

A Coupling of Extreme-Value Theory and Volatility Updating with Value-at-Risk Estimation in Emerging Markets: A South African Test

Anthony J. Seymour

University of Cape Town, South Africa

Daniel A. Polakow

*University of Cape Town and
Cadiz Holdings, South Africa*

This research is aimed at a formal appraisal of recent advancements in stochastic volatility modeling and extreme-value theory to application of value-at-risk computation in particularly volatile markets. Established methods such as historical simulation are prone to underestimating value-at-risk in such developing markets. Two contemporary methods of value-at-risk calculation are tested on a representative portfolio of South African stocks. The first method incorporates extreme value theory. The second model includes both extreme value theory and volatility updating (via GARCH-type modeling). The combined GARCH-type time-series approach and extreme value theory model is found to provide significantly better results than both straightforward historical simulation as well as the extreme value model. In no instance, however, were results on these VaR methods as good as those obtained when the same methods were tested in developed markets. This research highlights noteworthy improvements to value-at-risk estimation efficacy in volatile emerging markets, and also stresses the need for further work into the estimation of value-at-risk in this context (JEL D81,G10).

Keywords: backtesting, extreme value theory, GARCH, historical simulation, RiskMetrics, value-at-risk.

I. Introduction

Value-at-risk (VaR) has emerged as one of the better known and utilized risk measures in the financial industry. Its utility lies in a simplicity of interpretation – VaR is able to convey the risks to which an institutional portfolio or treasuries book is exposed in an intuitive manner. Although a single, robust and unambiguous definition of VaR may be given (for

(Multinational Finance Journal, 2003, vol. 7, no. 1 & 2, pp. 3–23)

©*Multinational Finance Society*, a nonprofit corporation. All rights reserved.

DOI: 10.17578/7-1/2-1

example, see Embrechts, Kluppelburg and Mikosch [1997]) there remain several methods of VaR calculation, all of which have the potential to afford vastly different VaR estimates. Implementation of any VaR system is coupled with the dilemma of deciding which of the available methods is the most appropriate. The most important consideration is the ability of the method to reflect accurately the probability of loss. For a methodology to be considered sound, the number of violations of estimates provided at various confidence levels should be consistent with those levels. For example, the estimates provided at the 95% confidence level should only be violated on approximately 5% of occasions. At higher confidence levels (say 99.99%), analysts require an estimate of extreme-risk while standard parametric maximum-likelihood techniques are ill-suited to model fitting in the extreme tails-regions of probability distributions. To this end, advancements in extreme value theory (EVT) were a welcome feature of proposed VaR methods in the 1990's (Danielsson and De Vries [2000], Embrechts, Resnick and Samorodnitsky [1998], Embrechts, Resnick and Samorodnitsky [1999], Login [1997], McNeil [1998], McNeil and Frey [2000]).

A large body of literature exists describing and examining various VaR algorithms (Jorion [1996]). Certain methods are known to afford better VaR estimates at higher confidence levels (for example, see Mahoney [1996]), while other methods have recently proposed taking into account the non-stationary nature of volatility in VaR calculations (for example, see J.P. Morgan [1996]). It is worthwhile noting, however, that most of these studies have been tested on developed markets such as those in the US and Europe without consideration for the volatility peculiarities of emerging markets.

In this paper we argue that studies based on developed markets provide little basis for the selection of a VaR estimation method in anomalously volatile emerging markets. Furthermore, we argue that any method likely to succeed at VaR estimation in such markets should accommodate features that better take into account the idiosyncratic nature of emerging markets. We present empirical evidence of these claims in the form of an empirical study on the South African market.

This paper is structured as follows: Section II provides an overview of the methods that were tested, highlighting the advantages and disadvantages of each. Section III provides more detail on the EVT based methods that were tested. Section IV describes in detail the empirical study performed using data from the South African stock

market. The final section (section V) presents a short synthesis of the findings and suggestions of directions for further research.

