VaR-x: Fat tails in financial risk management

Ronald Huisman, Kees G. Koedijk, and Rachel A. J. Pownall

To ensure a competent regulatory framework with respect to value-at-risk (VaR) for establishing a bank's capital adequacy requirements, as promoted by the Basle Committee on Banking Supervision, the parametric approach for estimating VaR needs to incorporate the fat tails apparent in the return distributions of financial assets. This paper provides a simple method to obtain accurate parametric measures by including a specific measure VaR for the tail fatness of an asset's return distribution: VaR-x. Evidence is provided for the accuracy of these VaR-x estimates by comparing different parametric VaR estimators for bi-weekly returns on US stocks and bonds.

1. INTRODUCTION

The quest for reliable risk management techniques has grown in response to higher volatility and instability on global financial markets, compounded by the enormous growth in trading activity and international exposure. One need only think of the losses made from recent currency and stock market crashes, as well as those resulting from the perilous positions taken, for example, by Barings, Daiwa, Orange County, and Metallgesellschaft. Value-at-risk (VaR) is one such risk management technique developed to improve the management of downside risk. It aims to summarize risk, by estimating the worst expected loss over a chosen time horizon within a given confidence interval. The methodology behind value-at-risk is therefore based on the probabilities associated with large negative returns and hence highlights how financial institutions have had to become more concerned with managing this downside risk. Only through the use of such risk management methods can the exposure towards large negative movements in financial markets be controlled and reduced. However, their benefit rests primarily on the accuracy of the value-at-risk estimates.

The VaR estimate is found from the probability distribution of the expected returns. This implies that one needs to make assumptions concerning the actual form of the expected return distribution. This can be done by assuming that the distribution of the expected returns equals the empirical distribution based on past observations or by assuming that the returns are drawn from a specific statistical distribution. The exact form of these analytical distributions is determined by various parameters, estimated using past data, and which have more recently also allowed for the use of conditioning methodologies, such as generalized autoregressive conditionally heteroskedastic (GARCH) processes.

A parametric approach has been the preferred method, since it enables simple conversion to take place (between quantiles and time horizons), and is hence

more pragmatic under the framework of the Basle Committee. It also enables conditionality in the data to be easily incorporated into the VaR estimate, making forecasts of VaR more appropriate. The crucial assumption therefore for an accurate estimation of the VaR is that the distribution in the left tail, reflecting the negative returns, is well represented by the specified distribution. Any discrepancy between the parametric distribution and the empirical distribution can result in large errors in the estimation of VaR.

For simplicity and convenience, asset returns are often assumed to be normally or lognormally distributed. However, the return distributions on many assets have been shown to exhibit fatter tails than the normal distribution.[1] This means that the assumption of normality results in an underestimation of the VaR on moving further into the tails. It is the exact nature of this extra mass in the tails of the distribution which is crucial when trying to capture the VaR of an asset. Other fatter-tailed distributions such as Pareto and sum-stable distributions have in the past proved difficult to implement. The normal distribution has therefore been retained as the most convenient proxy for an asset's actual distribution. However, the fatter the tails of the asset return distribution under consideration, the larger the discrepancy with the normal distribution, and the larger the errors made in VaR estimation. These errors become magnified for the million-dollar positions that mutual funds, for example, typically hold.

There is thus a need for simple methodologies to estimate VaR which capture the tail fatness apparent in return distributions. In this paper, we present such a simple technique. We show that VaR-x estimates, VaR estimates obtained from assuming the Student *t*-distribution as a fit to the empirical distribution, are better able to capture the extra risk involved for distributions exhibiting a higher probability of large negative returns. Since we are looking at downside risk, we are interested in the negative returns associated with the left tail of the distribution. The tail parameter of the Student *t*-distribution, reflected by its number of degrees of freedom, is set equal to the tail index for the left tail, and is a direct measure of the amount of fatness in the tail of the return's distribution. This method offers many advantages over the normal distribution. First, fat tails are captured. Second, focusing only on the left tail means that we do not need to assume distributions are symmetric. The tail fatness may vary between the two tails of the return distribution and hence allows for the possibility for skewness in the distribution. This provides us with a simple and more accurate estimator than would otherwise be obtained from assuming normality.

