Values at Risk

A simplified approach to the conditional estimation of value at risk, by

Giovanni Barone-Adesi *and* **Kostas Giannopoulos**

To estimate VaR requires an estimation of portfolio volatility. But the historical volatility of a bank portfolio is an ill suited measure of its current volatility, because investment weights may change rapidly and individual securities' volatility may shift over time. Moreover, to consolidate the volatilities of individual components into portfolio volatility requires a correlation matrix of returns, itself subject to shifts over time. And, even if the correlation matrix was constant, the effort required to estimate it in a multivariate time-series framework would pose a computational challenge.

A simple procedure for estimating current portfolio volatility is to construct the hypothetical return series the portfolio would have earned if it had been kept constant at its current weights in the past. (Options may be accounted for by substituting the products of their current delta multiplied by the volatility of their underlying assets.)

The resulting time series of portfolio returns is then analysed to identify the best fitting time-series model. Accurate point estimates of current volatility are then produced and VaR is computed from them.

Let R_t be the Nx1 vector $(R_{1,t}, R_{2,t}, \ldots, R_{n,t})$ where $R_{i,t}$ is the return on the ith asset over the period (t-1,t) and let W be the Nx1 vector of the portfolio weights over the same period. The historical returns of our current portfolio holdings are given by:

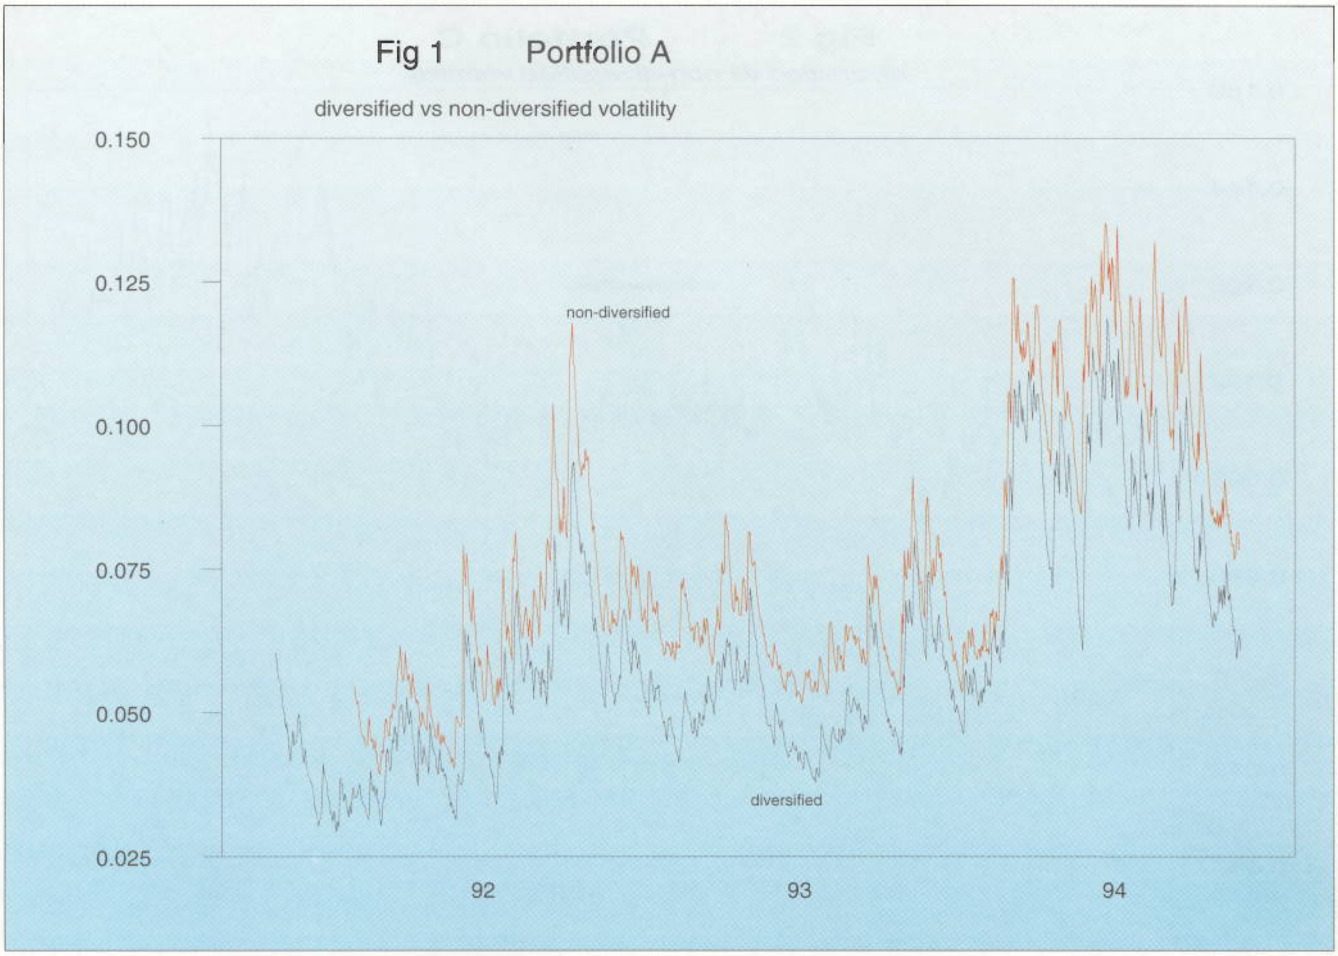

$$Y_t = W^T R_t \quad \textbf{(1)}$$

W represents investment holdings, whether actual or hypothetical, and the series Y the path of portfolio returns through history. Following Markowitz (1956) the portfolio's risk and return trade-off can be expressed in terms of the statistical moments of the multivariate distribution of the weighted investments as:

$$E(Y_t) = E(W^T R) = m \quad \textbf{(2.1)}$$

$$var(Y_t) = W^T \Omega W = \sigma^2 \quad \textbf{(2.2)}$$