II. Outline of Tested VaR Estimation Procedures

In order to understand the various methods of VaR calculation, it is necessary to examine a mathematical definition of VaR. If X is defined to be the loss on a portfolio, and p the confidence level, then VaR is simply defined by:

$$P(X \leq VaR) = \pi.$$

This is equivalent to:

$$F_X(VaR) = \pi,$$

where $F_x = P(X \leq x)$ is the probability distribution function of the random variable X . It is thus clear that if the probability distribution function of X is known, then VaR is given by:

$$F_X^{-1}(\pi) = VaR,$$

where $F_x^{-1}(\bullet)$ denotes the inverse of $F_x(\bullet)$. Note that these mathematical interpretations of VaR are consistent with the definition given in words by Hull (1999): ‘We are p percent certain that we will not lose more than V dollars in the next N days’. Here, V is the VaR of the portfolio. This definition is unambiguous in that it demonstrates that VaR is a function of p , the confidence level, and N , the time horizon.

Thus, although there are widely varying approaches to calculating VaR, all methods have the common goal of estimating the profit and loss distribution function of the portfolio under consideration. This is because VaR is simply a quantile of this distribution. All methods will therefore be described in terms of the way in which the profit and loss distribution function is estimated.

The methods that were examined were:

- a) Historical simulation
- b) RiskMetrics™
- c) EVT-based extensions of the above methods

A. Historical Simulation

The method of historical simulation is the simplest method of obtaining a profit and loss distribution of a portfolio. The method requires a database of returns for the stocks comprising the portfolio, and uses these to generate a set of simulated returns for the portfolio. The returns can be used to obtain an empirical distribution function for the losses on the portfolio by converting the returns to losses (simply via a change of sign), and sorting the losses into an ordered set $X_1 \leq \dots \leq X_N$. The empirical probability distribution for the losses on the portfolio is defined by:

$$\mathbb{P}(X \leq x) = \begin{cases} 0 & x < X_1 \\ i/N & X_i \leq x \leq X_{i+1} \\ 1 & X_N \leq x \end{cases}$$

Thus for a confidence level of $p = 0.95$, and a data set of 100 observations, VaR would be given as the 95th loss in the ordered set. The main disadvantage to this approach is that it is impossible to obtain an out-of-sample VaR estimate. In other words, the smallest probability that can be obtained such that $\mathbb{P}(X \leq \text{VaR})$ is $1/N$. Thus for a sample size of 100, it would not be possible to estimate VaR at a confidence level higher than 99%. The advantage of the historical simulation approach is that it accurately reflects the historical probability distribution of the market variables. However, since relatively few data points enter the tail of the distribution, VaR estimates obtained from these observations are likely to be inaccurate.

B. RiskMetricsTM

The details of the RiskMetricsTM methodology can be found in the RiskMetricsTM technical document, published by J.P Morgan (1996). In its most basic form, the RiskMetricsTM methodology involves specifying the type of distribution for the profit and loss variables, and using analytical techniques to calculate the volatility (standard deviation σ) of the portfolio. Generally, the profits/losses on a portfolio are assumed to be normally distributed with a mean of zero, thus allowing the probability distribution function to be fully specified once the volatility is known. The value of σ used is a one-step ahead predicted value obtained via an exponentially weighted moving average (EWMA) volatility model.

The most common criticism against this type of approach is that the assumption that the profits/losses on a portfolio are normally distributed is unrealistic. It is known that financial return series exhibit leptokurtosis, or 'heavier tails' than a normal distribution. In essence, this means that any VaR calculation technique based on a normal distribution function will tend to give VaR estimates that are too low. The advantage of the RiskMetrics™ methodology is that the current volatility background is captured via the EWMA model allowing VaR estimates to take into account changing volatility, a phenomenon that is well documented in financial time-series.

C. Extreme Value Theory Based Methods

The RiskMetrics™ approach to VaR estimation requires the assumption of a normal distribution for portfolio returns, which can lead to a serious underestimation of VaR, owing to the 'heavy tailed' nature of financial return series. A possible alternative to this method is to fit a fat tailed distribution to all the data. The drawback of this type of strategy is that the distribution is constructed to fit the central, common observations best, and is therefore ill-suited to the extreme observations with which VaR estimation is concerned. The problem facing anyone attempting a VaR analysis is one of estimating the probabilities of extreme or rare events with limited data. Dowd (1999) points out that this type of problem is not unique to risk management, occurring in other fields such as hydrology where engineers have had to determine how high sea walls should be to contain flood probabilities within reasonable limits. The branch of statistics that has arisen out of a study of these types of problems is extreme value theory (EVT). In essence, EVT allows one to determine the nature of the tail of a distribution without the need for strong assumptions concerning the distribution from which the data are drawn. A comprehensive treatment of extreme value theory is given by Embrechts et al. (1997).