One attempt to capture the extra probability mass in the tails has been to estimate a GARCH process. The unconditional distribution of a GARCH process does reveal fatter tails; however, it has been shown that the distribution of conditional residuals is still not normal (see Bollerslev 1987). This results in the VaR still being underestimated at high quantiles for fat-tailed assets. The

---

[1] See among others Fama and Roll (1968), Rogalski and Vinso (1978), Boothe and Glassman (1987), Taylor (1986), Jansen and de Vries (1991), Loretan and Phillips (1994), and Huisman *et al.* (1997, 1998).

appealing feature, however, of incorporating conditional volatility is that it allows for a changing distribution over time. Implementing this into a VaR framework means that the VaR estimates are made conditional. This is done in J. P. Morgan's RiskMetrics, for example, and VaR-x can easily be adjusted to capture conditional volatility.[2]

The plan of the paper is as follows. VaR and VaR-x are introduced in the following section. The data used and the results are presented in Sections 3 and 4, respectively. Conclusions are then drawn in the final section.

2. VALUE-AT-RISK METHODOLOGY

Exposure to downside risk can be summarized in a single number by an estimate of the VaR. This is formally defined by Jorion (1996) as "the worst expected loss over a target horizon within a given confidence level." Following Jorion, we define W_0 as the initial investment and R as the expected return over the target horizon. W^* is defined as the lowest portfolio value at the given confidence level c, that is, the value of the portfolio should not fall below W^* with probability c. VaR is defined as the dollar loss relative to the expected mean value of the portfolio

$$\text{VaR} = \text{E}(W) - W^*. \tag{1}$$

Defining R^* as the expected return associated with the portfolio value W^*,

$$W^* = W_0(1 + R^*), \tag{2}$$

gives us the VaR measured as the dollar loss relative to the mean,

$$\text{VaR} = -W_0(R^* - \mu), \tag{3}$$

where μ is the expected return on the portfolio for the target horizon. The crux of being able to provide an accurate VaR estimate is in estimating the cutoff return R^*. In this paper, we focus on the cutoff return estimated using historical data, and thereby assume that these are representative for the expected return measure.

The statistical methods developed to best estimate these cutoff returns can be divided into two types: parametric and nonparametric. The most obvious nonparametric approach uses the historical distribution itself to compute an empirical estimate of the VaR directly. In the parametric case, one tries to fit the historical distribution by a statistical distribution whose characteristic parameters are derived from the historical data. We shall therefore briefly review the standard ways to estimate VaR before presenting the methodology behind VaR-x. The crucial difference between VaR and VaR-x is that the latter

---

[2] See Koedijk and Pownall (1998) for an implementation of a conditional VaR-x approach and comparison of results with RiskMetrics using data on Asian emerging markets.

incorporates the tail fatness apparent in financial returns into the VaR estimate, thereby improving the quality of the estimates in a simple and efficient way.

2.1 Methods to Estimate the Cutoff Return and VaR

The cutoff return is defined as the worst possible realization R^* for a confidence level c, and is found from the following integral for the distribution of expected returns $f(r)$:

$$1 - c = \int_{-\infty}^{R^*} f(r)\,\mathrm{d}r. \tag{4}$$

2.1.1 Empirical VaR Empirical VaR involves determining the point R^* from a histogram of the empirical distribution based on historical returns. R^* is that point below which are the fraction $1 - c$ of the returns. This number is then plugged into (3) to get the empirical VaR estimate.

The empirical VaR measure has some serious disadvantages to both financial institutions and regulators. In order to obtain accurate estimates a large data sample of the empirical distribution is required. The VaR estimate is therefore subject to the frequency and length of the data sample. A further drawback is the inability to allow for conditionality of the parameters over time. To overcome these flaws, a parametric approach, such as the normal approach, is often adopted. Since the distribution is approximated by a parametric distribution, parameters can be allowed to change over time. Estimation risk on the VaR estimate itself is also reduced, particularly for higher quantiles.[3] Furthermore, the parametric approach has the advantage of not being dependent on the chosen quantile, facilitating the ease with which comparisons between the VaR estimates across various institutions can be made. Parametric conversion, however, will only hold in practice if the parametric approach accurately reflects the distribution at all quantiles in the tail. Indeed, it has been the case that institutions have notoriously chosen confidence levels and time horizons to suit them. The choice of parametric distribution is therefore crucial.

2.1.2 Normal VaR The simplest parametric approach is to assume that the expected returns are normally distributed with the mean and variance estimated using past data on returns. VaR estimates are then obtained by equating $f(r)$ in (4) to the p.d.f. of the normal distribution. The simplicity of this method also explains its popularity. However, to obtain accurate VaR estimates for higher confidence levels, say more that 95%, the parametric distribution should correctly approximate the distribution in the tails. Since it is commonly known that the distributions of returns on financial assets often exhibit fatter tails than the normal distribution, one could expect a large discrepancy to exist between the tails of the normal distribution and the tails of the actual distribution. Such a discrepancy could lead to serious errors in VaR estimates. These estimates could thus be improved upon by incorporating tail fatness; one such technique that incorporates the fat tails is VaR-x, proposed in the following section.