where Ω is the unconditional variance-covariance matrix of the returns of the N assets.

A simplified way to find the portfolio's risk-and-return characteristics is by estimating the first two moments of Y.

$$E(Y_t) = m \quad \textbf{(3.1)}$$

$$var(Y_t) = E(Y_t - m)^2 = \sigma^2 \quad \textbf{(3.2)}$$

Hence, if historical returns are known, the portfolios mean and variance can be found as in (3.1), (3.2). This is easier than (2.1), (2.2) and still yields the same results.

The method in (3.1) and (3.2) can easily be deployed in risk management to compute the value at risk at any given time t. For every change in W, however, the series of past returns, Y, needs to be reconstructed and σ^2, the volatility of the new position, needs to be re-estimated as in (3.2).

Time-varying risk

The portfolio risk-return estimates given rely upon a very strong assumption; that the series of returns, Y, is stationary. That means that both m and σ^2 do not change over the measurement period. Yet several studies have concluded that asset variances and covariances are not constant but change over time, eg Christie 1982.

A number of solutions have been proposed, covering how best to estimate current variances and covariances. Perhaps the most popular method is exponential smoothing (ES), proposed by JP Morgan. More sophisticated (but also computationally more demanding[1]) is Garch, based on the work of Engle (1982) and Bollerslev (1986).

Because of the huge dimensions that a variance-covariance matrix may have, both methods seek first to partition this matrix into (N-1)N/2 off-diagonal elements and then to capture the joint dynamics of the second moments for each possible pairwise combination of investment holdings. The volatility of current investment holdings is then computed as in (2.2). The problems both methods face stem from the way they partition the variance-covariance matrix. Unless certain preconditions are satisfied, there is no guarantee that the resulting variance-covariance matrix comes from a NxN multivariate distribution. Hence the portfolio variance estimates are very likely to be biased.

