Pakistan Journal of Applied Sciences 1(1):24-32, 2001

The Effect of a Dynamic Exponential Decay Factor on Volatility and VaR

Marius Botha[1], Gary van Vuuren[2], Paul Styger[1]
[1]Potchefstroomse Universiteit vir Christelike Hoër Onderwys,Department of Economics, Private Bag X6001, Potchefstroom, 2520, South Africa
[2]Old Mutual Asset Managers, PO Box 878, Cape Town, 8000, South Africa

Abstract: Value at Risk models rely heavily on the accuracy of price volatility estimates to measure capital that could potentially be lost over given time horizons. Volatility measured by equally weighting instrument price returns is often erroneous, but employing exponential envelope weighting largely circumvents the inaccuracies. The steepness of the envelope applied to returns depends upon a single number, λ, assumed unique and constant in any given market. This paper examines the historical evolution of λ for South African-specific data and shows it to vary constantly and significantly. The effects of these changes on volatility, correlation and ultimately, VaR, are also examined.

Key words: Dynamic exponenetial decay factory

Introduction

Value at Risk (VaR) is defined as the potential loss in value of a static financial portfolio due to changes in market rates or alternatively as a measure of the maximum potential change of a financial instrument portfolio with a given probability over a pre-set horizon (Reed, 1996). Stringent capital adequacy requirements which are based upon VaR (Basle, 1996) imposed by the Bank for International Settlements (BIS), coupled with the possibility of financial ruin from unexpected market turbulence, has fuelled the drive for the pursuit of an ever-increasing accuracy of institutional VaR.

Overestimating VaR results in the retention of too much idle, uninvested capital and underestimating VaR could spell disaster if a market crash occurred and the institution were unable to meet its obligations due to liquidity complications. Narrowing the gap between too much and too little capital in reserve is, therefore, widely regarded as one of the chief aims of any sound financial institution (Bernstein, 1999).

Many methods VaR models have evolved and of these three enjoy particular popularity, namely the historical, stochastic simulation and variance-covariance methods. The historical method uses past changes in market rates as a proxy for possible future changes whilst the stochastic simulation method preserves correlations between instrument returns and shocks market rates. The variance-covariance technique, popularised by JP Morgan's RiskMetrics™, uses historic price return volatilities to estimate VaR (JP Morgan, 1996). Of these three techniques, the variance-covariance method has emerged as the most popular due to its relative simplicity, ease of implementation and minimal demand upon computer resources (Shimko, 1996).

In this paper the accuracy of the variance-covariance technique for calculating VaR is examined. Specifically, various techniques used to estimate volatility, as well as modifications to these techniques to improve accuracy, were explored. The improved volatility and correlation estimates lead directly to a more accurate VaR.

Section 2 explores the constituents of the VaR equation employed in the variance-covariance method, including a brief discussion of the two most popular volatility estimation techniques.

Section 3 discusses the principles underlying the volatility and correlation estimation methods as well as inherent flaws and provides possible modifications employed to overcome these limitations.

Section 4 presents the results obtained from the investigation. These results explore the effect of a dynamic exponential envelope on historical price returns as well as volatility and correlation estimates. The resulting effects on VaR are also critically examined in this section.

Section 5 provides conclusions that may be reached as a result of this investigation and suggests possibilities for future study.

Materials and Methods

The VaR equation: The central equation used to calculate VaR using the variance-covariance method for a single financial instrument, measured at time t, is given by

$$\mathrm{VaR}_t = CI \cdot \sqrt{T} \cdot N \cdot \sigma_t \tag{1}$$

in which CI is the confidence interval measured off an assumednormaldistribution of profits and losses, $\sqrt{T}$ is the square root of the unwind period required to reverse a financial position, N is the nominal value of the instrument under consideration and σ_t is the price volatility forecast of that instrument. For a portfolio of instruments, the RiskMetrics™ method for the calculation of VaR uses the variance/covariance matrix to take account of the correlations between instruments and isgiven by

$$\mathrm{VaR} = CI \cdot \sqrt{T \cdot \left(\overset{\rho}{\mathbf{V}}\, \mathbf{C}\, \overset{\varpi}{\mathbf{V}}^{\perp}\right)} \tag{2}$$

where **V** is the weighted volatility matrix of instrument returns, is the correlation matrix between price returns and the **C** superscripts $\perp$ represents the relevant transposed matrix. This equation relies heavily on certain fundamental assumptions (Zangari, 1996, Dowd, 1998) about the behaviour of financial instrument returns. Many of these assumptions do not hold under conditions other than those found in stable markets (Smithson, 1996a, 1996b).

Confidence interval: Choosing a confidence level gives a cut-off point to delineate the left-hand tail from the rest of a normal distribution. The confidence level may be chosen to be any value desired, depending on the risk tolerance level of the institution. If it were chosen to be 99%, the VaR estimate would cover all but the largest 1% of losses. A 95% confidence level would cover all but the highest 5% of all losses, and so on.

Ultimately, what is really important from a risk management point of view is the behavior of the tails of the distribution of interest. The estimation of VaR at very high confidence levels is non-trivial. The higher the confidence level, the more important the need to model precisely the changes in portfolio values that would occur for large changes in market rates. At high confidence levels it is almost never sufficient to approximate the change in an option portfolio by taking only its delta and gamma into account. In addition, the higher the confidence level, the more important is the need to incorporate fat tails into the assumed distribution of changes in market rates (Picoult, 1998).

Unwind period: The holding period appropriate in any given market is the length of time it takes to ensure orderly liquidation of positions in that market. Three main factors affect the choice of unwind (or holding) period. These are:

1. The liquidity of the markets in which a financial institution operates. If positions can be liquidated relatively quickly, that position's VaR should be based on a short unwind period. If liquidations take time, longer periods are appropriate. Since orderly liquidation takes different times in different markets, institutions are often faced with little choice but to operate with common holding periods across all markets (Dowd, 1998).
2. The justification of the normal approximation. Typically, portfolios contain options so the portfolio return is not normal. The normal approximation provides a reasonable fit to the real position only if the observation (and hence holding period) is kept short (Dowd, 1998).
3. To accommodate changes in the portfolio itself. Portfolio managers are more likely to change the portfolio the longer the unwind period. If the portfolio makes persistent losses, management will shift out of loss-making positions and into others. The VaR analysis, however, requires that the portfolio remains the same over the holding period (Dowd, 1998).