---

[3] See Jorion (1996) and Kupiec (1995) for a greater insight into the recognition of the estimation error in VaR estimates.

2.1.3 VaR-x   The evidence that distributions of returns on financial assets have fatter tails than indicated by the normal distribution has meant that the normal approach underestimates the true value-at-risk at high quantiles. It therefore becomes apparent that in order to capture the full risk from fat-tailed assets a parametric distribution that is fatter in the tails should be used. The Student t-distribution, which also nests the normal distribution, is one obvious choice to model $f(r)$ in (4), but its parametrization has proved tedious and inconsistent in the past.[4]

The Student t-distribution exhibits fatter tails than the normal distribution. The amount of tail fatness is reflected in the number of degrees of freedom. In order to capture tail fatness correctly, one should correctly specify the exact number of degrees of freedom to be used. This has proved to be difficult for exchange rate returns (see Boothe and Glassman 1987, Huisman *et al.* 1998), but recent advances in extreme value theory makes the issue less complex.

Extreme value theory looks specifically at the distribution of the returns in the tails. Since VaR focuses predominantly on this area in the tail, extreme value theory can bring some valuable insight into improving VaR estimation (see Danielson and de Vries 1997). The tail fatness that a tail of a distribution exhibits is reflected by the tail index. It measures the speed with which the tail under consideration approaches zero. The fatter the tail, the slower the speed, and the lower the tail index given. A nice feature of the tail index is that it equals the number of moments that exist for a distribution. For example, a tail index estimate equal to 2 reveals that both the first and second moments exist, i.e. the mean and the variance, but that higher moments are infinite. All moments exist for the normal distribution, so that its tail index equals infinity by definition. Here also lies the link with the Student t-distribution. The number of its degrees of freedom reflects the number of existing moments, and the tail index can thus be used to set the number of degrees of freedom.[5]

To obtain tail index estimates, we use the estimator presented by Huisman *et al.* (1997). Unlike other tail index estimators, the estimator of Huisman *et al.* is shown to produce almost unbiased estimates in relatively small samples. This provides us with a superior estimator than previously used, and, as we shall see below, allows us to obtain robust tail index estimates from a yearly sample of daily data. Danielson and de Vries (1997) also use extreme value theory to obtain VaR estimates, but their approach has the drawback that an extremely large sample of data is required.[6]

Specifying k as the number of tail observations, and ordering their absolute values as an increasing function of size, we obtain the tail estimator proposed by Hill (1975). This is denoted below by γ and is the inverse of the tail index α. Let

---

[4] Alternative distributions to capture the tail fatness are, for example, the Pareto and sum-stable distributions or a mixture of two normal distributions.

[5] Huisman *et al.* (1998) use this method to fit the unconditional distribution of exchange rate returns.

[6] Danielson and de Vries (1997) typically require about 100,000 observations, obtained from high-frequency data.

x_i be the ith increasing order-statistic, i.e. $x_i \geqslant x_{i-1}$, based on the absolute values of the observations. The Hill estimator then reads:

$$\gamma(k) = \frac{1}{k} \sum_{j=1}^{k} \ln x_{n-j+1} - \ln x_{n-k} \,. \tag{5}$$

Following the methodology of Huisman *et al.* (1997), we can use a modified version of the Hill estimator (1997) to correct for the bias in small samples. The bias of the Hill estimator stems from the fact that the bias is a function of the sample size. A bias corrected tail index is therefore obtained by observing the bias of the Hill estimator as the number of tail observations increases up until κ, where κ is equal to half of the sample size:

$$\gamma(k) = \beta_0 + \beta_1 k + \epsilon(k) \quad (k = 1, \ldots, \kappa). \tag{6}$$

The optimal estimate for the tail index is the intercept β_0. The α estimate is just the inverse of this estimate, and it is this estimate of the tail index that we shall use in order to parametrize the Student t-distribution.

The procedure for obtaining the VaR-x estimates is therefore as follows. First, the tail index referred to by α is estimated using the Huisman *et al.* estimator for the left tail of the empirical return distribution. The focus on the left tail directly reflects the downside risk. Furthermore, the mean μ and the variance σ^2 of the return distribution are estimated. In the second step, the tail index estimate α is then used to equate the number of degrees of freedom in the Student t-distribution. Read the value S^* off the standard Student t- distribution with α degrees of freedom using appropriate tables provided in standard textbooks (see e.g. Bain and Engelhardt 1987) or statistical software. This value then needs to be converted, since the standard Student t-distribution with α degrees of freedom has a preset mean equal to zero and a variance equal to $\alpha/(\alpha - 2)$. The value S^* is then transformed into the real cutoff return $R^* = -\theta S^* + \mu$, where θ is a scale factor given by

$$\theta = \frac{\sigma}{\sqrt{\alpha/(\alpha - 2)}} \,. \tag{7}$$

The value R^* then equals the cutoff return needed to calculate the VaR-x measure for the confidence level c. Plugging the expression for R^* into (3), we obtain the VaR-x estimate for the VaR relative to the mean μ as

$$\text{VaR-x} = W_0 \theta S^*. \tag{8}$$

In the following sections, we shall apply all the above techniques to calculate the VaR for $100 million investments in both US stocks and bonds.