Fig 2 Portfolio C

diversified vs non-diversified volatility

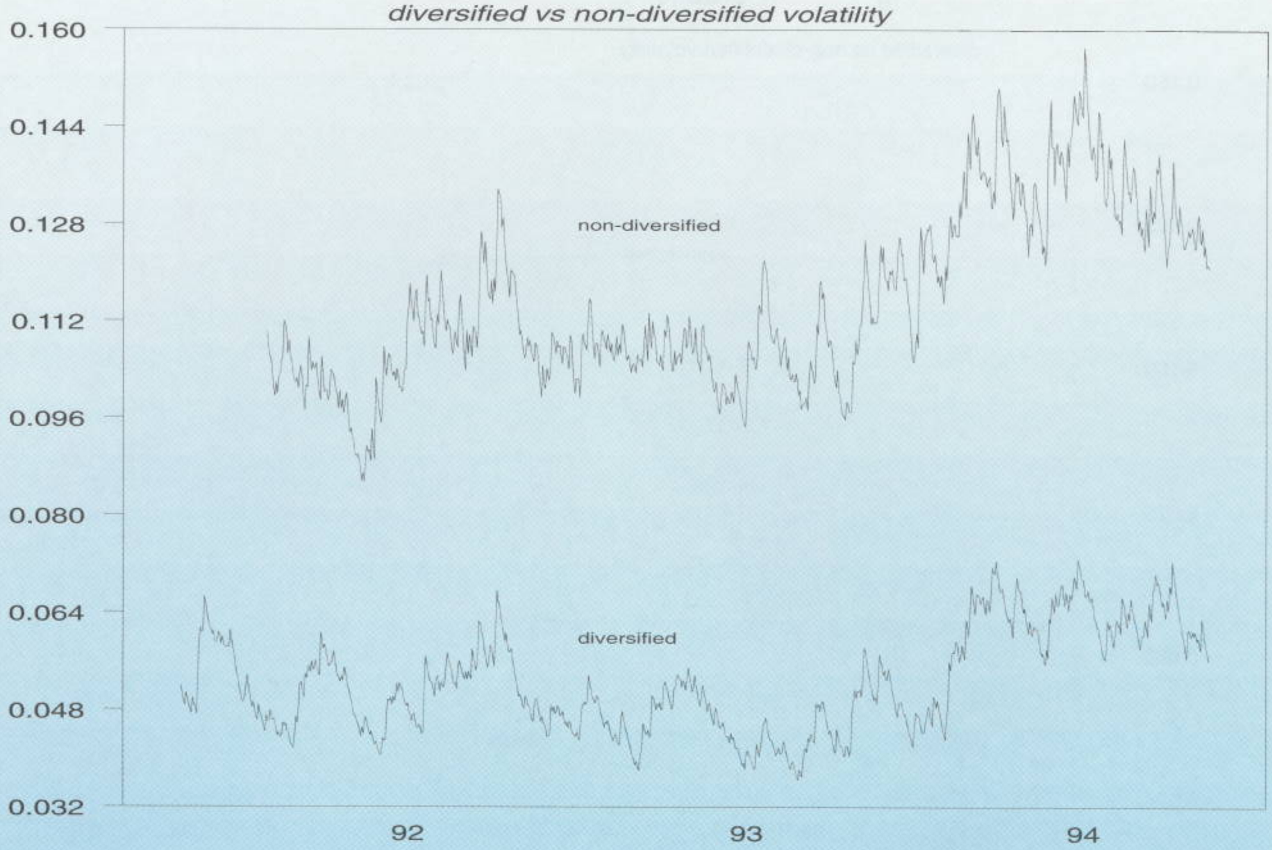

A simplified approach

Our simplified approach to computing portfolio VaR aims to overcome both of the above problems –non-stationarity and dimensionality – while remaining unbiased on volatility. We believe that past returns contain all necessary information about the current portfolio's risk-return trade-off. And, in order to estimate portfolio volatility, it is sufficient to study the portfolio's returns rather than those of its components.

It is very likely that the volatility of most individual assets included in a portfolio will change over time, particularly if returns are measured with high frequency, ie daily. And, if the constant volatility hypothesis is rejected estimates computed by (2.2) or (3.2) cease to be reliable.

We can compute portfolio Y's volatility as time-varying by treating past returns as time series in their own right. This approach has many advantages. It is simple, easy to compute and overcomes the dimensionality and bias problems that arise from estimating the covariance matrix. At the same time, the portfolio's past returns contain all the necessary information about the dynamics that govern the aggregate current investment holdings and we should really make the best use of this information[2]. For example it might be possible to capture the time path of portfolio volatility using a Garch model. This hypothesis is based on the fact that most high-frequency security returns have been found to contain volatility clusters.

