VaR stress tests for highly non-linear portfolios

J.H.J. Einmahl, W.N. Foppen,
O.W. Laseroms, C.G. de Vries

VaR stress tests for highly non-linear portfolios

John H.J. Einmahl[†]
Walter N. Foppen[‡]
Olivier W. Laseroms[§]
Casper G. de Vries[¶]

April 2002

Abstract: Risk management techniques approximate derivative portfolios using a second order Taylor series. If one retains the option pricing formula combined with the semi-parametrically estimated tail transition probabilities, one captures the non-linearity of the portfolio risk, where the quadratic approximation fails.

[†]Tilburg University & EURANDOM
[‡]Rabobank International
[§]MathMetrics
[¶]Erasmus University Rotterdam & EURANDOM

1

1 Introduction

Banks regularly estimate the downside risk on their trading portfolios for the purpose of internal risk management and external supervision. The Delta-Gamma cum normal distribution approach, as advocated by e.g. the RiskMetrics product, is the industry standard for assessing the risks of a portfolio containing derivatives. This Taylor expansion based method tries to the capture the non-linear behavior of a portfolio containing derivatives by a quadratic approximation. This approximation is accurate in the vicinity of the current price of the underlying asset, which is the usual case in Value-at-Risk exercises. The approximation, however, can be very inaccurate when the price of the underlying asset drifts away from the current price. For Value-at-Risk exercises like stress testing, see the BIS (2000) report, large changes in the price of the underlying asset are important. We quote from the BIS report (2000): "...VaR has been found of limited use in measuring firms' exposures to extreme market events...". There are two problems: One problem is that the probability by which extreme market events do occur is not well captured by the normal distribution. The other problem is the high curvature of a derivative portfolio with respect to the underlying in extreme situations.

To overcome these problems, we present a method that uses closed form option pricing formulas, such as the normal distribution based Black-Scholes formula, in order to capture the high curvature at the edges, which may be completely missed if the quadratic approximation is used. In order to estimate the probability that the derivative portfolio will be in this area, we do however not use the assumption of normality. Instead we rely on a fit of the tail of the distribution of the underlying that is commensurate with the empirical distribution. Thus while the pricing uses standard normality based formulas, the transition probabilities for the underlying asset prices are different. This hybrid procedure may seem to fly in the face of theoretical consistency. But other research has shown that the Black-Scholes pricing formula is usually within the 95% confidence area, even if the underlying follows a non-normal stochastic process, see e.g. Mahieu and Schotman (1998). This may in part be due to the usage of the Black-Scholes pricing formula as the basic input for giving actual price quotes, modified in the tails to capture smiles and smirks. The transition probabilities of the underlying asset features the heavy tails found in practice.

2 The VaR of a portfolio of options on a single underlying asset

Representative for the standard way of obtaining VaR estimates is the RiskMetrics Group (1996) methodology. It assumes that the log-returns of the underlying stock are normally distributed with mean zero and a volatility that changes over time. On basis of this assumption and a quadratic approximation of the portfolio returns as a function of the returns on the underlying asset, the first four moments of the distribution of the portfolio returns can be estimated. Subsequently, a Johnson distribution is fitted on these moments. This distribution determines the VaR estimate.

Two observations regarding this fully parametric approach can be made. First, around the center there is little need to approximate the distribution of the returns by a specific parametric model, since the empirical distribution contains sufficiently many observations in this area. Second, in the tail area there is ample evidence that the normal model is not appropriate, see e.g. Campbell et al. (1997). Based on these two observations we propose instead a semi-parametric approach, as in e.g. Caserta et al. (1998), whereby the tail part is modeled semi-parametrically and the linearly interpolated empirical distribution function is used in the center. We assume in line with empirical evidence that the distribution function of asset returns is heavy tailed, i.e. exhibits power decline. The class of distributions with this property, formally the class of regularly varying distributions (see e.g. Danielsson et al. (2001)), exhibits to a first order a Pareto-type tail[1]:

$$
\begin{aligned}
P(X \le -x) &= x^{-\frac{1}{\gamma}} L(x),\ \gamma > 0, \\
P(X \ge x) &= x^{-\frac{1}{\widetilde{\gamma}}} \widetilde{L}(x),\ \widetilde{\gamma} > 0,
\end{aligned}
\tag{1}
$$

for $x > 0$. The estimate of the distribution function of returns thus consists of a trimmed empirical distribution to which we attach estimated heavy tails at both ends.

