DC Wentzel and E Maré*

Extreme value theory – An application to the South African equity market

1. INTRODUCTION

The assumption of normal returns remains ubiquitous in much of modern finance. In the pricing of conditional liabilities, for example, we frequently assume that logarithmic returns are normally distributed. Many risk management metrics, such as Value at Risk (VaR), and performance metrics such as Sharpe ratios also assume normally distributed movements.

There is however evidence, that financial market returns display leptokurtotic behaviour (also known as heavy tailed behaviour). This was already noted by Fama in 1965. More recent work includes, for example, Campbell, Lo and Mackinlay (1997), Duffie and Pan (1997), Focardi and Fabozzi (2003), Huisman, Koedijk and Pownall (1998), Johansen and Sornette (1999) and Longin (1996).

Consider, for example, recent events in the South African equity market. During the month of June 2006 the market experienced high levels of volatility. The largest downward move in the FTSE/JSE TOP40 index was of the order of five standard deviations, roughly a 7% decrease. If returns were normally distributed this event would only be expected to occur approximately once every 52 000 years …

In general the distributions of market movements have heavier tails than allowed for under the normality assumption. In this paper we focus on a method that is particularly suitable to modelling large movements with little history, namely extreme value theory (EVT).

We begin with a short description of extreme value theory followed by an application to the South African FTSE/JSE TOP40 equity index. We also provide comparisons between movements predicted by the normality assumption and those predicted using extreme value theory. We provide some "backtesting" results to highlight the differences.

2 EXTREME VALUE THEORY

Predictive power is one of the main goals in modelling. This normally takes the form of fitting some equation or known distribution to observed data points. For example, when trying to predict the probability of a wave being six feet high during high tide the customary approach would be to fit the normal distribution to the heights of all the waves on the beach in question. The probability can then easily be calculated. The probability that it reaches a height of forty feet is typically very small and this is said to lie in the tail of the distribution.

The disregard of extreme values in the statistical literature and frequently in practise can probably be traced back to comments made by Fourier around 1824. His comments resulted in what is known as the three standard deviations rule, which suggests that all observations more than three standard deviations from the mean can be neglected.

However, should we link a potential catastrophe to a forty foot wave, accurate modelling of the tail becomes important as catastrophes could lead to large losses. It would be ill-advised to assume normality, and ignore the potential for a catastrophe. We therefore need a method that can adequately model the tail.

Extreme value theory is a statistical technique that is especially useful in modelling the tail of a distribution. It has been applied in many areas including hydrology, oceanography and structural engineering. More recently it has found a number of applications in finance (See, for example, Embrechts, Klüppelberg and Mikosh (1997), McNeil (1997, 1999) and Gençay, Selçuk and Ulugülyağci (2002) for applications in insurance and financial risk management.) EVT focuses on the observations that lie in the tail and attempts to fit a distribution to these observations. It then provides the functionality to extrapolate into unknown areas of the tail.

In the financial markets we typically consider returns. A return that is out of the norm, e.g. a 7% movement in the FTSE/JSE TOP40 index, could result in a large loss or a large gain depending on the position held. Such a movement typically lies in the tail of the FTSE/JSE TOP40 index's return distribution.

There are two well known methods of applying EVT, the Block Maxima approach and the Peaks Over Threshold approach. The group of models for threshold exceedances are more modern and powerful than the Block Maxima models (McNeil, Frey and Embrechts, 2005) and so we focus on this approach and its application to the returns on the FTSE/JSE TOP40 index.

3. APPLYING EXTREME VALUE THEORY TO THE FTSE/JSE TOP40 INDEX

We used 10 years of daily FTSE/JSE TOP40 closing index data, covering the period 1996 to 2006. The

*Department of Mathematics and Applied Mathematics, University of Pretoria 0001, Republic of South Africa.
Email: eben.mare@up.ac.za

analysis was conducted using the statistical package R combined with the Extreme Values in R (EVIR) package. Although we used log returns in our analysis below, the tabled results have been converted to actual return figures. Furthermore, we converted the negative log returns to positive log returns to ease the calculation in R and hence we talk about positive decreases. In all cases we worked with daily return figures.