Volatility and correlation: Volatility forecasts of the instruments comprising a portfolio are required as inputs into the variance-covariance method for calculating VaR for that portfolio. The quality of VaR estimates depends directly on the caliber of these forecasts and if they have been inadequately estimated, inaccurate VaR values can result.

Methods of obtaining estimates of price volatilities and correlations between returns are numerous, detailed and require the determination of several input parameters (see, for example, JP Morgan, 1996, Shimko, 1997 and Leong, 1996). Slight inaccuracies in the determination of these parameters may be small in themselves, but the cumulative combination of many such errors can be amplified into considerable, and unacceptable, discrepancies when estimating the volatility.

Methods of estimating volatility are presented in the next section.

Volatility Estimation:

Equally weighted moving average: The equally weighted moving average technique for determining price volatility is defined as the standard deviation of the returns as measured by standard statistical theory and as given in Equation (3). For a set of T returns, the equally weighted (standard deviation) volatility is given mathematically by:

$$\sigma^2 = \frac{1}{T} \sum_{t=1}^{T} (r_t - \bar{r})^2 \tag{3}$$

where r_t is the price return on day t and $\bar{r}$ is the average of all T returns (JP Morgan, 1996). The $1/T$ factor effectively weighs each point in the series equally, i.e. it ignores the temporal order of the observations. *Each* $(r_t - \bar{r})^2$ term is multiplied by $1/T$ so each point, no matter how long ago it occurred in the past, is deemed to be as important as any other, no matter how recently it occurred. Market shocks that come in to and fall out of the measurement sample alter volatility values with equal abruptness (Wei, 1989 and see Figs. 3 and 4) and may not reflect true variations in current volatility. If market conditions are relatively stable when a large perturbation falls out of the return series (as the sample period rolls forward), volatility decreases abruptly for no *contemporary* reason. More accurate techniques, capable of reacting faster to changing market conditions and recovering more quickly from large fluctuations, must be sought.

Exponentially weighted moving average: A more promising alternative to the equally weighted method described above is the exponentially weighted moving average (EWMA). This technique takes the volatility forecast, σ_t, to be a weighted average of the previous period's forecast volatility, σ_{t-1}, and the current squared return, r^2_t, thereby effectively giving less weight to events more distant in time than contemporary ones. Mathematically, this is described by:

$$\sigma_t^2 = \lambda \sigma_{t-1}^2 + (1-\lambda) r_t^2 \tag{4}$$

where the weight λ is chosen in such a way as to minimise the error between the forecast and measured volatilities over a sample period (JP Morgan, 1996) and all other symbols have the same definition as given for Equation (3).

This approach also captures the dynamic features of volatility by allowing it to vary from one period to another and it explains volatility clustering since a higher volatility in one period is likely to lead to a higher-than-average volatility in the next (Dowd, 1998). Equation (4) may be written as an exponentially weighted moving average of a data series:

$$\sigma_t^2 = (1-\lambda) \sum_{t=1}^{T} \lambda^{t-1} (r_t - \bar{r})^2 \tag{5}$$

where again, the symbols have the same definition as those given for Equation (3). The latest observations carry the highest weight in the volatility estimate and the most distant observations the lowest due to the λ^{t-1} weighting factor (as contrasted with the equal weighting applied by

the equally weighted moving average technique in Equation (3)). Since $0 \geq \lambda \geq 1$, the term λ^{t-1} decreases exponentially as t increases (i.e. moves into the past, so $t = 5$, for example, occurred 5 days ago).
This approach has two important advantages over the equally weighted model. Firstly, volatility reacts faster to shocks in the market as recent data carry more weight than data in the distant past. Secondly, after any large shock, the volatility declines exponentially as the weight of the shock observation falls as λ^{t-1} (JP Morgan, 1996).
This technique, employed in the RiskMetrics™ methodology, is relatively simple to implement and requires the input of only one parameter value, λ. The closer the value of λ to unity, the shallower the exponential envelope placed on observations and so the smoother the data series becomes (Alexander, 1997).
The determination of λ, however, is itself a daunting task and requires experienced risk management in order to ascertain, *inter alia*, inputs for its calculation (as well as which instruments to include and minutiae such as sample periods and sizes of increments in the determination of(λ) (JP Morgan, 1996). These inputs are discussed in greater detail in the next section.

Correlation: The EWMA model may be used to construct the covariance of and correlation between price return forecasts in the same manner as performed for the volatility forecasts. The covariance estimator is given by:

$$\sigma_{12t}^2 = (1-\lambda)\sum_{j=1}^{T}\lambda^{j-1}(r_{1t}-\bar{r}_1)(r_{2t}-\bar{r}_2) \tag{6}$$

where σ is the covariance between two return series r_{1t}, r_{2t} at time t and all other symbols have their usual definition.
Analogous to the expression for a variance forecast, the covariance may also be written in recursive form (JP Morgan, 1996). The one-day covariance forecast between any two return series, r_{1t} and r_{2t} made at time t is

$$\sigma^2_{12,t+1|t} = \lambda\sigma^2_{12,t|t-1} + (1-\lambda)r_{1t}\cdot r_{2t} \tag{7}$$

The most recent returns and covariances are weighted the most and the most distant, the least, in a similar manner to the calculation for the volatility.
The correlation, ρ_{12}, between two return series is given by:

$$\rho_{12,t+1|t} = \frac{\sigma^2_{12,t+1|t}}{\sigma_{1,t+1|t}\,\sigma_{2,t+1|t}} \tag{8}$$

where the σ's are as defined in Equation (7) (Alexander, 1997).
Determination of exponential decay factor: Volatility and correlation forecasts, based on the EWMA model, require that an appropriate value of the decay factor, λ, be chosen (JP Morgan, 1996). The RiskMetrics™ methodology uses 480 time series as inputs from world markets (i.e. global foreign exchange rates, swap rates, equity indices, money market rates, etc.) generating 480 variance and 114 960 covariance forecasts. Each time series has its own optimal decay factor, later combined into a single decay factor, $\tilde{\lambda}$.
Two aspects to the general problem of choice of decay factor arise, namely its theoretical constraints and its practical determination.
Theoretical constraints: The choice of λ's must be such that they are consistent with their respective covariance matrix, generated from the return series. For the covariance matrices to be properly defined, the constituent variances cannot be negative, the covariances must be equal (i.e. the matrices must be symmetric) and the correlation, ρ, between return series must be in the range $-1 \leq \rho \leq +1$. Although it is possible in principle to choose λ's that are consistent with their respective covariance matrix, the task is exceedingly burdensome in practice and necessitates the application of certain restrictions on the λ's. The RiskMetrics™ methodology applies a single optimal decay factor to the entire covariance matrix (JP Morgan, 1996).