The tail parameters are estimated by means of Hill's method. The number of observations used to estimate the tail parameters is determined by the bootstrap method in Danielsson et al. (2001). This also determines where the empirical distribution function is trimmed, which ensures that our estimator of the distribution function will be continuous and hence increasing.

The portfolio value is determined via the (normal based) pricing formula as a function of the value S of the underlying asset. We apply the Black-Scholes formula to each option in the portfolio and denote the pricing function by V. If we want to estimate the VaR at level $1 - \alpha$, we need the value of x for which $P(V(S) < x) = \alpha$. This x can be estimated by solving the equation

$$
\int_0^\infty 1_{[V(s) < x]} d\hat{F}(s) = \alpha,
\tag{2}
$$

where $\hat{F}$ is the semi-parametric estimator of the distribution function of next period's asset price. The estimator $\hat{F}$ follows directly from the current asset price and the estimator of the distribution function of the returns. The estimate of the VaR then equals the current portfolio value minus x. The value of x can be found using a bisection method, since $P(V(S) < x)$ is increasing in x. To determine the above expression for given x, we need to find the prices of the underlying asset for which the portfolio has a value smaller than or equal to x. These prices can be found efficiently by determining the local extrema of $V(S)$. We note that between two local extrema the function is monotonous and thus $V(S) - x$ can only have one root. The local extrema are reached at those asset prices where the delta of the portfolio equals zero.

---

[1] Here L (and $\widetilde{L}$) represents a slowly varying function, i.e. a function satisfying $\lim_{x \to \infty} L(tx)/L(x) = 1$, $t > 0$.

3 An example

We compare the alternative procedures for estimating the VaR for a given example portfolio of derivatives. Consider the hypothetical portfolio where 15,000 plain vanilla put options with strike price 107.50 are written (price per option: 0.090), 10,000 plain vanilla call options with strike 100 (price per option: 12.628) and 40,000 with strike 112.50 are bought (price per option: 1.437) and 40,000 call options with strike 110 are written (price per option: 3.061). The time to maturity is 8 trading days for all options.

The current asset price equals 112.50, the volatility used to price the options is taken to be 17%. The risk-free interest rate is assumed to be 4%. The current value of the portfolio then equals 59,981. For estimation we use the returns of the Amsterdam AEX Index from January 1983 until December 1994. The 1491 data from January 1995 until November 2000 are used for out-of-sample comparison. Figure 1 shows the value of the portfolio as a function of the underlying stock and its approximation by the RiskMetrics methodology. As can be seen directly from the figure, the quadratic approximation performs badly in the (left) tail area and this has a considerable influence on the VaR estimates.
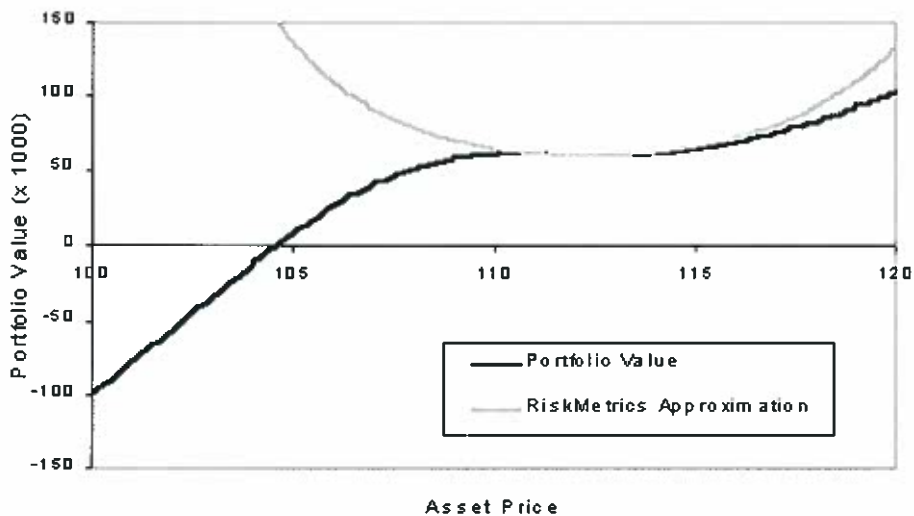