The topic of the estimation of tails of the profit and loss distribution functions via EVT has been the subject of extensive recent research. Dowd (1999) attributes the first application of EVT to the problem of VaR estimation to Longin (1997). Since EVT based methods address the shortcomings inherent in both historical simulation and the RiskMetrics™ methodology, it was decided to test the applicability of modern EVT methods to the calculation of VaR in the South African market, and compare these results with those of the established

methods. Of the published work, the papers by Danielsson and De Vries (2000), and McNeil and Frey (2000) were selected for further analysis. These methods were chosen based on the fact that they can be viewed as extensions of established methodologies, thus providing scope for the comparison of 'traditional' VaR estimation techniques with an EVT based approach. We discuss each of these methods separately, together with documentation of their findings within an international VaR setting.

In evaluating the effectiveness of a particular VaR estimation procedure, a decision must be made regarding the criterion for success or failure of the method. The approach adopted in this study is based on the evaluation methods of Danielsson and De Vries (2000), and McNeil and Frey (2000). This is primarily to facilitate a direct comparison between the results obtained in this study, and those obtained by the aforementioned authors. These authors examine the frequency of violations of VaR estimates at various confidence levels, which is in essence the method advocated by the Basle Committee on Banking Supervision (1996). An overview of alternative methods is given by Raatikainen (2002).

III. Description of Tested EVT Based Methods

The method of Danielsson and De Vries (2000) (denoted D&DV) can be viewed as an extension of historical simulation. The approach essentially involves obtaining an empirical distribution function via historical simulation and using EVT to fit a smooth curve through the tail of this distribution. It thus appears to be an attractive method for implementation in South Africa, as historical simulation is already in use by many of the larger South African institutions and the EVT estimator can be applied without much extra effort.

A noteworthy aspect of D&DV approach is that it does not take into account the stochastic nature of volatility. The reason for adopting the D&DV method is based on empirical observations of extreme returns. In particular, D&DV impute that extremes (or spikes) exhibit no volatility clustering, and no strong correlation. D&DV argue that since extreme returns occur infrequently, and do not appear to be related to a particular level of volatility, an unconditional approach is better suited to VaR estimation than conditional volatility forecasts.

To demonstrate the effectiveness of their method, D&DV use data

TABLE 1. Violations of D&DV VaR Estimator

Confidence Level	Expected	Method		
		Risk Metrics™	Historical Simulation	D&DV
0.95000	50.00	52.45 (7.39)	43.24 (10.75)	43.14 (11.10)
0.97500	25.00	30.26 (4.41)	20.50 (7.22)	20.84 (7.35)
0.99000	10.00	16.28 (3.13)	7.66 (3.90)	8.19 (3.86)
0.99500	5.00	10.65 (2.73)	3.69 (2.39)	4.23 (2.55)
0.99725	2.50	7.29 (2.27)	1.90 (1.57)	2.35 (1.72)
0.99900	1.00	4.85 (2.06)	0.95 (1.03)	1.06 (1.13)
0.99950	0.50	3.55 (1.81)	0.75 (0.89)	0.59 (0.82)
0.999725	0.25	2.72 (1.66)	0.75 (0.89)	0.33 (0.62)
0.99990	0.10	2.00 (1.45)	0.75 (0.89)	0.12 (0.35)
0.99995	0.05	1.58 (1.29)	0.75 (0.89)	0.06 (0.23)

Note: The number of violations of the VaR estimates are given as averages of the violations for the 500 portfolios with standard errors in parenthesis. In addition, the average number of violations of RiskMetrics™ and historical VaR estimates at various confidence levels are presented. These violations can be compared with the expected number of violations for each confidence level.

on 6 randomly selected US stocks in addition to the J.P. Morgan bank stock price as the basis for portfolio analysis. The data set for each VaR estimation was set to 1500 trading days. For each of these portfolios, VaR was calculated by applying 500 random portfolio weights to 1500 days of returns for the individual stocks to obtain 1500 days of returns for the portfolios, as in the method of historical simulation, and using the EVT-derived estimator. The VaR estimate was compared to the realized portfolio return on the following day, and this procedure was repeated for 1000 different consecutive days. The total number of trading days required for this back-test was thus 2500. The results of D&DV back-test on six randomly selected US stocks in addition to J. P. Morgan are shown in table 1. Table 1 also displays the results of

backtests of the basic RiskMetrics™ methodology and historical simulation.