Practical determination of optimal λ ($\tilde{\lambda}$):
The definition of the time $t+1$ forecast of the variance of the return, r_{t+1}, made one period earlier is

$$E_t\left[r^2_{t+1}\right] = \sigma^2_{t+1|t}$$

that is, the expected value of the squared return one period earlier. In a similar manner, the definition of the time $t+1$ forecast of the covariance between two return series, $r_{1,t+1}$ and $r_{2,t+1}$ made one period earlier is

$$E_t\left[r_{1,t+1}r_{2,t+1}\right] = \sigma^2_{12,t+1|t}$$

The above results hold for any forecast made at time $t+j$ where $j \geq 1$ (JP Morgan, 1996).
If the variance forecast error is defined as

$$\varepsilon_{t+1|t} = r^2_{t+1} - \sigma^2_{t+1|t}$$

it follows that the expected value of the forecast error is zero, i.e.

$$E_t\left[\varepsilon_{t+1|t}\right] = E_t\left[r^2_{t+1}\right] - \sigma^2_{t+1|t} = 0$$

A requirement for the choice of λ suggests itself from the above formulation; that of choosing λ in such a way as to minimise the average squared errors. When applied to daily forecasts of *variance* this leads to the daily root mean squared prediction error, given by

(variance) (9)

$$RMSE_v = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(r_{t+1}^2 - \sigma_{t+1|t}^2(\lambda)\right)^2}$$

where the forecast value of the variance has been written explicitly as a function of λ. In practice, the optimal value of λ is obtained by searching for the smallest *RMSE v* over various values of λ. This involves seeking the decay factor that produces the best forecasts or minimises the forecast measures (JP Morgan, 1996).

Similar results to those for the variance may be derived for the covariance forecasts. The covariance forecast error, using the same technique as described above, is

$$\varepsilon_{12,t+1|t} = r_{1,t+1} \cdot r_{2,t+1} - \sigma_{12,t+1|t}^2$$

such that

$$E_t\left[\varepsilon_{12,t+1|t}\right] = E_t\left[r_{1,t+1} \cdot r_{2,t+1}\right] - \sigma_{12,t+1|t}^2 = 0$$

(JP Morgan, 1996) and

$$RMSE_c = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(r_{1,t+1} \cdot r_{2,t+1} - \sigma_{12,t+1|t}^2(\lambda)\right)^2}$$

(covariance) (10)

The method of determining λ from Equation (9) and (10) involves minimising each *RMSE* and *RMSE* for a series of i data sets. The RiskMetrics™ methodology uses 480 such data sets from world-wide financial time series, thereby obtaining 480 optimal decay factors. These are combined to obtain a single λ (called the optimal λ or $\tilde{\lambda}$) in such a way that $\tilde{\lambda}$ is a weighted average of individual optimal decay factors (see Footnote 9). The weights are a measure of individual forecast accuracy.

Practical issues: The value of the South African $\tilde{\lambda}$ was calculated (Botha and Van Vuuren, 1999) using the technique provided in detail by JP Morgan (1996).

The price and yield series used as input to determine $\tilde{\lambda}$ were: the (spot) interest rate, the 1, 3, 6, 9 and 12-month money market rates, the 2, 3, 4...20 year Swap rates, the Johannesburg Stock Exchange (JSE) All-Share price Index, the South African Gold price Index and the US$/ZAR and the GB£/ZAR foreign exchange rates (Blaauw, 1997, Botha and Van Vuuren, 1999). Three years (about 760 trading days from 1 January 1997 to 1 January 2000) in total of historical data were analysed.

One year of return history was used to estimate $\tilde{\lambda}$ and this rolled forward one day at a time to give two full years of historical $\tilde{\lambda}$ values. Thus, the first calculated value of $\tilde{\lambda}$ (on 1 January 1998) uses data from 1 January 1997 to 31 December 1997. The second value (on 2 January 1998) uses data from 2 January 1997 to 1 January 1998, and so on. Shorter periods than one year were used to calculate $\tilde{\lambda}$, but values thus obtained tended to be unstable as these periods took insufficient market information into account. Longer periods than one year were also used to estimate historical $\tilde{\lambda}$, but these did not yield substantially different or more accurate values for $\tilde{\lambda}$ than those obtained using one year of data.

Yield volatilities, σ_y, were converted into price volatilities, σ_y, using $\sigma_p = Dy\,\sigma_y$ where D is the modified duration and y the current yield of the relevant instrument (Hull, 1997).

Application aspects: Reuters, in conjunction with JP Morgan, provide free, internet-based, daily and monthly exponentially weighted price return volatilities and correlations. JP Morgan provides these data for most countries involved in the global economy as well as country-specific $\tilde{\lambda}$'s. Whilst data for larger economies such as the US and European countries are extensive, significant gaps do exist for smaller economies. The South African data set, for example, contains only partial information on money market rates and has omissions in several of the capital market rates. Institutions operating within these smaller economies and desirous of current, accurate volatility and correlation data are faced at this point with numerous options. VaR values could, for example, be obtained using the reduced data set available from the Internet. With several of the RiskMetrics vertices entirely absent from this data set, significant interpolation between existing vertices is necessary and VaR thus obtained subject to large margins of error (Botha, 2000). An institution may incorporate missing vertices into the data set using price data available locally for those vertices. By applying an exponential envelope (for example with daily $\tilde{\lambda}$ = 0.94 for South African data as given by JP Morgan) to the daily return series for each missing vertex, volatilities could be calculated and manually entered into the data set. A problem with this technique involves the reliability of the supplied $\tilde{\lambda}$. Many emerging markets are assigned a $\tilde{\lambda}$ of 0.94, despite the uniqueness of individual economies. A third option available to institutions would be to calculate $\tilde{\lambda}$ using local economic and financial data as the basis for the calculation as discussed in Section 3.3.3. Using this $\tilde{\lambda}$, an exponential envelope could be applied to price return data to calculate volatilities and correlations between all vertices required. This option provides the most satisfactory and accurate value of a local institution's VaR, but other aspects, such as the day-to-day reliability of this (constant) $\tilde{\lambda}$, must also be considered.