The Peaks Over Threshold approach attempts to fit the generalized Pareto distribution (GPD) to the observations that exceed some predetermined threshold (McNeil et al 2005). The distribution function of the GPD is given by

$$G_{\xi,\beta(x)} = \begin{cases} 1-\left(1+\xi x/\beta\right)^{-1/\xi}, & \xi \neq 0, \\ 1-\exp(-x/\beta), & \xi = 0, \end{cases}$$

where the parameters ξ and β are referred to, respectively, as the shape and scale parameters. We refer to the threshold as u. The methods used to fit the GPD have been thoroughly documented (Embrechts et al (1997)) and we used the maximum likelihood method in R to fit the distribution. The maximum likelihood method is the most general fitting method used in statistics.

The choice of threshold, however, remains subjective (McNeil 1999), so we performed exploratory data analysis in R to aid in the threshold choice.

The data analysis commences with a QQ-plot against the exponential distribution.

Quantiles above the exponential quantiles indicate that the data exhibits a heavier tail than that of an exponential distribution. As expected our data indicates a heavier tail. We now proceed to choose our threshold level.

Mean excess plots are used to decide on a suitable threshold level. In general if the plot has a positive gradient and follows a straight line above a certain level, we can assume the data follows a GPD above the particular level (McNeil et al 2005). We then test this level as a threshold. The following two figures are mean excess plots where we omitted firstly the three greatest (Figure 2) and then the hundred greatest (Figure 3) returns.

Clearly choosing a threshold based on analysing the figures is difficult. It is important to keep in mind that by choosing a lower threshold a better fit is obtained but one risks losing the extreme behaviour of the tail. Similarly choosing too high a threshold will result in an inadequate fit. We settled on a threshold of 0,0165 as the plot in Figure 3 tends to behave in a linear fashion with a positive slope beyond this level. Figure 4 shows the results of our fit as we plot the observed returns against the fitted tail based on our estimated parameters (ξ = 0,30, β = 0,00759 and u = 0,0165).