To illustrate our procedure we collected daily data for assets with different risk exposure and we constructed three hypothetical portfolios, as in *table 1*. We then employed Garch methodology and stress analysis on the portfolio return to study its riskiness. We generated constant-weighted portfolios for the period 1 November 1991-15 November 1994. Futures contracts have been rolled to create a single series. Missing observations and bank holidays have been set equal to a smoothed value[3]. When a futures contract was rolled, the first observation was considered as missing and so was set equal to its smoothed value. The three portfolios we constructed had the weights as shown in table 1.

For a portfolio diversified across a wide range of assets (such as portfolio C), the non-constant volatility hypothesis is an open issue. The LM test (regressing the squared residuals of an autoregressive process against their own lagged values) can be used to verify whether there are any Garch effects. Diagnostic test results allow us to conclude that the implemented Garch parameterisation, although it has been very general and simple, has removed the Garch effects from the portfolio.

Figures 1 and *2* illustrate how the daily annualised standard deviation of portfolios A and C behaved over the tested period. The upper line shows the volatility of an undiversified portfolio, wherein all pair-wise correlation coefficients are 1.0. (The

TABLE 1 PORTFOLIO COMPOSITION (%AGE WEIGHTS)

Portfolio	Bonds		Equity		Commodities	
A	Italian	40				
	German	30				
	L Gilt	30				
B	Italian	30	S&P 500	15		
	German	20	FT-SE 100	15		
C	Italian	15	S&P 500	15	Oil	10
	German	15	FT-SE 100	15	Cocoa	8
	Gilt	10			Copper	6
					Alumin	6