3. DATA

We use data from US stock and bond indices from January 1980 until August 1998, using bi-weekly returns to provide results that can easily be set against the 10-day regulatory framework adopted by the Basle Committee. The use of two different assets exhibiting different tail index alphas enables us to gauge the effect on the value-at-risk estimates from a variation in tail fatness. We use data on the S&P 500 Composite Return Index for the US and the 10-Year Data-stream Benchmark US Government Bond Return Index, both obtained from Datastream. The bi-weekly data excludes the crash of October 1987, so that we can estimate the value-at-risk consistent with normal market conditions. Summary statistics are presented using lognormal returns for the sample of stock and bond returns in Table 1.

Over the period, stocks have had an average return of 17.33% per annum, nearly twice the 10.25% return on government bonds. The volatility was however much lower for US government bonds, with the variance around a third of that prevailing on the S&P 500. Both assets appear to exhibit significant skewness as well as excess kurtosis. According to the kurtosis statistic, the extra probability mass in the tail areas of the stock returns appears to be high, and, since the distributions appear skewed, the two tails may differ dramatically. For the VaR-x estimates, we hence take any skewness in the tails into account by taking the tail index estimator of the left tail only. The effect on the VaR depends on the exact structure of the distribution of negative returns.

The degree of fat-tailedness is estimated in terms of α, calculated using the estimator developed by Huisman et al. (1997). From the gamma estimates of the

TABLE 1. Summary statistics for stocks and government bond returns. This table contains the statistics on the S&P 500 Composite Return Index and the 10-Year Datastream US Benchmark Government Bond Index for the period January 1980 until August 1998 using 486 bi-weekly total returns. The alpha estimate is calculated using a modified version of the Hill estimator for the tail indexes and is presented for the left tail.

	S&P 500 Composite Return Index	US 10-Year Government Bond Return Index
Annual Mean %	17.329	10.247
Max Return	0.153	0.089
Min Return	−0.183	−0.061
Annual St Deviation	0.146	0.086
Annual VaRiance	0.021	0.007
Skewness	−0.641	0.503
Kurtosis	9.399	5.163
Gamma Left Tail	0.233	0.143
Standard Error	0.050	0.030
Alpha Left Tail	4.285	7.009
Observations	118	121

left tail (standard errors given below), the alpha estimates are 4.29 and 7.01 for the stock and bond returns respectively. For normally distributed returns the alpha estimate tends to infinity, so we can see that both distributions exhibit a fatter left tail than the normal. Owing to this leptokurtosis, the frequency of large negative returns is greater than that reflected by the normal distribution; hence the greater the downside risk, the fatter the tails, with the equity returns exhibiting more downside risk than the bonds.

4. VALUE-AT-RISK ESTIMATES

Value-at-risk by definition should be highly sensitive to the degree to which the distribution is fat tailed: the fatter tailed the distribution, the higher the value-at-risk for a given confidence level. As we have seen, stocks have more downside risk than bonds, represented by a lower alpha estimate for the left tail index, and thus have a higher than normal probability of extreme returns. Thus, for higher confidence levels, we would expect an empirical VaR estimate to be larger than that predicted from using the parametric approach assuming normality. The higher the confidence level, and thus the quantile chosen for the VaR estimate, the greater the effect of extreme values in the asset's return distribution. This has the important implication that the existence of a fat-tailed return distribution implies that at high confidence levels the parametric-normal VaR underestimates the exposure to market risk, with the difference likely to become larger for higher confidence levels chosen and fatter tails.

To see by how much the estimates for value-at-risk are affected by the evidence that stocks have a lower alpha tail index estimate than the bonds, we compute the VaR estimates using the various approaches discussed in Section 2. In Table 2 the value-at-risk is estimated for a $100 million investment in the two assets, using both the parametric-normal approach (equally weighted moving average method for calculating volatility) and the empirical approach.

We can see how the VaR estimates increase, the higher the confidence level taken. The structure of the difference between the empirical and the parametric-normal VaR estimates is indeed what would be expected for a fat-tailed distribution. For the S&P 500 Composite Return Index, we see that for low probability levels the distribution exhibits a so-called thin waist, since the parametric-normal VaR is larger than the empirical VaR. Moving further into the tails, the VaR estimate assuming normality becomes smaller than the empirical VaR. This means that at low probability levels the parametric-normal VaR overestimates the VaR and then, as we move to higher probability levels, the parametric-normal approach underestimates the VaR.