RiskMetrics™ is seen to have provided the most accurate VaR estimates at the lowest confidence level, but consistently under-predicted VaR at subsequent levels. Historical simulation performed well until its probability limit (1/1500), and the EVT-based estimator is impressive in its agreement with the expected number of violations, especially at the higher confidence levels. Based on these results, one may conclude that for the data set employed the D&DV estimator seems appropriate.

The method of McNeil and Frey (2000) (denoted M&F) can be viewed as an extension of the basic RiskMetrics™ methodology. RiskMetrics™ assumes that the returns on a financial asset X_t follow a process of the form:

$$X_t = \sigma_t Z_t,$$

here the Z_t are assumed to be independent and normally distributed with mean zero and unit variance, and σ_t is a one step ahead predicted volatility, estimated via an EWMA model. M&F specify a similar model of the returns process, given by:

$$X_t = \phi X_{t-1} + \sigma_t Z_t,$$

where σ_t is estimated via a generalized autoregressive conditional heteroskedastic (GARCH) model, first proposed by Bollerslev (1986). The essential difference, however, between the method of M&F and RiskMetrics™ lies in the specification of the distribution function of the noise process $F_z(z)$. Instead of assuming a normal distribution for Z_t , M&F use extreme value theory to estimate the tail of $F_z(z)$, and hence VaR. A noteworthy feature of the method of M&F is that it is in contrast to the method of D&DV who ignore the stochastic nature of volatility.

M&F tested their method using a similar back-testing technique to that used by D&DV. The size of the data set used to fit the GARCH and AR(1) models was 1000 days, and the method was tested on individual assets (in contrast to D&DV test on portfolios of assets). The return series of the following assets were examined: S&P, DAX, BMW, USDGBP, and gold. VaR was calculated at the 95, 99, and 99.5% confidence levels for the purposes of the back-test. The VaR estimation techniques that were tested were the conditional and unconditional

TABLE 2. M&F Back-Testing Results

	S&P	DAX	BMW	USDGBP	Gold
Length of Test	7414	5146	5146	3274	3414
0.95 Quantile					
Expected	371	257	257	164	171
Conditional EVT	366	258	261	151	155
Conditional Normal	384	238	210	169	122
Conditional t	404	253	245	186	168
Unconditional EVT	402	266	251	156	131
0.99 Quantile					
Expected	74	51	51	33	34
Conditional EVT	73	55	48	35	25
Conditional Normal	104	74	86	56	43
Conditional t	78	61	52	40	29
Unconditional EVT	86	59	55	35	25
0.995 Quantile					
Expected	37	26	26	16	17
Conditional EVT	43	24	29	21	18
Conditional Normal	63	44	57	41	33
Conditional t	45	32	18	21	20
Unconditional EVT	50	36	31	21	11

(D&DV) EVT methods, as well as methods in which the white noise process is assumed to have a particular distribution, i.e. normal or Student's t (the normal case is similar to the RiskMetricsTM method, while the latter can be regarded as the fitting of a heavy tailed distribution to the data). The results of their back-test are summarized in table 2.

In the analysis of their results, M&F point out that in 11 out of 15 cases, their approach is closest to the mark. The conditional normal approach was found to fail 11 times, and unconditional EVT, 3 times. The unconditional EVT estimate cannot respond quickly to changing volatility and tends to be violated several times in a row in stress periods. It is apparent from these results that approaches to tail estimation which ignore the conditional heteroskedasticity exhibited by most financial return series are unlikely to be suitable for VaR calculation.

TABLE 3. Stocks Selected to Form the Basis of Portfolio Analysis

Code	Short Name	Sector
AMS	AMPLATS	Resources
BOE	BOE	Financial
DDT	DIDATA	Industrial
INT	INVSTEC	Financial
LGL	LIBERTY	Financial
NED	NEDCOR	Financial
RCH	RICHEMONT	Industrial
SBC	SBIC	Financial
SOL	SASOL	Resources

IV. Application of VaR Calculation Methods to the South African Market

We now turn our attention towards the applicability of these VaR methods to the South African market. The published studies noted above demonstrate that the established methodologies perform well at certain confidence levels when tested on return series generated by US and European financial assets, but this is obviously no indication of applicability to other markets. It is the central objective of this study to test established calculation methods on return series generated by a representative portfolio of South African financial stocks, and to determine if the modern EVT-based procedures offer any improvement.

