributions are the most accurate. Both methods come very close to the specified probability of 1%. The Monte Carlo simulation with the EWMA updating scheme matches the 99% confidence interval almost precisely. The results indicate that models that do not account for fat tails tend to underestimate VaR numbers at the 99% confidence interval. Five of the six methods generate somewhat too conservative VaR estimates at the 95% confidence interval. The Monte Carlo simulations based on the equal weighting scheme come closest to the 5% fraction implied by a 95% confidence interval. Note that the Monte Carlo approach based on the mixture distributions with EWMA produces the lowest fraction of tail events in this case.

6 Conclusions

This section of the paper analyzes the six different approaches used to estimate the value at risk. Two methods were based on the variance-covariance approach with equally and exponentially weighted moving averages, and two methods were based on historical simulation with different historical period lengths. Both types of models are commonly used by financial institutions to compute VaR. Furthermore, since many financial return distributions display excess kurtosis, a new method based on mixtures of normal distributions was applied to incorporate fat tails in the VaR estimates.

A comparison of the various methods revealed that the resulting VaR numbers may differ extremely for identical portfolios. With linear fx portfolios, differences sometimes exceeded 200% when the methods were compared with the EWMA-based variance-covariance approach as the benchmark. Even average differences fell into the 25 to 59% range. The results indicate that it may be highly misleading to compare VaR numbers across financial institutions if the reported numbers are based on different methods. However, it has to be pointed out that the Value at Risk concept itself is an extremely useful tool for financial institutions with regard to their in-house risk management. Provided VaR calculations are based on a single methodology, comparisons across trading desks, risk categories, etc. provide valuable information for risk management purposes.

Backtesting was used to investigate the performance of the various methods with respect to specified confidence intervals. The results are consistent with the conjecture that methods that do not incorporate excess kurtosis tend to underestimate VaR at the 99% confidence interval. On the other hand, the same methods tend to overestimate VaR at the 95% confidence interval. For both confidence intervals one particular version of the Monte Carlo simulations which is based on mixtures of normal distributions and incorporates fat tails performed best.

References

- Alexander, C. (1996 a).** Volatility and correlation forecasting. In Alexander, C. eds. 1996 b: 233–260.
- Alexander, C. eds. (1996 b).** The Handbook of Risk Management and Analysis. New York, Toronto, Singapore: John Wiley & Sons.
- Campbell, J. Y., A. W. Lo and MacKinlay, A. C. (1997).** The Econometrics of Financial Markets. Princeton University Press.
- Dowd, K. (1998).** Beyond Value at Risk. The New Science of Risk Management. New York, Toronto, Singapore: John Wiley & Sons.
- Duffie, D. and Pan, J. (1997).** An overview of value at risk. The Journal of Derivatives (Spring): 7–49.
- Figlewski, S. (1994).** Forecasting volatility using historical data. New York University Working Paper S 94–13.
- Hendricks, D. and Hirtle, B. (1997).** Bank capital requirements for market risk: The internal models approach. FRBNY Economic Policy Review (December): 1–12.
- Hull, J. and White, A. (1998).** Value at Risk when daily changes in market variables are not normally distributed. The Journal of Derivatives (Spring): 9–19.

- Jorion, P. (1997).** Value at Risk: The new Benchmark for Controlling Market Risk. Chicago, London, Singapore: IRWIN.
- Kroner, K. F. (1996).** Creating and using volatility forecasts. *Derivatives Quarterly* (Winter): 39–53.
- Müller, U.A., Dacorogna, M. M. and Pictet, O.V. (1996).** Heavy tails in high-frequency financial data. Working Paper. Zurich, Switzerland: Olsen & Associates.