Botha and Van Vuuren (1999) investigated the effect of a large market fluctuation on $\tilde{\lambda}$ and showed that it was unstable in turbulent market conditions. The current study builds on and extends this work in several aspects.

In order to simplify and isolate individual factors contributing to $\tilde{\lambda}$, Botha and Van Vuuren (1999) used a single instrument's price returns and introduced large (±6%) artificial market fluctuations of unrealistically short duration (2 days for complete market recovery). The current work uses some 30 financial instruments and some three years of actual historical market data thereby evaluating a true, historical $\tilde{\lambda}$ for the South African market as a whole. Botha and Van Vuuren (1999) used only a single instrument and obtained broad, generic results of this effect. The current study assumes a portfolio of 30 long positions of different instruments (the same as those used to determine $\tilde{\lambda}$) each with a single unit invested in the portfolio. Exponentially weighted correlations were not considered in the original study and are now also taken into account in accordance with Equations (7) and (8). The results of this investigation are presented in the following section.

Results and Discussion

Effect of changing market conditions on $\tilde{\lambda}$: Fig. 1 shows the evolution of $\tilde{\lambda}$ and the price return history of the JSE all share index on the same time scale. The first point of the $\tilde{\lambda}$ curve begins one year after the JSE return history begins since that point is calculated from the previous year of historical price index data. From that day onwards, the calculation rolls forward one day at a time, always using the previous one year of historical data to determine $\tilde{\lambda}$ on that day. Fig. (2) shows the evolution of the US$/ZAR foreign exchange price rate changes and $\tilde{\lambda}$, again on the same time scale. Although there are several other instruments involved in the calculation of $\tilde{\lambda}$, only these two are presented for brevity and between them show the bulk of the important features worth noting.

The first important feature to emerge is that $\tilde{\lambda}$ is not constant over time, but rather highly variable. Whilst it is true that $\tilde{\lambda}$ varies within a fairly narrow range (0.91 to 0.98 or some 8%), it is also true that these changes can occur dramatically within short (usually a single day) time scales. A second significant feature is the 'ghosting' effect or sudden changes in the value of $\tilde{\lambda}$ that occur with little accompanying change in volatility of underlying instruments. Consider the left end of the one-year time scale marked **1** indicated in Fig. (1), i.e. the brief but significant 10th anniversary of the October 1987 market crash, which occurred in October 1997. Exactly one year later (or the length of the sample period used to calculate $\tilde{\lambda}$), at the right end of this time scale, $\tilde{\lambda}$ jumps 4% from 0.936 to 0.973. This is a similar effect to that observed in many statistical series: abrupt changes when large fluctuations drop out of the measurement sample (and see Section 3.1). Similar events occur at the *ends* of the time scales labelled **2, 3** and **4** in Fig. (1), all occurring exactly one year after events that precipitated abrupt changes in the value of $\tilde{\lambda}$ at those times. The Asian Crisis, as it has come to be known, responsible for the September 1998 event, caused $\tilde{\lambda}$ to drop 2.5% in that month. A year later, in September 1999, $\tilde{\lambda}$ jumps up 3.7% as the Asian crisis event drops out of the measurement sample. The value of $\tilde{\lambda}$, then, depends as much on the underlying constituents' *present* volatility as it does on their volatility at the start of the sample period (in this case, one year).

The US$/ZAR currency fluctuations which began in June 1998 (indicated by the left hand end of the time scale labelled 1 in Fig. (2) result in a 1.1% drop in $\tilde{\lambda}$, but these do not cause nearly so large change in $\tilde{\lambda}$ (only 0.5%) when they leave the measurement sample at the right hand end of the time scale. The Asian crisis is again responsible for the large changes in $\tilde{\lambda}$ at either end of the time scale labelled 2. December 1998 fluctuations in US$ currency, due in part to the currency crisis in Brazil at that time, were responsible for large changes in $\tilde{\lambda}$ at either end of the time scale labelled 3. The October 1997 anniversary also contributed to the large jump in $\tilde{\lambda}$ at the right end of the time scale labelled 4.

Many variables are responsible for the behaviour of $\tilde{\lambda}$ and to segregate out individual effects is non-trivial. Some general 'rules' do, however, emerge from these results. It is clear that $\tilde{\lambda}$ reacts suddenly to severe fluctuations in financial markets. When these occur, $\tilde{\lambda}$ decreases, adjusting the shape of the exponential envelope so as to weigh more recent data even *more* heavily than before. When these fluctuations drop out of the measurement sample, $\tilde{\lambda}$ experiences an increase on the same time scale with the effect of now weighing more recent data *less* heavily than before.

Having evaluated historical values for $\tilde{\lambda}$, they were then applied to market returns for the 30 instruments discussed above in order to determine the effect of a changing $\tilde{\lambda}$ on these instruments' price return volatilities.

Effect of changing market conditions on σ: Fig. (3) shows the volatility of the JSE All Share price index returns calculated using (i) the equally weighted moving average technique (σ) (ii) the EWMA method with $\tilde{\lambda}$ constant at 0.94 ($\sigma(\lambda = 0.94)$) and (iii) the EWMA method with the daily (variable) calculated $\tilde{\lambda}$ measured using 100 days (about 5 trading months) of historical data measured from 1 January 1998 ($\sigma(\tilde{\lambda})$). Rolling this period forward one day at a time generates succeeding values.

Comparison of volatilities calculated using the equally weighted and exponentially weighted methods show familiar results, i.e. volatilities using the equally weighted method are higher than those measured using the EWMA method during periods of low volatility because of the ghosting effect and lower during periods of high volatility because of averaging.

Comparison of volatilities calculated for the two EWMA methods shows that $\sigma(\lambda = 0.94)$ can be as much as 18% higher than that measured using $\sigma(\tilde{\lambda})$ during periods of high market activity (e.g. Sep 1998). At periods of low volatility the values are very similar.

Fig. (4) shows the volatility of the US$/ZAR FX rate returns calculated using the same techniques used for Fig. (3) and also measured using 100 days of historical data. Comparison of volatilities calculated for the two EWMA methods shows that $\sigma(\lambda = 0.94)$ can be as much as 16% higher than that measured using $\sigma(\tilde{\lambda})$ during periods of high market activity (e.g. mid-July 1998). At periods of low volatility the calculated volatility values are again very similar.