A collection of nine stocks was chosen to form the basis of portfolio analysis based on the requirement for a historical data series of each stock (of at least 10 years), and then on the representative nature of the stocks given the constraints of an adequate history. The selection is displayed in table 3, and forms an illustrative portfolio in the general equity category of the South African Unit Trusts (i.e. commonly 80% of top and average performing unit trusts would have 75% of these stocks).

The following sections describe the backtests of the methods of historical simulation, RiskMetricsTM, and the EVT-based methods of D&DV and M&F.

A. Historical Simulation

The back-test of the historical simulation method of VaR calculation was carried out as in D&DV. The daily returns for the 9 stocks

TABLE 4. Average Number of Violations of Historical Simulation Estimates

Confidence Level	Expected	Observed
0.95000	49.35	119.28 (5.95)
0.97500	24.68	70.52 (5.79)
0.99000	9.87	30.89 (2.47)
0.99500	4.94	20.07 (1.67)
0.99750	2.47	9.48 (1.02)
0.99900	0.99	4.73 (0.93)
0.99950	0.49	2.96 (0.21)

Note: Standard errors of observations are indicated in parenthesis.

comprising the hypothetical portfolio were arranged into a 2486×9 matrix. A random 9×500 matrix, whose columns summed to one, was generated, which represented 500 random portfolio weights. Multiplication of these two matrices afforded a 2486×500 matrix, which represented 2486 days of returns for 500 different portfolios. This matrix formed the core of the analysis. Beginning at day 1500, for each portfolio, the previous 1500 days returns were used to compute VaR estimates for various confidence levels. These were compared to the realised return on the following day. This was repeated 986 times. The results of the analysis are summarized in table 4.

The results demonstrate how inadequate the method of historical simulation is with the VaR estimates being violated approximately 3 times as often as the expected values. When tested in developed markets (e.g. US, Europe), historical simulation has been shown to offer reasonable VaR estimates for at least the 95% confidence level. That is clearly not the case here, and affirms that the South African market has its own volatility peculiarities that need to be taken into account. It also shows that companies allocating risk capital based on historical-simulation based estimates are taking inadequate steps to protect themselves against extreme market risk.

TABLE 5. Average Number of Violations of D&DV Simulation Estimates

Confidence Level	Expected	Observed
0.95000	49.35	127.31 (8.18)
0.97500	24.68	72.98 (6.55)
0.99000	9.87	29.13 (2.35)
0.99500	4.94	14.85 (1.94)
0.99750	2.47	6.87 (1.43)
0.99900	0.99	2.68 (0.59)
0.99950	0.49	1.30 (0.53)
0.999725	0.25	0.92 (0.27)
0.99990	0.10	0.08 (0.27)
0.999995	0.05	0.00 (0.00)

Note: Standard errors of observations are indicated in parenthesis.

B. Danielsson and De Vries Tail Estimator

The backtest of D&DV method was carried out as in the historical simulation method, except that VaR at each step was calculated using the tail estimator proposed by D&DV. The tail observations were set to consist of the highest 10% of observed losses as proposed by M&F. The results of the backtest are displayed in Table 5.

It is clear that up to the 99% confidence level, the number of violations of the D&DV VaR estimates, and that of historical simulation are essentially the same. This is not surprising, since the method of D&DV involves fitting a smooth curve through the tail of the empirical distribution function, and thus in a region where there is a higher concentration of empirical observations, the two methods are likely to have results that coincide. One would expect the two methods to diverge where data is sparse – the domain of interest in EVT. This is exactly what is observed at confidence levels higher than 99%.

In order to assess the utility of this method, some additional analysis

of the data was undertaken. Since the back-test involved calculating VaR approximately 1000 times for 500 different portfolios (i.e. 500 000 VaR calculations), it was clearly impossible to analyze the data in every single VaR calculation. It was decided to select a single vector of portfolio weights, and to examine the return series of this portfolio in two different 1500 trading day windows, labeled a and b in figures 1 and 2 below. The windows of portfolio returns were obtained by generating 2500 days of returns via the multiplication of the vector of portfolio weights with the matrix of stock returns, and selecting two starting points within this set of data from which to select 1500 days of data. In particular, Window 'a' was chosen to begin at return number 70 (i.e. 19 August 1991) in the total of 2500, while Window 'b' was chosen to begin at return number 500 (i.e. 5 May 1993). The selection of the two windows was arbitrary, although they do exhibit contrasting behavior in terms of volatility. The return series from Window 'a' shows little change in volatility, whereas the series from Window 'b' displays large jumps, or spikes, corresponding to a period of high volatility. The purpose of this additional data analysis was simply to give an indication of the mechanics of the method, and not to provide for any statistical inference. The two return series are shown in figures 1 and 2 respectively.