The magnitude of the error from using the normality assumption is a reflection of the amount of tail fatness, which of course is much more significant for the stock price index. The extent of the discrepancy from using the assumption of normality for the S&P 500 and the Government Bond Index is depicted in Figure 1.

TABLE 2. Comparison of value-at-risk estimates. The value-at-risk estimates have been calculated for the two asset classes using the empirical approach (historical data) and both parametric approaches. The normal VaR estimates assume normally distributed returns, whereas the VaR-x estimates assume a fatter-tailed distribution denoted by the Student t-distribution, and use the alpha estimates for the left tail as given in Table 1. The relative VaR estimates, expressed in millions of dollars, have been calculated for a position of $100 million in the particular asset, and for a range of confidence levels.

| Confidence level (left tail)	S&P 500 Composite Return Index			US 10 Year Government Bond Return Index		
	Empirical VaR ($100m)	Normal VaR ($100m.)	Student-t VaR-x ($100m)	Empirical VaR ($100m)	Normal VaR ($100m)	Student t VaR-x ($100m)
95	4.3288	4.7176	4.3563	2.5347	2.7505	2.6812
95.5	4.6174	4.8626	4.5463	2.5986	2.8350	2.7819
96	4.7584	5.0211	4.7613	2.6912	2.9274	2.8942
96.5	4.9716	5.1967	5.0086	2.7648	3.0298	3.0212
97	5.3585	5.3943	5.2993	2.9283	3.1450	3.1678
97.5	5.7897	5.6214	5.6512	3.0463	3.2774	3.3414
98	5.9831	5.8903	6.0950	3.1674	3.4342	3.5546
98.5	6.3505	6.2240	6.6909	3.2729	3.6287	3.8315
99	7.5190	6.6722	7.5825	3.5503	3.8900	4.2273
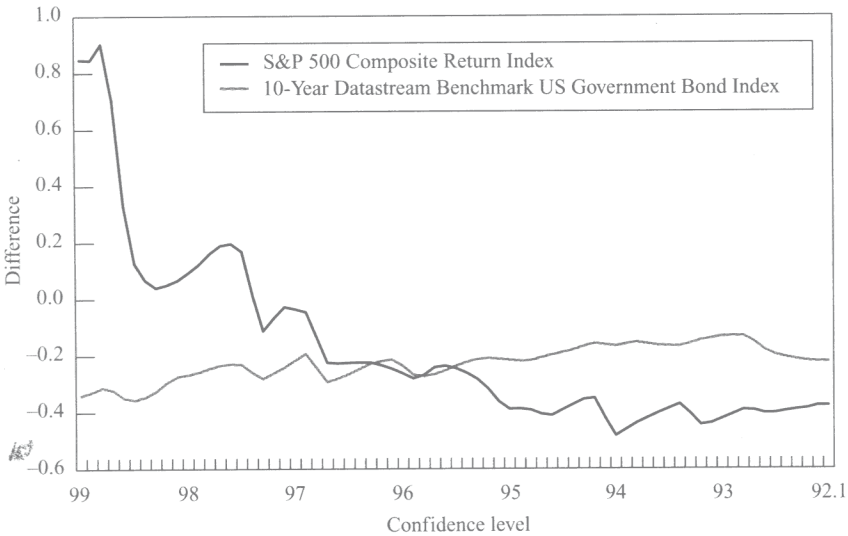
FIGURE 1. Value-at-risk estimates. The graph depicts how much the parametric-normal VaR estimates differ from the empirical VaR estimates for the two assets over a range of confidence levels. The parametric-normal approach assumes normally distributed returns and the empirical approach uses the observed frequency distribution. The difference is the error generated by using the assumption of normally distributed returns and is estimated for a $100 million position in the particular asset.
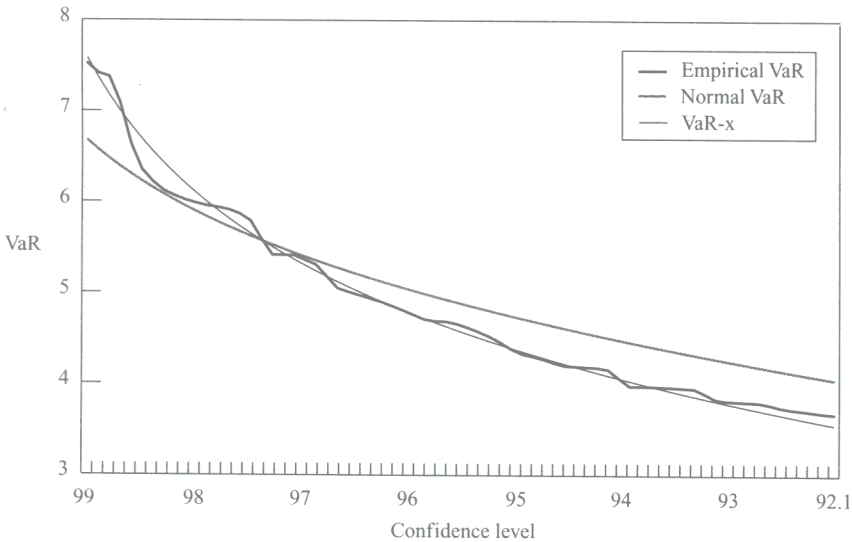
FIGURE 2. VaR-x for the S&P 500 Composite Return Index. The graph depicts how the VaR-x estimates, using the Student-*t* distribution, compare to the parametric-normal VaR estimates and the empirical VaR estimates for the S&P 500 over a range of confidence levels. The VaR-x uses the modified Hill estimator for the tail index as the parameter in the Student-*t* distribution, the parametric-normal approach assumes normally distributed returns and the empirical approach uses the observed frequency distribution. The difference is given in million dollars for a position of $100 million in the particular asset.