Having evaluated the historical volatilities of the 30 instruments, the correlation between their exponentially

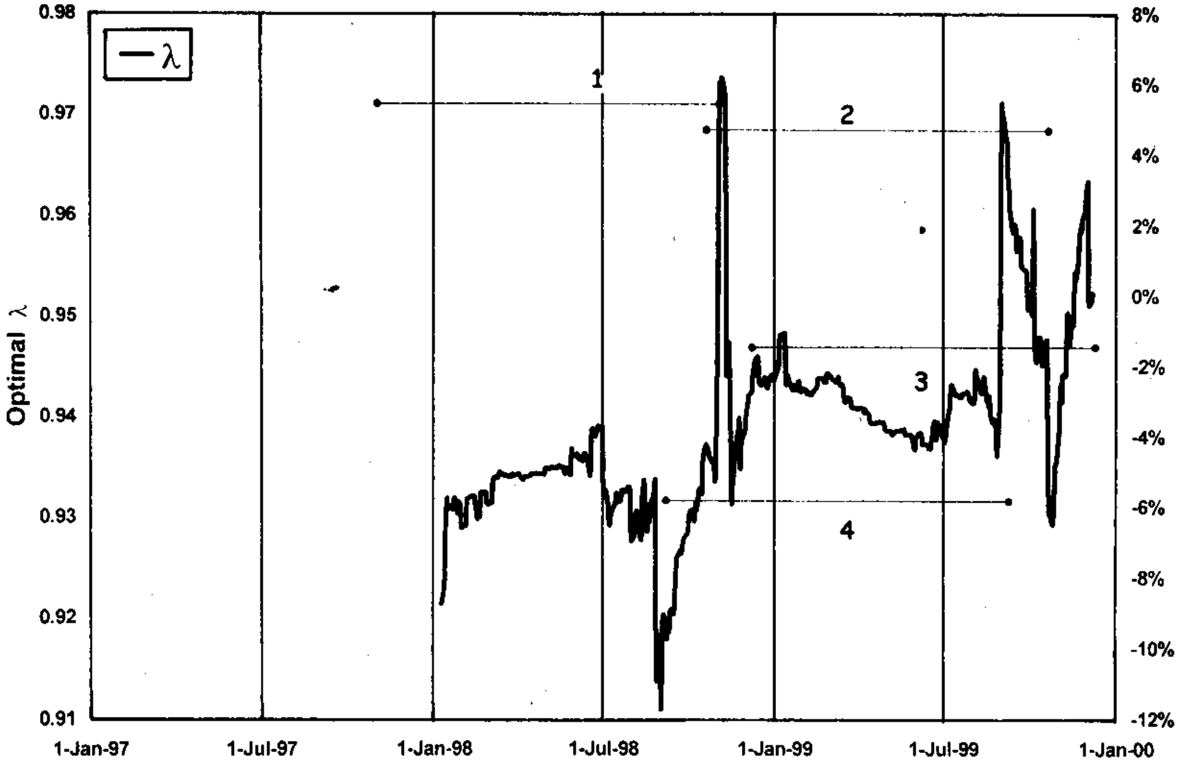

Fig. 1: Evolution of λ and the price return history of the JSE all share index on the same time scale from 1 January 1997 to 1 January 2000.

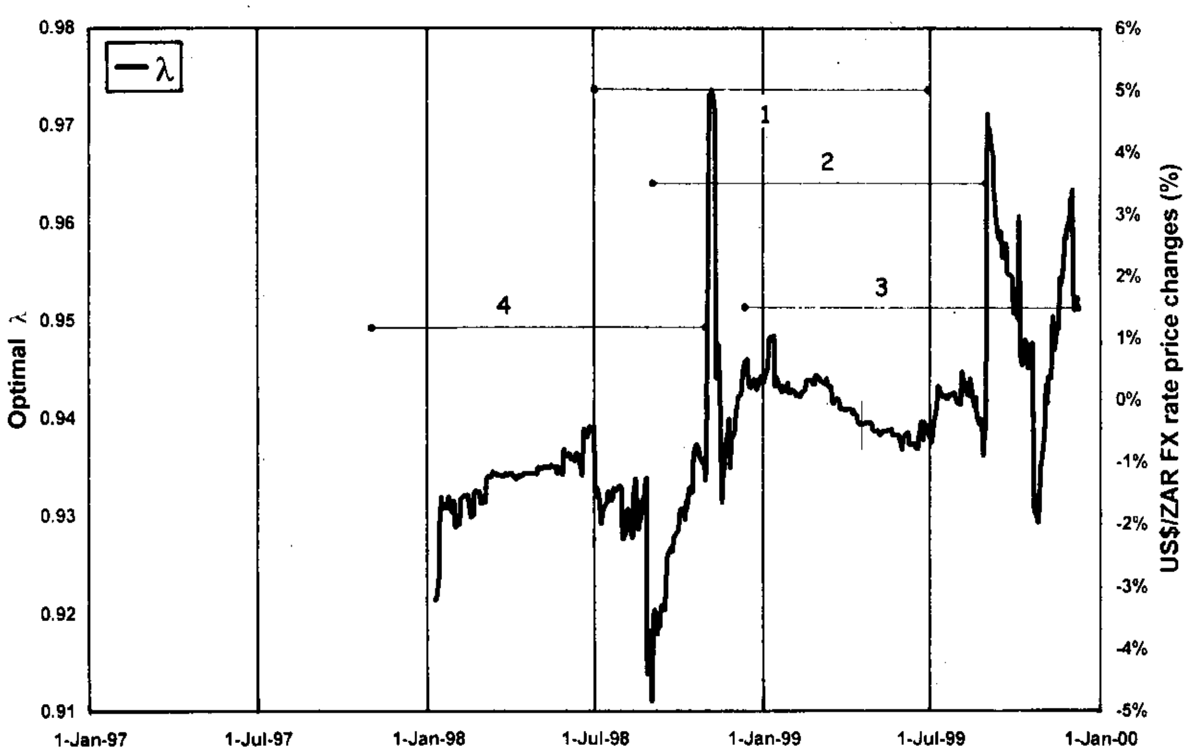

Fig 2: Evolution of the US$/ZAR foreign exchange rate and λ, again on the same time scale as in Figure (1).