Fig 3 Portfolio C

Stress analysis

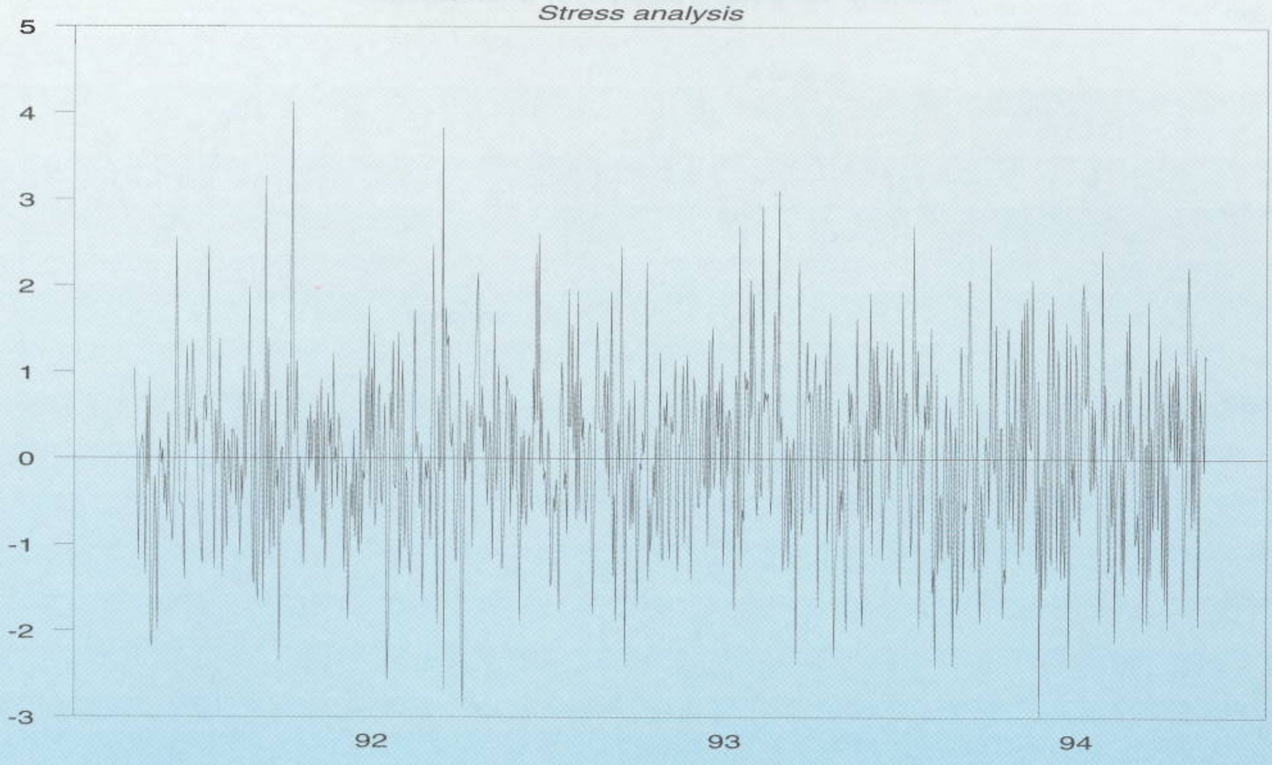

volatility charted is simply the weighted average of conditional volatilities of assets in the portfolio.) Because portfolio C is diversified, its volatility oscillates less (between 3.65% and 7.29%) than portfolios A and B.

There are three useful products of our methodology. The first is a simple and accurate measure for the volatility of the current portfolio. This is achieved without using computationally intense multivariate methodologies. The second is the possibility of comparing a series of volatility patterns similar to *figures 1* and *2* with the historical volatility pattern of the actual portfolio with its changing weights. This comparison allows an evaluation of the manager's ability to time volatility. Timing volatility is an important component of performance, especially if expected security returns are not positively related to current volatility levels. Finally, the possibility of using the Garch residuals on the current portfolio weights allows for the implementation of meaningful stress-testing procedures. We will focus on stress testing and the evaluation of correlation risk because of their importance in risk management models.

Stress analysis

The series innovations affecting the volatility of one of the portfolios is exhibited in *figure 3*. It is apparent that the distribution of the innovations is not normal, with values reaching up to four standard deviations for the most diversified portfolio, C. Negative innovations are more modest, ranging up to four. Worst-case scenarios for stress analysis may be built applying the largest outliers in the innovation series to the current Garch parameters. This exercise simulates the effect of the largest historical shock on the current market conditions. Thus, to stress our portfolios, it is not necessary to choose between the largest shocks for the different securities because the most interesting shocks are a direct by-product of the Garch estimation of portfolio volatility.

Correlation and diversification

Conditional VaR models which use the quadratic equation (2.2) to update portfolio volatility, eg RiskMetrics, need first to estimate all the possible pair-wise covariances. In a widely diversified portfolio, containing say 100 assets, there are 100 variances and 4950 conditional covariances to be estimated. Furthermore, any model used to update the covariances must keep the multivariate features of the joint distribution. With a large matrix like that, one is unlikely to get unbiased estimates for all the 4950 covariances and at the same time to guarantee that the joint multivariate distribution still holds. Clearly, errors in covariances as well in variances will affect the accuracy our portfolio's VaR estimate and lead to the wrong risk management decisions.

Our approach estimates conditionally the volatility of only one, univariate, time series – the portfolio's return – overcoming all of the above problems, yet measuring in full the changes in assets' variances and covariances. Moreover, it discloses the impact that the overall changes in covariances have on the portfolio volatility. It can tell us in what proportion an increase/decrease in the portfolio's VaR is due to changes in asset variances or correlations. We call this type of analysis correlation stability.

Each correlation coefficient is subject to changes at any time. However, changes across the correlation matrix might themselves not be correlated and therefore their impact on the overall portfolio risk may be diminished. Our conditional VaR approach allows to attribute any changes in the portfolio's conditional volatility to two main components; changes in asset volatilities and changes in asset correlations. If h_t is the portfo-

Fig 4 Correlation stability

volatility ratio (diversified over undiversified)