As predicted, the difference is larger for stocks, whose return distribution exhibits a fatter tail (a lower alpha estimate), and becomes much larger for confidence intervals above the 96% level. In the example, the assumption of normality means that the VaR is underestimated above the 96% level, and greatly underestimated at the 99% level. We therefore conclude that the assumption of normality appears inappropriate for estimating VaR at high quantiles for a distribution with an alpha estimate of around 4.

In Table 2, we indeed see that taking a $100 million position in the S&P 500 Composite Return Index generates a relative VaR estimate at the 95% probability level of $4.72 million, using the parametric assumption of normality, compared with the $4.33 million using the empirical distribution. The average bi-weekly return is $0.67 million, and the value-at-risk is stated relative to this mean. This means that, assuming normally distributed returns, 95% of the time we would not expect to achieve a bi-weekly loss of more than $4.72 million. However, at the 99% confidence interval, the VaR becomes $6.67 and $7.52 million respectively. The large discrepancy of over $0.75 million between the two approaches shows just how important it is to find as accurate a measure as possible for the VaR. Indeed, since the 99% level is the level required by the Basle Committee, it becomes apparent just how inappropriate the assumption of normality in the tails is.

The VaR-x estimates, which incorporate the fat tails, are also given in Table 2. The α estimates from Table 1 are used to parametrize the Student t-distribution. For the S&P 500 Composite Return Index, the VaR-x estimates provides a much more accurate estimate when compared with the empirical distribution for the whole range of quantiles than the parametric-normal VaR. This is illustrated in Figure 2, where all three estimates for a range of confidence levels are plotted.

For the US Government Bond Index, the difference between the two parametric approaches is much less (see Figure 3), indicating that an alpha estimate of around 7 already gives similar results to those under normality.

This approach therefore provides us with an estimator that more accurately reflects the VaR estimates for the whole range of confidence levels, and is thus a more accurate estimate for assessing the downside risk as measured by value-at-risk. We have seen that the estimator performs well for a range of quantiles up to and including the 99% level, and therefore allows for simple parametric conversion to be adhered to. Indeed, time aggregation for various holding periods is merely a simple extension to the framework, so that the estimates provided can easily be converted for different quantiles and time horizons, as required by the regulatory bodies.

To see how forecasts of the two approaches perform over time, we carry out the following out-of-sample test. We have estimated the rolling 10-day absolute VaR forecasts at the 99% level for the S&P 500 Composite Return Index, using
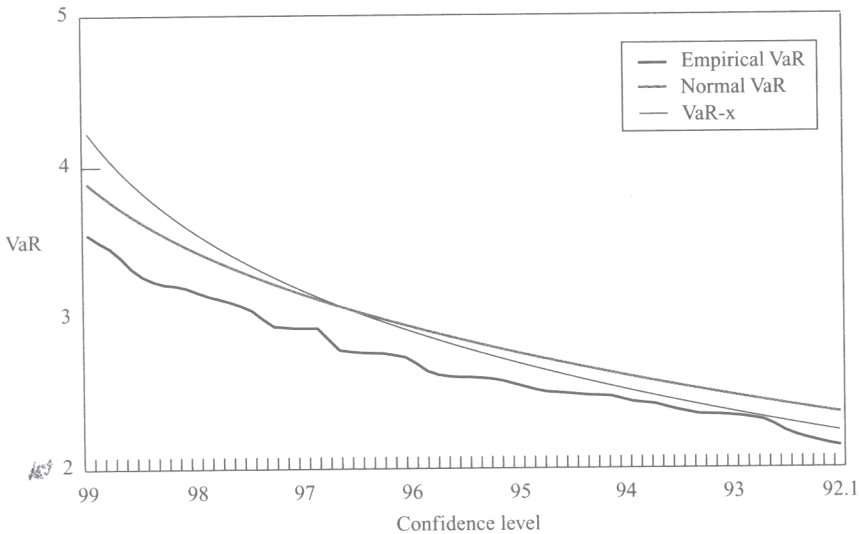