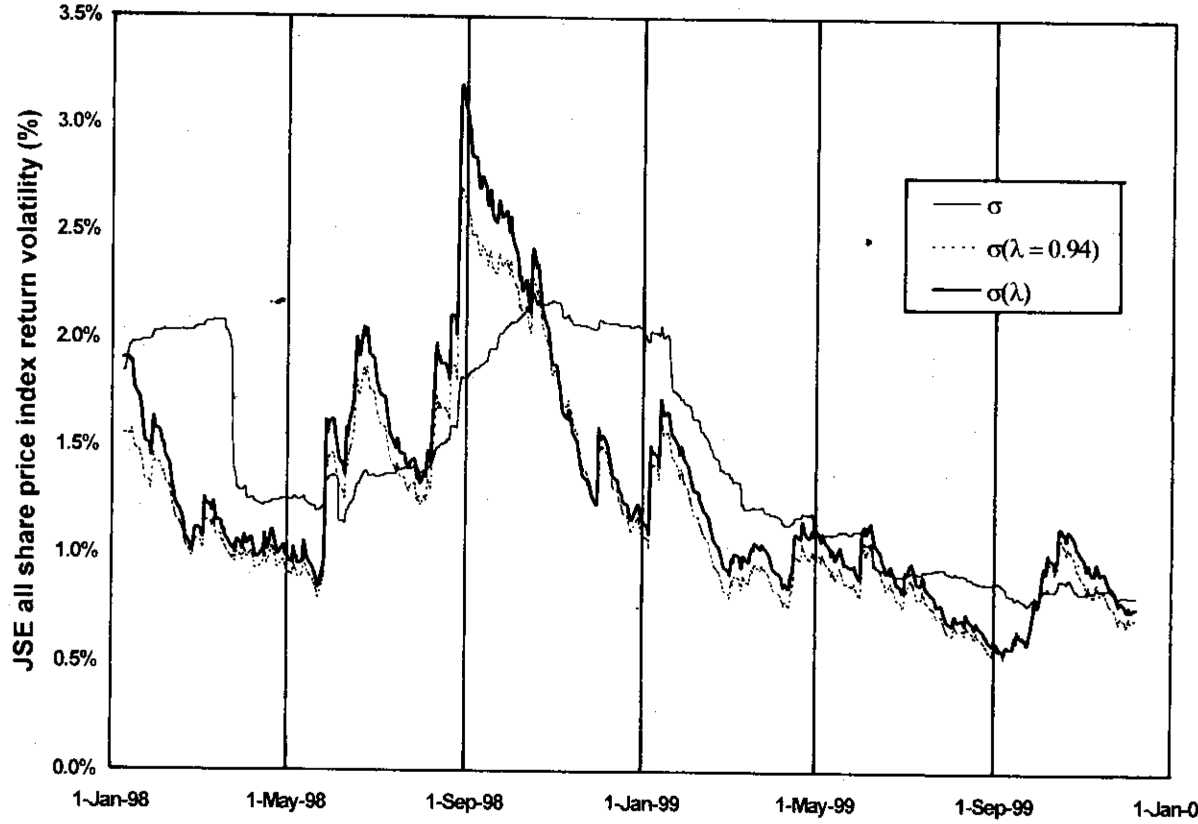

Fig. 3: Evolution of the volatility of the JSE All Share price returns calculated using (i) the equally weighted moving average technique (ii) the EWMA method with static $\tilde{\lambda}$ and (iii) the EWMA method with the dynamic $\tilde{\lambda}$ measured using 100 days of historical data measured from 1 January 1998.

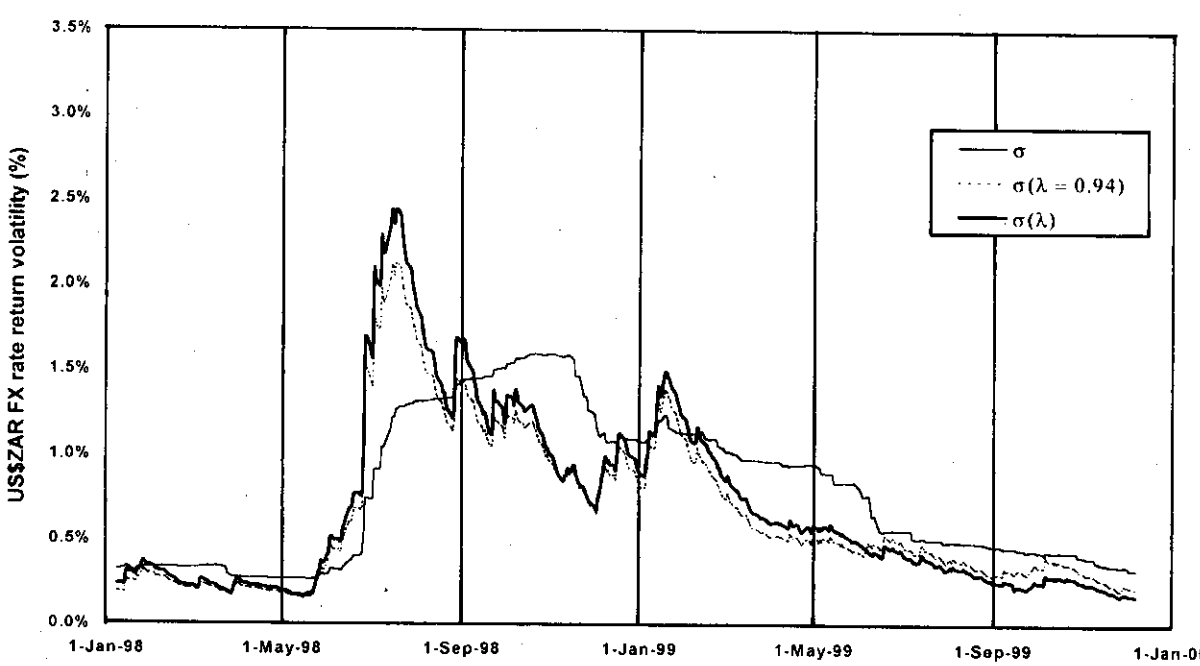

Fig. 4: Evolution of the volatility of the US$/ZAR FX rate price returns calculated using (i) the equally weighted moving average technique (ii) the EWMA method with static $\tilde{\lambda}$ and (iii) the EWMA method with the dynamic $\tilde{\lambda}$ measured using the same duration and time period as used for Fig. 3.