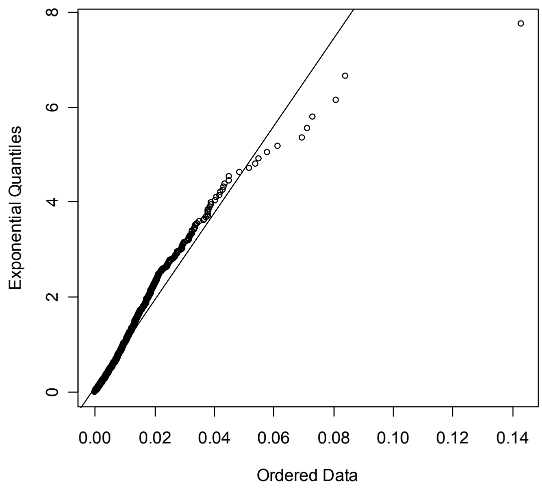

Figure 1: QQ plot against exponential distribution

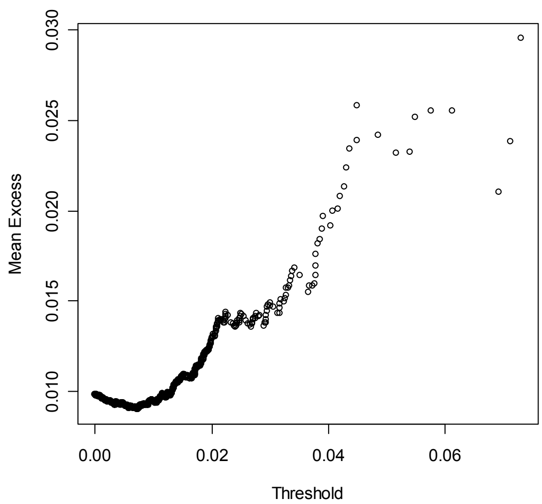

Figure 2: Sample mean excess plot, omitting 3 largest observations

4. DISTRIBUTIONS OF FTSE/JSE TOP40 RETURNS

We now focus on some results obtained by applying EVT. For ease of illustration, results are expressed as a probability in the form of once in x years for a specified decrease in the FTSE/JSE TOP40 Index. We also calculate the 1st percentile to serve as a comparison against "normal" movements. This assumes 250 trading days per annum. The results obtained are compared to the equivalent results under the normal distribution assumption and are shown in Table 1.

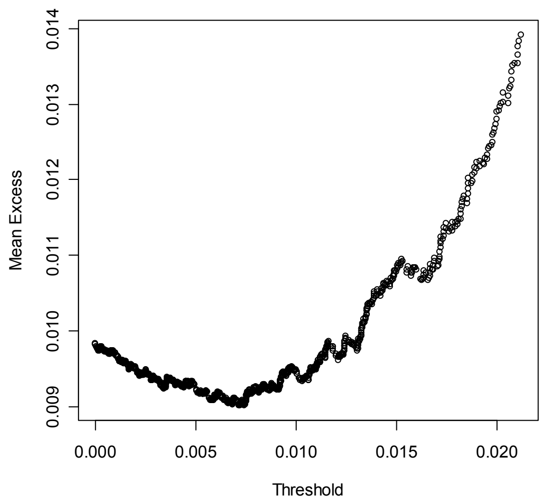

Figure 3: Sample mean excess plot, omitting 100 largest observations

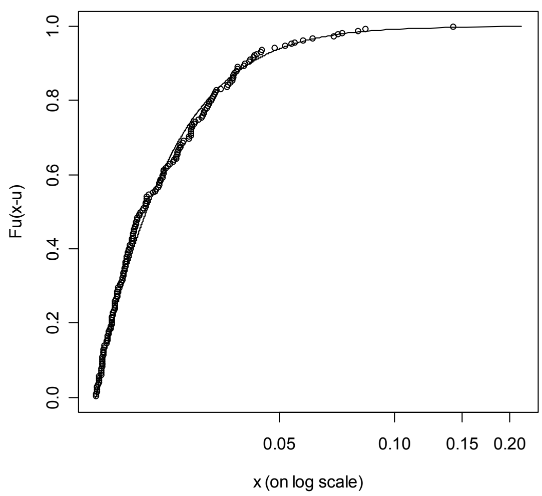

Figure 4: Performance of fit in the tail area

The differences are not difficult to spot. Under "normal" market conditions, for a movement we expect once every 100 days (1st percentile), the results appear similar. As soon as we expand the movement to, say, three standard deviations, the differences become more pronounced. The normal distribution basically implies that market decreases of 10% (or more) will not happen in our lifetime, perhaps not even in the lifetime of our solar system, yet a 13% index movement did occur in 1997. Analyses based on the assumption of normal returns could therefore be flawed as a result of the normal distribution's inability to cater for the observed heavier tails.

We are not advocating the abolishment of the normal assumption, as EVT also has its own built-in flaws. We merely highlight that large movements may (and do) occur with relatively high probability contrary to the belief of many who find moves typically above 5 standard deviations unimaginable.

To verify the substantiality of our results we performed some "backtesting" to assess the goodness-of-fit of the EVT model.

5. COMPARISON AND BACK-TESTING THE MODEL

In this section we present the results of a comparison between EVT, the normal assumption and the observed distribution and comment on the results obtained when testing EVT and the normal assumption (based on 1996 to 2001 movements) against the movements that occurred in the period 2001 to 2006.