- 1 Both authors are economist in the Financial Market Analysis Division of the Oesterreichische Nationalbank. They would like to thank Helmut Elsinger, University of Vienna, as well as Gerald Krenn and Diane Moore, Financial Market Analysis Division (Oesterreichische Nationalbank), for their helpful comments. The opinions expressed in the section "Studies" are those of the individual authors and may differ from the views of the Oesterreichische Nationalbank.
- 2 The risk factors r_t are typically measured as logarithmic returns $r_t = \ln(p_t/p_{t-1})$ or as arithmetic returns $r_t = (p_t - p_{t-1})/p_{t-1}$. By using a Taylor series expansion it can be shown that for small r_t both expressions are approximately equal. All calculations in this paper are based on arithmetic returns.
- 3 Hendricks and Hirtle (1997), p. 1.
- 4 Hendricks and Hirtle (1997), p. 8.
- 5 For a comprehensive discussion of variance-covariance approaches and historical simulation methods, see, for example, Dowd (1998) or Jorion (1997).
- 6 The assumption of zero means is quite common since the means of most daily financial return series are very close to zero and are hard to estimate precisely. For more details and a comprehensive paper on this issue, see Figlewski (1994).
- 7 The allowed range of lambda is between zero and one.
- 8 A variance-covariance approach in conjunction with variances (covariances) based on exponentially weighted moving averages assumes conditional normality.
- 9 For this and other issues concerning the estimation of variance-covariance matrices, see Alexander (1996) or Kroner (1996).
- 10 For other possible ways of implementing this methodology, see Hull and White (1998).
- 11 The data were retrieved from Datastream.
- 12 The return distributions (for various frequencies) of major exchange rates are studied in Müller, Dacorogna and Pictet (1996).
- 13 Fat-tailed distributions may, for example, arise from jump diffusion processes, stochastic volatility or Markov switching. For a discussion, see Duffie and Pan (1997).
- 14 This is obvious because this estimator produces the same variances and covariances, respectively, for every possible ranking of the observations contained in the time window.
- 15 For a discussion, see Campbell, Lo and MacKinlay (1997), chapter 12.
- 16 This is, of course, not the best lambda for each individual time series, since the lambdas may be estimated separately for each time series. For an empirical justification of choosing 0.94, see RiskMetrics™ Technical Document (1996).
- 17 It is shown that a mixture of normal distributions model such as equation (7) produces distributions with fatter tails than a normal distribution with the same variance. For a discussion, see for example Duffie and Pan (1997), Hull and White (1998), or Campbell, Lo and MacKinlay (1997).
- 18 The variance of the mixture of normal distributions is given by $pu^2\sigma^2 + (1-p)v^2\sigma^2$.
- 19 The plots for the other portfolios are quite similar and do not change the conclusions.
- 20 The results for the 95% level are quite similar and therefore not reported.