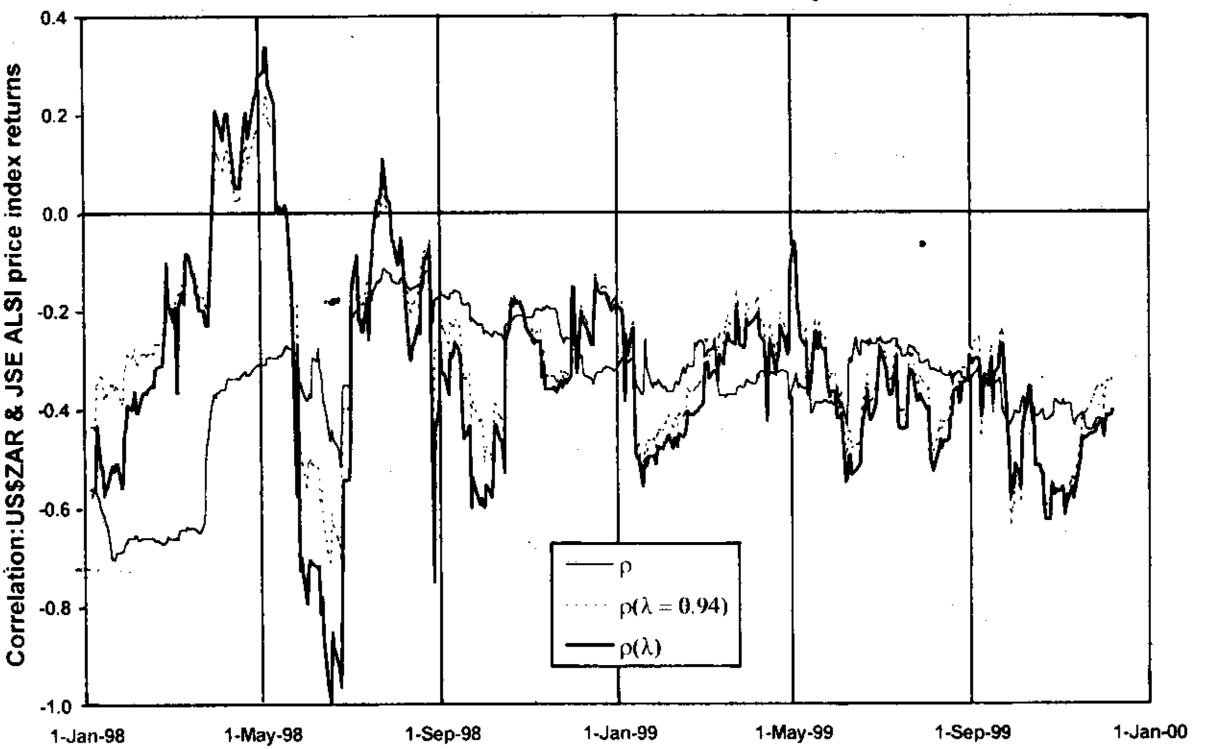

Fig 5: Evolution of the correlation between the US$/ZAR FX rate returns and the JSE All Share index returns calculated using (i) the equally weighted moving average technique (ii) the EWMA method with static and (iii) the EWMA method with dynamic measured using the same duration and time period used for Fig. 3 and 4

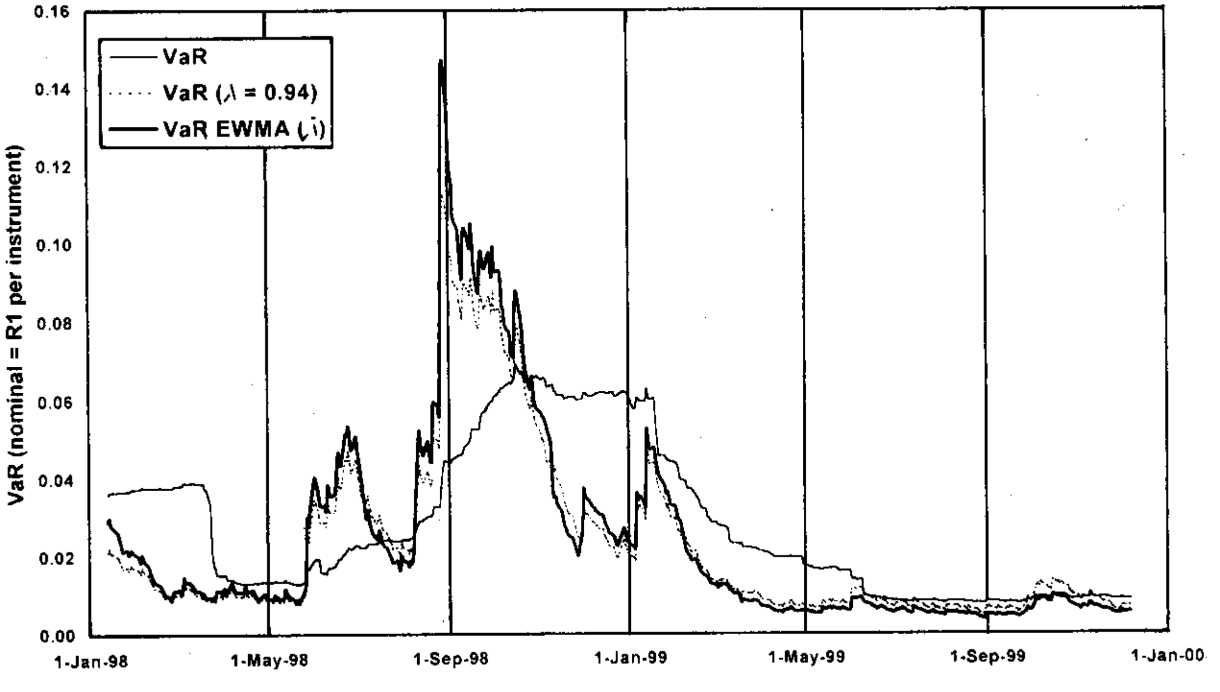

Fig. 6: Evolution of VaR (using the above three techniques) for the portfolio of instruments discussed in Section 3.3.4. R1 nominal value was 'invested' for each instrument (all long positions) and the portfolio remained 'static' (i.e. no instruments were bought or sold) throughout the measurement period. $CI = \sqrt{T} = 1$, and Equation 2 was used

weighted returns was investigated. The technique used to measure VaR in this study was the variance/covariance technique and hence only when *both* volatilities and correlations have been evaluated can a meaningful VaR be measured.