We compared the probability of an average move within a number of intervals using EVT, the normal distribution and the actual observed probability based on frequency analysis. The intervals are based on a multiple of standard deviations to represent the downward movements in the FTSE/JSE TOP40. This allows us to compare the distributions for a specified multiple of standard deviations. We assumed the typical move in each interval to be the average of the borders of the interval. Table 2 shows the results obtained. The probabilities are given as a decrease of the relevant magnitude occurring once every x years. For instance in Table 2 the Peaks Over Threshold approach estimates a decrease of magnitude 6,9% to 8,3% to occur once every 3,2361 years, and the normal distribution estimates a decrease of magnitude 1,4% to 2,8% to occur once every 0,0657 years. The table also shows the percentage difference between the observed probability and the probabilities obtained under EVT and the normal distribution.

It can be seen quite clearly that EVT (represented in the table by "Peaks Over Threshold") is not a good predictor when we model events that do not occur in the tail of the distribution. However, the differences become relatively smaller as we increase intervals. This clearly shows the ability of EVT to better model large movements. The table also shows that large moves are far more likely to occur than is assumed under the normal assumption.

Finally we tested the models by fitting them over market movements during the period 1996 to 2001 and compared the results with the actual observed probabilities obtained from the market movements during 2001 to 2006. The results are presented in Table 3.

Table 1: EVT and normal probabilities for extreme events in the FTSE/JSE Top40 index

	EVT probability	Normal probability
1st percentile	3,75% decrease	3,14% decrease
5% decrease	Once in 0,9 years	Once in 40,5 years
10% decrease	Once in 7,7 years	Once in 2,8 x 10^{11} years
25% decrease	Once in 183,2 years	Once in 2,5 x 10^{92} years

Table 2: Comparison between observed probabilities and probabilities obtained through EVT and the normal distribution for the period 1996 to 2006

	Probabilities given as once every X years				
	Fitted and Observed probabilities based on 1996 to 2006 movements				
Interval (decreases)	Observed	Peaks Over Threshold	% Difference	Normal	% Difference
0,0% to 1,4%	0,0045	0,0104	-134%	0,0135	-204%
1,4% to 2,8%	0,0514	0,0876	-70%	0,0657	-28%
2,8% to 4,2%	0,2260	0,3222	-43%	0,7915	-250%
4,2% to 5,5%	1,0800	0,8294	23%	26,3082	-2 336%
5,5% to 6,9%	2,4300	1,7469	28%	2614,14	-107 478%
6,9% to 8,3%	3,2400	3,2361	0%	8,32E+05	-25 681 323%
8,3% to 9,7%	-	5,4826	-	9,05E+08	-
9,7% to 11,1%	-	8,6977	-	3,59E+12	-
11,1% to 12,5%	-	13,1208	-	5,52E+16	-
12,5% to 13,9%	9,7200	19,0212	-96%	3,53E+21	-4,0E+20
13,9% to 15,3%	-	26,7008	-	1,01E+27	-

Table 3: Comparison between observed probabilities for the period 2001 to 2006 and probabilities obtained through EVT and the normal distribution for the period 1996 to 2001

	Probabilities given as once every X years				
	Probabilities	Fitted probabilities based on 1996 to 2001			
Interval (decreases)	based on observed 2001 to 2006 movements	Peaks Over Threshold	% Difference	Normal	% Difference
0,0% to 1,3%	0,01	0,02	-49%	0,01	-7%
1,3% to 2,5%	0,04	0,07	-74%	0,04	-3%
2,5% to 3,8%	0,23	0,19	19%	0,24	-6%
3,8% to 5,0%	1,62	0,42	74%	2,97	-83%
5,0% to 6,3%	4,86	0,82	83%	75,58	-1 455%
6,3% to 7,6%	4,86	1,47	70%	4,24E+03	-87 062%
7,6% to 8,8%	-	2,43	-	5,45E+05	-
8,8% to 10,1%	-	3,80	-	1,68E+08	-

The table shows that the normal distribution is inadequate in modelling large movements and, whilst EVT doesn't provide a perfect fit, it certainly gives more accurate results when we move into the tail of the distribution. More importantly the normal distribution rated 5 to 6 percent decreases as almost impossible during our lifetime whereas EVT showed them as being highly probable. Looking at the actual movements during 2001 to 2006, decreases of the order of 5 to 6 percent occurred at least once.