Asset Price

Figure 1 - Portfolio value as a function of the underlying asset price and RMG's approximation

The competing VaR estimates at various confidence levels are given in Table 1. The quadratic approximation (Quadratic) does not capture the behavior of the portfolio for low values of the underlying stock price and therefore these VaR estimates are out of bound. The estimates based on our non-linear Black-Scholes cum heavy tail method (Hybrid tail) do have the right order of magnitude.

Confidence Level	Hybrid tail	Quadratic	Out of sample loss
95.0%	703	559	702
99.0%	3,070	564	5,830
99.5%	12,041	564	15,737
99.75%	33,676	564	23,202

Table 1 - VaR estimates and out of sample results for various confidence levels

In Table 2 we give the number of exceedances of the VaR values that we counted in the out of sample dataset. E.g., for the 95% confidence level we should see approximately $0.05 \times 1491 \approx 75$ exceedances. We also provide 95% confidence bounds for the number of exceedances, based on the binomial distribution.

Confidence Level	Hybrid tail	Quadratic	Expected	Lower bound	Upper bound
95.0%	67	540	74.55	59	91
99.0%	21	534	14.91	8	23
99.5%	11	534	7.46	3	13
99.75%	1	534	3.73	1	8

Table 2 - Exceedances of the VaR estimates with expected values and confidence bounds

We see that the exceedances of our hybrid Black-Scholes cum heavy tail method estimates are between the bounds for all the tabulated confidence levels. The quadratic approximation cum normal transition probabilities method dramatically underestimates the frequency of high losses, leading to a large number of VaR exceedances in the out of sample exercise.

4 Portfolios on multiple underlying assets

In this section we briefly present the case where we have a portfolio of derivatives on several underlying assets. The quadratic approximation again gives rise to problems at the edges. We consider a portfolio on two different underlying assets and we include for each asset the same options as in the univariate example. Both underlying assets have the same spot price as in the univariate example, i.e. 112.50, and the corresponding options have the same strike prices as reported in de univariate example. Moreover, the same volatility and risk-free interest rate as in the univariate example are used to price the options. The value of the portfolio thus amounts to $2 \times 59,981 = 119,962$. In Figure 2 we see respectively the value of this portfolio as a function of the underlying asset prices and its quadratic approximation. Clearly the situation is similar to the univariate case.
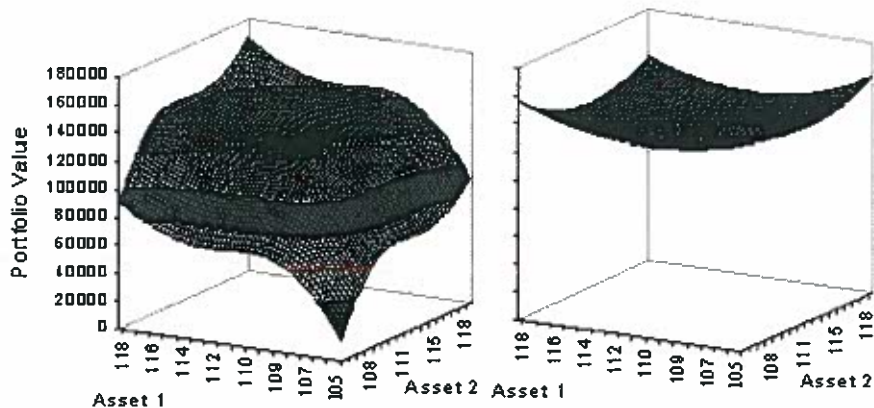
Figure 2 - Portfolio value as a function of the underlying asset prices (left) and RMG's approximation