A COMPARISON OF VALUE AT RISK
 APPROACHES AND THEIR IMPLICATIONS
 FOR REGULATORS

Annex

Chart 1

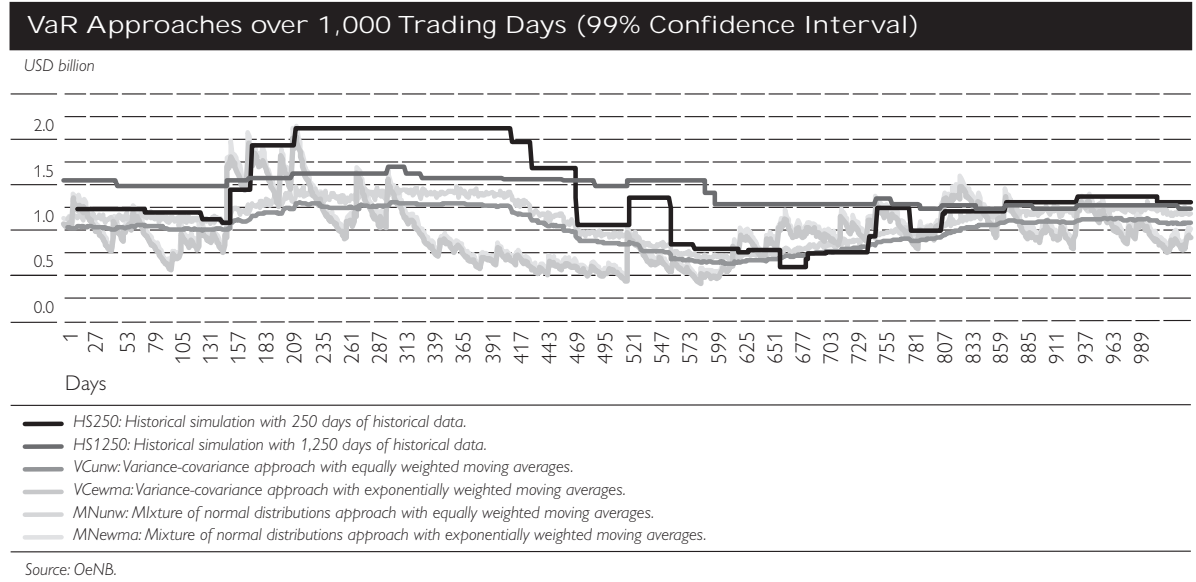


Chart 2

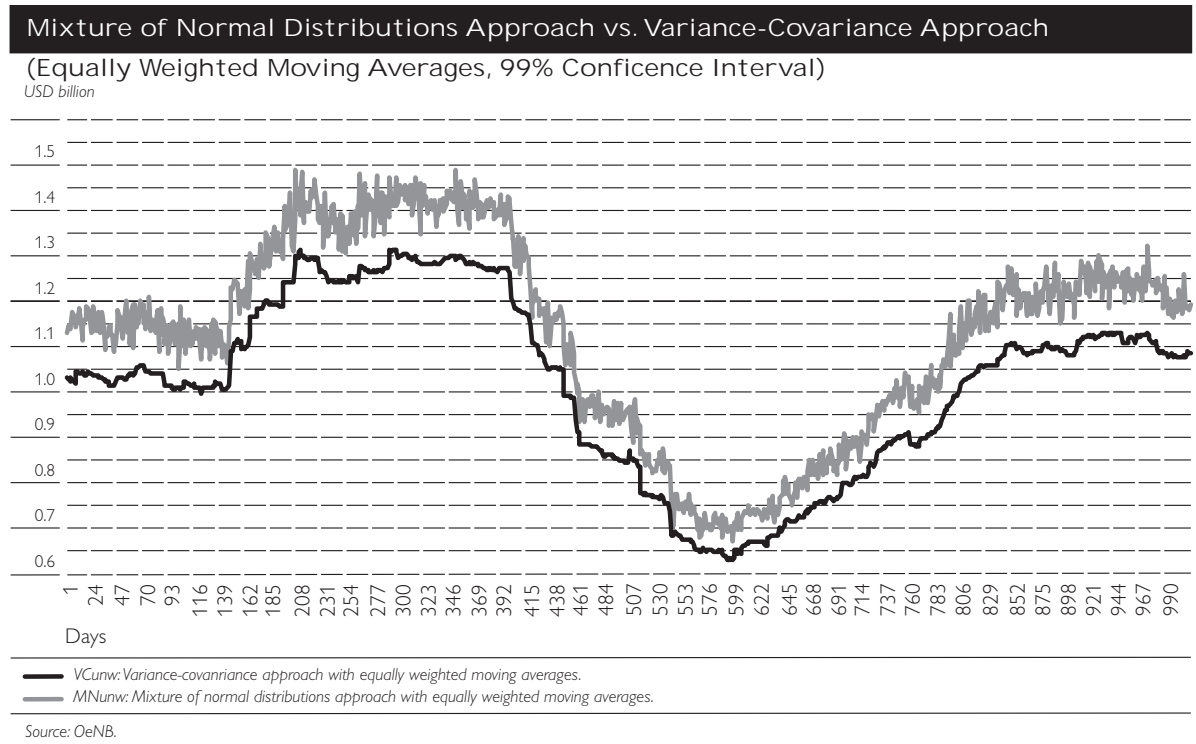



Chart 3

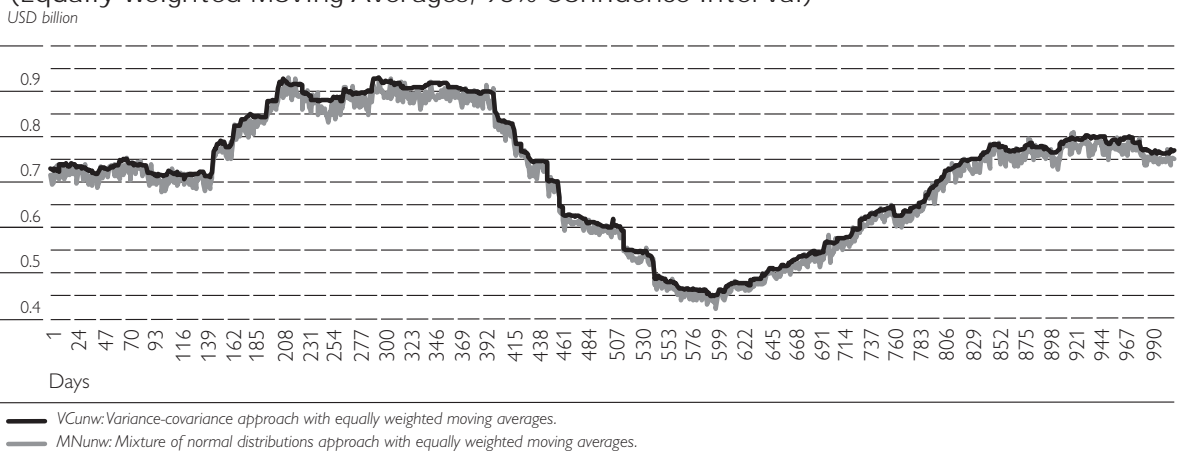
Mixture of Normal Distributions Approach vs. Variance-Covariance Approach
(Exponentially weighted Moving Averages, 99% Confidence Interval)



Source: OeNB.

Chart 4

Mixture of Normal Distributions Approach vs. Variance-Covariance Approach
(Equally Weighted Moving Averages, 95% Confidence Interval)



Source: OeNB.