Since the essence of D&DV method is the fitting of a smooth curve through the tail of the empirical distribution function, it was decided to plot the fitted curves for the Windows 'a' and 'b' to obtain a sense of how well D&DV tail estimator actually approximates the tails. Figures 3 and 4 show the tail of the empirical distribution function for the two windows, along with fitted tails. In the case of Window 'b', the fit appears to be satisfactory, but less so in the case of Window 'a'.

The final aspect of D&DV method that was examined was the actual VaR estimates that were produced. Table 6 shows the VaR estimates that were calculated using the data from Windows 'a' and 'b'. Also included are VaR estimates calculated using the same data, but via the method of historical simulation.

Up to the 99% confidence level, historical simulation and D&DV method yield similar VaR estimates. At higher confidence levels they diverge. The most noteworthy feature of Table 6 is the very high VaR estimates produced by D&DV method, particularly in the case of Window 'b'. This, coupled with the fact that not once was the highest confidence level VaR estimate ever violated, suggests that D&DV method may be prone to overestimating VaR in the extreme tail region.




FIGURE 1.—Return Series Derived From Window ‘a’




FIGURE 2.—Return Series Derived From Window ‘b’




FIGURE 3.—Empirical Distribution Function With Fitted Tail (Window ‘a’)




FIGURE 4.—Empirical Distribution Function With Fitted Tail (Window ‘b’)

TABLE 6. VaR Estimates Produced by the D&DV Tail Estimates

Confidence Level	Window ‘a’		Window ‘b’	
	Historical Simulation	D&DV	Historical Simulation	D&DV
0.95000	0.0109	0.0098	0.017	0.0155
0.97500	0.0137	0.0137	0.0234	0.0238
0.99000	0.0203	0.0214	0.0397	0.0418
0.99500	0.0234	0.0299	0.0541	0.064
0.99750	0.0254	0.0418	0.0687	0.0981
0.99900	0.0298	0.0651	0.0873	0.1724
0.99950	0.0315	0.091	0.1358	0.2642
0.999725	–	0.1272	–	0.4047
0.99990	–	0.1981	–	0.7113
0.999955	–	0.277	–	1.0897

Note: The VaR estimates of the historical simulation (HS) technique are noted for the same two windows. The interpretation of the units of VaR is, using for example the first VaR figure in the table below (0.0109), that there is a 95% chance that the loss of a portfolio will not exceed 1.09% in a day.

On the whole, the VaR estimation procedure does not produce consistent results – VaR is underestimated at the lowest confidence levels, while overestimated at the highest. Thus not too much emphasis can be placed on any reasonable performance in between these two extremes (e.g. at the 99.99% level).

C. McNeil and Frey's Tail Estimator

The method used to backtest M&F's VaR estimation procedure was the same as that used in the historical simulation case, and in the test of the D&DV estimator to ensure consistency throughout the report. This back-test approach differs to that adopted by M&F in that M&F use a data window of 1000 trading days to estimate VaR and in addition, only consider return series generated by single assets, as opposed to returns on a portfolio. M&F state that the method obviously also applies to the time series of profits and losses generated by portfolios of financial instruments and can therefore be used for the estimation of market risk measures in a portfolio context. In a separate paper, McNeil (1999) gives a brief overview of multivariate extreme value theory (MEVT), and outlines a simple bivariate POT model. It is pointed out that parametric models of this kind are viable only in a small number of dimensions. In higher dimensional space, there are simply too many parameters to estimate and in such situations collapsing the question to a univariate problem by considering a whole portfolio of assets as a single risk and collecting data on a portfolio level seems more realistic.

Each VaR calculation involves the estimation of GARCH(1,1) parameters, as well as determination of parameters for the fitted Generalized Pareto Distribution (GPD). With GARCH and GPD parameters in hand, VaR values for various confidence levels were calculated based on the formulae described by M&F. Also calculated were VaR estimates based on the assumption of the white noise process Z_t having a standard normal distribution. The reason for this was to allow a back-test of a RiskMetrics™ type approach to VaR estimation. The results of the back-test are displayed in table 7.