Comparing the EVT results with those obtained in Table 2, we notice some differences. These can partially be attributed to the fact that the period 1996 to 2001 was generally more volatile than the period 2001 to 2006. The other important factor stems from the fact that EVT is extremely sensitive to the underlying data used in the modelling; a small change in the estimated parameters can lead to a large difference in estimation. The sensitivity is also a consequence of the somewhat subjective choice of threshold as discussed in Section 3.

6. CONCLUSION

As seen from the analysis presented, returns on the South African equity market, here represented by the FTSE/JSE TOP40 index, display some serious departures from those obtained using the assumption of normal return distributions. Extreme value theory serves as a good reminder that actual market returns are not normally distributed. It highlights the implications of assuming normality and it can give a good estimate of the movements we may see when exogenous events cause market turmoil.

Our analysis was based on the FTSE/JSE TOP40 index as a proxy for the equity market; however, it could easily be applied to individual equities. As remarked by Fama (1965) and Focardi and Fabozzi (2003), there is much value to be extracted in the investment process by understanding the distribution of returns of a common stock. The normal distribution gives us a fair idea of return distributions for everyday events. EVT gives us an understanding of the best/worst case returns and the frequency thereof.

This theory will not replace the more convenient assumption of normally distributed asset returns but has its place in risk management. People are generally more comfortable working with the normal distribution owing to its friendly mathematical properties. Although using EVT presents a significant increase in work, the results obtained justify the efforts. Appreciation of the true distribution of returns not only presents us with trading opportunities but also a clearer picture of the risk involved in an investment decision.

ACKNOWLEDGEMENTS

The authors wish to express their appreciation to the editor for several useful comments and suggestions made which improved significantly on an earlier version of this paper. The authors also wish to thank the anonymous referees for their suggestions.

REFERENCES

Campbell JY, Lo AW and Mackinlay AC. 1997. *The econometrics of financial markets*. Princeton University Press.

Duffie D and Pan J. 1997. An overview of value at risk. *Journal of Derivatives*, 4: 7 – 49.

Mathematics, ETH, Swiss Federal Technical University.

Embrechts P, Klüppelberg C and Mikosch T. 1997. *Modeling extremal events for insurance and finance*. Springer.

Fama EF. 1965. The behaviour of stock-market prices. *Journal of Business*, 38(1): 34 – 105.

Focardi SM and Fabozzi FJ. 2003. Fat tails, scaling, and stable laws: A critical look at modeling extremal events in financial phenomena. *Journal of Risk Finance*, 5 – 26.

Gençay R, Selçuk F and Ulugülyağci A. 2002. High volatility, thick tails and extreme value theory in Value-at-Risk estimation. *Insurance: Mathematics and Economics* 33:337-356.

Huisman R, Koedijk KG and Pownall RAJ. 1998. Var-x: Fat tails in financial risk management. *Journal of Risk*, 1(1): 47 – 61.

Johansen A and Sornette D. 1999. Critical crashes, *Risk* 12(1): 91-94.

Longin FM. 1996. The asymptotic distribution of extreme stock market returns. *Journal of Business*. 69(3):383 – 408.

McNeil AJ, Frey R and Embrechts P. 2005. *Quantitative risk management: Concepts, techniques and tools*. Princeton University Press.

McNeil AJ. 1999. *Extreme value theory for risk managers. Internal modelling and CAD II.* RISK Books, 93-113.

McNeil A J. 1997. *Estimating the tails of loss severity distributions using extreme value theory*. ASTIN Bulletin, 27: 1117–137.