FIGURE 3. VaR-x for the 10-Year US Government Bond Index The graph depicts how the VaR-x estimates, using the Student-t distribution, compare to the parametric-normal VaR estimates and the empirical VaR estimates for the US Government Bond Index over a range of confidence levels. The VaR-x uses the modified Hill estimator for the tail index as the parameter in the Student-t distribution, the parametric-normal approach assumes normally distributed returns and the empirical approach uses the observed frequency distribution. The difference is given in million dollars for a position of $100 million in the particular asset.

both the parametric-normal and VaR-x methods. These forecasts, using daily returns, are plotted in Figure 4 against the actual rolling bi-weekly returns, of which some appear to look like multiple returns.

From Table 3, we can see that the parametric-normal approach provides VaR forecasts within a 99% confidence interval which are exceeded 1.99% of the time. Since a 99% confidence level requires the actual returns to exceed it only 1% of the time, the VaR-x forecasts have performed much better, with the forecast exceeded only 0.94% of the time. The consistently greater VaR-x forecasts provide evidence of the stability of α, and hence the VaR-x estimates. The VaR-x method using data on the government bonds showed a slight improvement, yielding 1.15% compared with 1.47% for the parametric-normal approach.

These results clearly show that the VaR-x method provides a more accurate estimate for the value-at-risk than the parametric-normal approach, illustrating the importance of including the tail fatness into the VaR estimate. Indeed, we see that by including a parameter for the distribution's fat-tailedness the estimate assesses the downside risk much more adequately than the assumption of normally distributed returns in the tails. The structure of the VaR-x parametric approach compared to that of the normal provides us with the phenomena of a larger value-at-risk as we move further into the tails. Indeed, as the normal distribution is nested in the Student t-distribution, as the alpha estimates become larger (less fat tailed), the VaR estimates will converge. This therefore provides us with a consistent parametric approach to modeling the
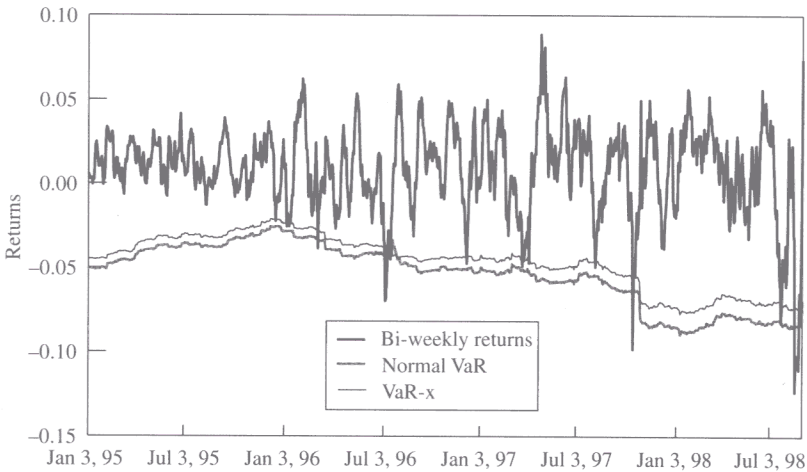

FIGURE 4. Rolling VaR-x and parametric-normal VaR estimates The graph shows how the forecasts of the VaR-x estimates, using the Student-t distribution, compare to forecasts from using the parametric-normal VaR approach for the S&P 500 Composite Return Index. We have used rolling observations of daily data, over the period January 1994 until August 1998 using 1216 rolling bi-weekly total returns, to provide forecasts of the value-at-risk at the 99% confidence level. The forecasts are based on yearly samples of daily data, and the alpha estimate is calculated for the left tail using a modified version of the Hill estimator.

TABLE 3. Exceedance statistics for rolling bi-weekly returns. This table contains the statistics on the S&P 500 Composite Return Index and the 10-Year Datastream US Benchmark Government Bond Index for the period January 1994 until August 1998 using 1216 rolling bi-weekly total returns. The forecasts are based on yearly samples of daily data (252 observations), and the alpha estimate is calculated for the left tail using a modified version of the Hill estimator.