The results clearly show the benefit of incorporating a volatility updating scheme into the calculation of VaR. At the 95% confidence level, the GARCH-based methods outperform the other two, with the RiskMetrics™-type approach providing the best results at this level. However, at higher confidence levels, the conditional normality assumption leads to serious under-predictions of VaR. As in the case of the analysis of D&DV method, further data analysis was undertaken to gain insight into the mechanics of M&F's method and to assist in the interpretation of the back-test results. The same two windows (see figures 5 and 6) that were examined in the discussion of D&DV method were used here, to allow a direct comparison.

Similar to the case of the VaR estimator of D&DV, an important

TABLE 7. Average Number of Violations of M&F’s VaR Estimates

Confidence Level	Expected	M&F’s EVT-GARCH Observed	Normal Observed
0.95000	49.35	77.91 (4.43)	69.77 (5.07)
0.97500	24.68	42.708 (3.24)	47.02 (3.24)
0.99000	9.87	21.55 (2.20)	31.63 (2.64)
0.99500	4.94	12.156 (1.95)	24.91 (2.29)
0.99750	2.47	6.394 (1.35)	20.16 (2.19)
0.99900	0.99	2.778 (0.99)	15.55 (2.36)
0.99950	0.49	1.79 (0.74)	12.66 (2.31)
0.999725	0.25	1.172 (0.67)	10.32 (2.29)
0.99990	0.10	0.648 (0.60)	7.97 (2.12)
0.999955	0.05	0.446 (0.55)	6.58 (1.93)

Note: Included are the VaR estimates for normal GARCH methods. Standard Errors of observations are indicated in parenthesis.




FIGURE 5.—Empirical Excess Distribution Function and Fitted GPD (Window ‘a’)




FIGURE 6.—Empirical Excess Distribution Function and Fitted GPD (Window ‘b’)

step in M&F’s method involves the fitting of a smooth curve to a distribution function of extreme observations and thus it is important to examine how good the fit is. Figures 5 and 6 expose how the estimated GPD curve is a good fit to the empirical distribution function in both cases.

The final aspect of M&F’s method that was examined was the actual VaR estimates themselves. Table 8 shows VaR estimates for various confidence levels using the data from Windows ‘a’ and ‘b’. Also included are the estimates provided by historical simulation and the method of D&DV, for comparative purposes. All three methods provide similar estimates at the lower confidence levels. At the highest levels, the large discrepancy between the values predicted by the D&DV estimator and that of M&F is evident.

Based on the backtest results and the magnitudes of the VaR estimates provided by the methods, the method of M&F performed the best. It is important to note, however, that while the method of M&F afforded the best results, the number of violations of the VaR estimates was higher than the expected number of violations at all confidence levels. This is in contrast to the results of the backtests performed by both D&DV and M&F who reported very good agreement between the expected number of violations, and the realised number of violations. This is indicative of the main point of this research – methods with proven success in foreign, developed markets are not necessarily appropriate for all markets. As far as finding an improved method specific to South Africa is concerned, the results obtained in this study

TABLE 8. VaR Estimates Produced by the M&F and D&DV Tail Estimators

Confidence Level	Historical Simulation	D&DV	M&F	Historical Simulation	D&DV	M&F
0.95000	0.0109	0.0098	0.0095	0.0170	0.0155	0.0177
0.97500	0.0137	0.0137	0.0126	0.0234	0.0238	0.0235
0.99000	0.0203	0.0214	0.0168	0.0397	0.0418	0.0321
0.99500	0.0234	0.0299	0.0200	0.0541	0.0640	0.0392
0.99750	0.0254	0.0418	0.0231	0.0687	0.0981	0.0470
0.99900	0.0298	0.0651	0.0272	0.0873	0.1724	0.0584
0.99950	0.0315	0.0910	0.0304	0.1358	0.2642	0.0679
0.99975	–	0.1272	0.0335	–	0.4047	0.0783
0.99990	–	0.1981	0.0376	–	0.7113	0.0934
0.99995	–	0.2770	0.0407	–	1.0897	0.1060

Note: The interpretation of the units of VaR is, using for example the first VaR figure in the table below (0.0109), that there is a 95% chance that the loss of a portfolio will not exceed 1.09% in a day.