We sketch how one can implement a complete pricing formula for the case of multiple underlying assets. Since the univariate method has no natural generalization, we propose a simulation based method. To this end we need an estimator of the multivariate distribution function of the underlying asset returns. Since the estimator of a univariate distribution function (see Section 2) performs well, we use it to estimate the univariate marginals of the multivariate distribution function. In order to model the dependence between the returns (as is observed in financial data) we make use of the multivariate normal copula (or dependence) function; see, e.g., Joe (1997) for a treatment of copulas and other dependence concepts. Since a copula function has uniform-(0,1) marginals, in the bivariate case the normal copula contains only the parameter ρ. The ρ has to be estimated from the data. We simulate returns for the assets by generating multivariate normal copula random vectors and transform these component-wise using the inverse of our univariate semi-parametric heavy tail distribution function estimator. A similar approach is also used in Hull and White (1998). With these simulated returns, we determine corresponding asset prices and the portfolio value. By sorting all portfolio values attained in this way we can estimate the VaR for the portfolio.

Table 3 gives our VaR estimates and the RMG's quadratic approximation VaR estimates for various confidence levels for the bivariate portfolio. For estimation we used the returns on the ING and Philips stock from March 1991 until September 1999. We notice that the estimates based on Historical Simulation and the Hybrid tail method are in the same order of magnitude for all confidence levels, whereas the VaR estimates resulting from RMG's method are out of bounds for all confidence levels. This is due to the fact that RMG's quadratic approximation is unable to capture the behavior of the portfolio at the edges.

Confidence level	Hybrid tail	Quadratic	Historical Simulation
95.0%	5,572	1,050	4,748
99.0%	63,591	1,179	67,971
99.5%	127,096	1,197	111,550
99.75%	193,509	1,207	156,959

Table 3 - VaR estimates at different confidence levels in the multivariate example

5 Conclusions

In summary, a quadratic approximation in estimating the Value-at-Risk can be quite misleading, if the current portfolio value is trapped in a local minimum. Moreover, the normal based transition probabilities underestimate the tail risk. Therefore it is better to use a complete pricing method, with adjusted transition probabilities, to take account of market incompleteness and observed pricing practices as well as the actual risks in the underlying values.

References

[1] Bank for International Settlements (April 2000). Stress testing by large financial institutions: current practice and aggregation issues, Committee on the Global Financial System, Basel.

[2] Campbell, J., A. Lo and C. MacKinley (1997). The econometrics of financial markets, Princeton University Press, Princeton.

[3] Caserta, S., J. Danielsson and C.G. de Vries (1998). Abnormal returns, risk, and options in large data sets, Statistica Neerlandica 52, 324-335.

[4] Danielsson, J., L. de Haan, L. Peng and C.G. de Vries (2001). Using a bootstrap method to choose the sample fraction in tail index estimation, Journal of Multivariate Analysis 76, 226-248.

[5] Hull, J.C. and A. White (1998). Value at Risk when daily changes in market variables are not normally distributed, Journal of Derivatives 5, 9-19.

[6] Joe, H. (1997). Multivariate models and dependence concepts. Monographs on Statistics and Applied Probability, 73. Chapman & Hall, London.

[7] Mahieu, P.J. and P.C. Schotman (1998). An empirical application of stochastic volatility models, Journal of Applied Econometrics 13, 333-360.

[8] RiskMetrics Group (1996). RiskMetrics - Technical Document, Fourth Edition, New York.