	Exceedance of VaR at 99% confidence level		
	Theoretical	Normal VaR	VaR-x
Number of Exceedances			
S&P 500 Composite			
Return Index:	9.64	19	9
US 10-Year Government Bond			
Return Index:	9.64	14	11
Percentage of Exceedances			
S&P 500 Composite			
Return Index:	1.00%	1.990%	0.942%
US 10 Year Government Bond			
Return Index:	1.00%	1.466%	1.152%

additional downside risk associated with fat-tailed assets, which can easily be extended to allow for further conditionality in the data (see Koedijk and Pownall 1998).

5. CONCLUSIONS

It is widely known that the distributions of financial asset returns exhibit fatter tails than the normal distribution. This implies that the downside risk of a portfolio containing fat-tailed assets, as measured by value-at-risk, is underestimated when VaR is estimated with the assumption of normally distributed returns. Furthermore, this suggests that parametric conversion for different confidence levels as adhered to by the regulatory framework of the Basle Committee will provide inaccurate estimates of the VaR. It is no wonder that it has been necessary to 'ad hoc-ly' multiply the VaR by 3 to provide a larger, more representative, number for the Basle capital requirements. It would be preferable to have a more accurate measure reflecting the true risk from extreme returns, and the avoidance, or reduction at least, of the Basle multiplication factor. In this paper we present such a measure: VaR-x.

This methodology provides us with a simple approach to finding an accurate estimator for the VaR. The tail fatness apparent in financial returns is incorporated more accurately into the VaR-x estimator by using the Student t-distribution as a proxy for the distribution of future returns. We show that for both US stocks and bonds the VaR-x estimates reflect the true downside risk apparent in financial returns much better than those from the standard VaR estimators. The approach is easily extended to include further time-varying

parameters, and hence the implications for risk management seem tremendous. Certainly the move towards building portfolios which exploit these departures from normality (see Bekaert *et al.* 1998) will only serve to underline the vital importance of including an additional measure for the downside risk into the risk management techniques of the future.

REFERENCES

Bain, L. J., and Engelhardt M. (1987). *Introduction to Probability Theory and Mathematical Statistics.* Duxbury Press.

Bekaert, G., Erb, C., Harvey, C., Viskanta, T. (1998). Distributional characteristics of emerging market returns and asset allocation. *Journal of Portfolio Management*, forthcoming.

Bollerslev, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics*, **69**, 542–547.

Boothe, P., and Glassman, D. (1987). The statistical distribution of exchange rates: Empirical evidence and economic implication. *Journal of International Economics*, **22**, 297–320.

Danielsson, J., and de Vries, C. (1997). Value-at-risk and extreme returns. Working paper, London School of Economics.

Fama, E. F., and Roll, R. (1968). Some properties of symmetric stable distributions. *Journal of the American Statistical Association*, **63**, 817–846.

Hendriks, D. (1996). Evaluation of value-at-risk models using historical data. *Federal Reserve Bank of New York Economic Policy Review*, April.

Hill, B. (1975). A simple approach to inference about the tail of a distribution. *Annals of Mathematical Statistics*, **3**, 1163–1174.

Huisman, R., Koedijk, K. G., Kool, C., and Palm, F. (1997). Fat tails in small samples. Working Paper, Limburg Institute of Financial Economics, Maastricht University.

Huisman, R., Koedijk, K. G., Kool, C., and Palm, F. (1998). The fat-tailedness of FX returns. Working Paper, Limburg Institute of Financial Economics, Maastricht University.

Jansen, D., and de Vries, C. (1991). On the frequency of large stock returns: Putting booms and busts into perspective. *The Review of Economics and Statistics*, **73**, 18–24.

Jorion, P. (1996). Risk: Measuring the risk in value at risk. *Financial Analysts Journal*, November/December.

Jorion, P. (1997), *Value at Risk: The New Benchmark for Controlling Derivatives Risk*. McGraw-Hill.

Koedijk, K. G., and Pownall, R. A. (1998). VaR-x: Risk management in emerging markets. Working Paper, Limburg Institute of Financial Economics, Maastricht University.

Kupiec, P. H. (1995). Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives*, **2**, 73–84.

Loretan, M., and Phillips, P. C. B. (1994). Testing the covariance stationarity of heavy-tailed time series: An overview of the theory with applications to several financial data series. *Journal of Empirical Finance*, **1**, 211– 248.

Rogalski, R. J., and Vinso, J. D. (1978). Empirical properties of foreign exchange rates. *Journal of International Business Studies*, **9**, 69–79.

Taylor, S. (1986). *Modeling Financial Time Series*. Wiley.

R. Huisman
Erasmus University, Rotterdam

K. G. Koedijk and R. A. J. Pownall
University of Maastricht