suggest that historical simulation and the associated ‘tail smoothing’ approach are not of great utility. Methods incorporating volatility updating show promise, but the conditional normality assumption of the basic RiskMetrics™ approach leads to a serious underestimation of VaR at higher confidence levels. The results therefore also indicate that any attempts to obtain an improved VaR estimation method should have an approach similar to that of M&F as a starting point. In this study, a very brief investigation was made into a possible modification of M&F’s method with the aim of obtaining a better agreement between the expected and realised number of violations of the VaR estimates. The two main features of M&F’s method are the fitting of a GPD to the tail of the white noise distribution function, and obtaining a one-step ahead volatility forecast via a GARCH model. Any attempts to modify the method could therefore focus on one, or both of these aspects.

V. Concluding Remarks

The results in this study indicate that methods shown to afford accurate VaR estimates in developed markets do not necessarily have global application. It is clear that emerging markets such as that of South Africa have unique characteristics that need to be considered when implementing a VaR calculation procedure. The fact that methods

incorporating a volatility updating scheme afforded the most accurate backtesting results suggests that volatility is one of the important factors to be considered. In addition, the relatively good performance of the EVT-based methods implies that the importance of considering the probabilistic properties of extreme or rare events cannot be overemphasized. While the methods suggested as appropriate for South Africa are technically complex, we are sure that individuals who are reliant on rigorous VaR results in developing markets outside of South Africa are likely to benefit from their use.

Lastly, it has been admitted that while the results of the McNeil and Frey backtest are far more reliable than that of other tested approaches, the results do suggest a small shortfall between desired and empirical estimates of loss. What possible refinements could be effected to address this issue?

It is apparent that the GPD curve appeared to be a very good fit to the data and suggests one possible avenue for further research. It is likely that a source of improvement would be the use of an alternative GARCH model to provide the one-step ahead volatility forecast. One obvious possible refinement is asymmetrical GARCH models (*sensu* Levy 2001). Currently, we in the process of conducting research into the applicability of utilising asymmetrical GARCH models in the South African context. The results of this research will form the basis of a short companion paper in the near future.

References

- Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31: 307-27.
- Basle Committee on Banking Supervision. 1996. Amendment to capital accord to incorporate market Risks. *Basle Committee Publications*, no. 24, January 1996, Bank for International Settlements.
- Danielsson, J., and De Vries, C. 2000. Value at risk and extreme returns. *Annales d' Economie et de Statistique*, 60 (October – December), 239 - 270
- Dowd, K. 1999. The extreme value approach to VaR: An introduction. *Financial Engineering News*:11.
- Embrechts, P.; Kluppelburg, C.; and Mikosch, T. 1997. Modelling extremal events for Insurance and finance. Springer, Berlin.
- Embrechts, P.; Resnick, S.; and Samorodnitsky, G. 1998. Living on the edge. *Risk* 11(1): 96-100
- Embrechts, P.; Resnick, S.; and Samorodnitsky, G. 1999. Extreme value theory

- as a risk management tool. *North American Actuarial Journal* 3: 30-41.
- Hull, J.C. 1999. Options, futures, and other derivatives. 4th Ed., Prentice Hall.
- Jorion, P. 1996. Value at risk: The new benchmark for controlling market risk. Chicago: Irwin Professional Publishing.
- Levy, G. 2001. An introduction to GARCH models in finance. *Financial Engineering News*: 22.
- Longin, F.M. 1997. From value-at-risk to stress testing: The extreme value approach. CERESSEC. Working paper 97-004.
- Mahoney, J.M. 1996. Forecast biases in value-at-risk estimations: Evidence from foreign exchange and global equity portfolios. *Federal Reserve Bank of New York*.
- McNeil, A.J. 1998. Calculating quantile risk measures for financial return series using extreme value theory. *ETH preprint*.
- McNeil, A.J. 1999. Extreme value theory for risk managers. *ETH preprint*
- McNeil, A.J., and Frey, R. 2000. Estimation of tail-related risk measures for heteroscedastic financial time series. *Journal of Empirical Finance* 7: 271-300.
- J.P. Morgan. 1996. Riskmetrics technical manual. J.P. Morgan Bank New York.
- Raatikainen, J. 2002. Testing forecasting ability of value-at-risk models. *Risk Management Seminar in Jyväskylä*. <http://www.jyu.fi/finance/seminaari02/backte~1.ppt>.