Effect of changing market conditions on ρ : Fig. 5 shows the correlation between the US$/ZAR FX rate returns and the JSE all share price index returns calculated using 100 days of historical data from 1 January 1998 using the three techniques indicated in the previous Figures. During periods of low volatility, correlations measured using all three techniques show similar results.
During the September 1998 emerging market meltdown, however, $\rho(\lambda = 0.94)$ = −0.445 and $\rho(\tilde{\lambda})$ = −0.626, a difference of some 40%. Note that Fig. (5) shows only the historical correlation between two instruments. All of the actual correlations between instruments were calculated and subsequently used for the calculation of VaR.
Having obtained all the constituents for the VaR calculation (Equation (1)), the effect of the temporal evolution of $\tilde{\lambda}$ on VaR may be investigated.

Effect of changing market conditions on VaR: Fig. 6 shows the effect of a changing $\tilde{\lambda}$ on VaR. VaR was calculated using Equation (2) with CI = 1 and T = 1 (since these are merely scaling effects and do not affect the shape of the curves, only the magnitude) and with a nominal value of R1 per instrument for each of the instruments indicated in Section 3.4.4 (hence the value of VaR (on the *y*-axis of Fig. (6)) may also be thought of as a percentage). VaR was calculated using (i) the equally weighted volatility in conjunction with the equally weighted correlations between return series, (ii) EWMA with $\tilde{\lambda}$ constant at 0.94 (and corresponding exponentially weighted correlations) and (iii) EWMA with $\tilde{\lambda}$ variable (and corresponding exponentially weighted correlations) for each day from 1 January 1998.
The VaR calculated according to each of the methods differs substantially. As with the volatilities, VaR is similar for each technique during periods of low volatility. During a high volatility burst, i.e. at the time VaR estimation is most important, the values differ significantly. During the September 1998 period, VaR($\tilde{\lambda}$) differs from VaR($\lambda = 0.94$) by as much as +31%. On the other hand, the VaR decreases more quickly for VaR($\tilde{\lambda}$) than it does for VaR(λ 0.94) as seen in Fig. (6) from March 1999 until January 2000. Since $\tilde{\lambda}$ is constantly changing, the exponential envelope responds by adjusting dynamically and reacting to new market information. This study examined the time dependence of the parameter $\tilde{\lambda}$, used to establish the shape of the exponential envelope applied to price returns. This application assumes $\tilde{\lambda}$ to be static and weighs recent return data more heavily than those from the distant past. Thus, the technique to estimate volatility that employs this weighting scheme enjoys the advantage of a faster reaction speed to changing market conditions than methods that weigh returns equally. An improvement in the volatility estimate leads directly to an improvement in the variance/covariance VaR technique since this method employs volatility directly.

It was found that volatile market data affects the value of $\tilde{\lambda}$ significantly and that the resultant effects on the volatility estimate are non-trivial. A dynamically changing − rather than a static $\tilde{\lambda}$, then, has an ultimate effect on VaR (through the changes it effects on volatility and correlation between instruments). By using a daily-calculated value of $\tilde{\lambda}$ adjustments may be made daily that take account of the constantly evolving marketplace, and ultimately lead to more accurate macro variables such as VaR.
The changing duration of the instruments used in this study was not taken into account and future research could incorporate this important aspect. A comparison of back-testing results (using the different techniques discussed) from an actual portfolio which was dynamically hedged and managed over a period of time, will help establish the validity of the arguments presented here.

References

Alexander, C., 1997. Volatility and Correlation Forecasting. In: Alexander, C. (Ed.), The Handbook of Risk Management and Analysis. Chichester: John Wiley and Sons.

Basle Committee on Banking Supervision, 1996. Amendment to the Capital Accord to Incorporate Market Risks. Bank for International Settlements.

Bernstein, P., 1999. Risk at the Roots. Balance Sheet, 7, 1: 18 - 20.

Blaauw, A., 1997. ABSA ALCO Department, Internal Report #5, Johannesburg: ABSA Bank, 13 - 18.

Botha, M., (Marius), 2000. Old Mutual Asset Managers, Internal Report #23, Cape Town, Old Mutual, 4 – 9.

Botha, M., (Martin) and Van Vuuren, G. W., 1999. The Silence of the Lambdas. In: Crabbe, M. (Ed), South Africa: A Risk Special Report, Risk, 12, 22 - 24.

Christoffersen, P. and Diebold, F., 1997. How Relevant is Volatility Forecasting for Financial Risk Management? Wharton Financial Institutions Center, University of Pennsylvania: working paper.

Diebold, F., Hickman, A., Inoue, A. and Schuermann, T., 1998. Scale Models. Risk, 11, 1, 104 - 107.

Dowd, K., 1998. Beyond Value at Risk. New York: John Wiley and Sons.

Drost, F. and Nijman, T., 1993. Temporal Aggregation of GARCH processes. Econometrica, 61, 312 – 321.

Hull, J. C., 1997. Options, Futures and Other Derivatives. 3rd Edition, University of Toronto: Prentice Hall International Inc. JP Morgan, 1996. RiskMetrics™ Technical Document, Fourth Edition, New York.

Leong, K., 1996. The Right Approach. In: Cooper, G. (Ed.), Value-at-risk: A Risk Special Supplement, June, 9 -14.

Picoult, E., 1998. Calculating Value-at-Risk With Monte Carlo Simulation. In: Dupire, B. (Ed.), Monte Carlo: Methodologies and Applications for Pricing and Risk Management, Risk Books, Somerset: Bookcraft Books, 209 - 230.

Reed, N., 1996. Variations on a Theme. In: Cooper, G. (Ed.), A Risk Special Supplement: Value at Risk. Risk, 6, 11, 2 - 4.

Shimko, D., 1996. VaR for corporates. Risk, 9, 6, 34 - 35.
Shimko, D., 1997. Accentuate the Positive. In:
Paul-Choudhury, S. (Ed.), Value at Risk for End-Users, Risk Special Supplement, March, 10 - 14.

Smithson, C. and Minton, L., 1996a. Understanding Various Ways to Calculate VaR. Risk, 9, 1, 54 - 56.

Smithson, C. and Minton, L., 1996b. The Debate on the Use of VaR. Risk, 9: 58 - 60.

Wei, W. W. S., 1989. Time Series Analysis: Univariate and Multivariate Methods. USA: Addison Wesley Publishing Company.

Zangari, P., 1996. An Improved Methodology for Measuring VaR. RiskMetrics™ Monitor, First quarter Reuters/JP Morgan.

Zangari, P., 1997. Catering for an Event. Risk, 10: 34 - 36.