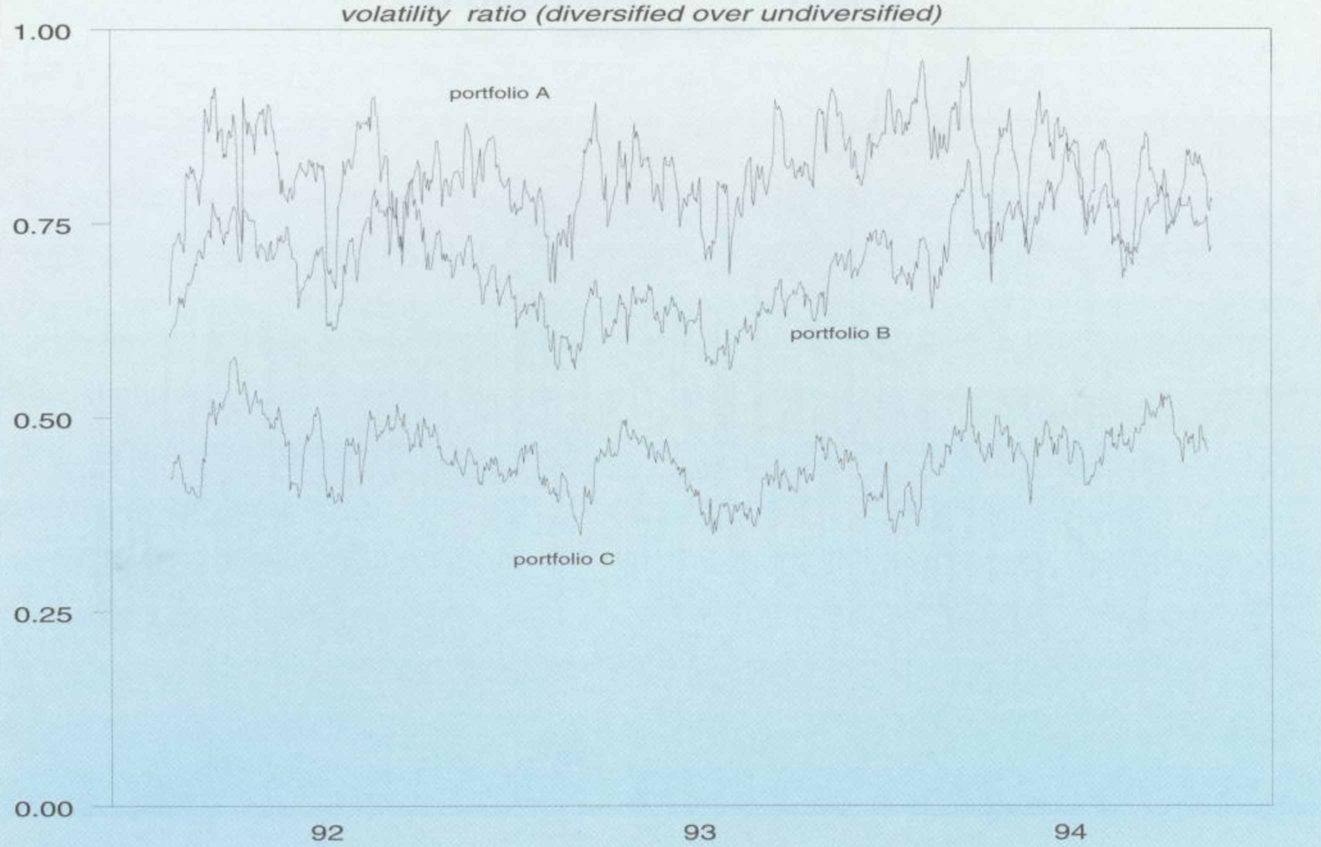

lio's conditional variance, as estimated in (4.2), its time-varying volatility is $\sigma_t=\sqrt{h_t}$. This is the volatility estimate of a diversified portfolio at period t.

By setting as equal to 1.0 all the pair-wise correlation coefficients in each period, the portfolio's volatility becomes the weighted volatility of its components. Conditional volatilities of the individual asset components can be obtained by fitting a Garch-type model for each return series. We note the volatility of this undiversified portfolio as s_t. The quantity $(1-\frac{\sigma_t}{s_t})$ tells us in what proportion the portfolio volatility has been diversified away because of non-perfect correlations. If that quantity does not change significantly over time, then the overall effect of time-varying correlations is invariant and we have correlation stability.

Figure 4 shows how correlation stability improves for more diversified portfolios. Portfolio A, containing only bonds, is subject to greater correlation risk because of the tendency of bonds to fall in step in the presence of large market moves. Risk managers who relied on average standard deviations would be surprised by the extreme volatility values our bond portfolio might produce in a crash.

Our conditional volatility estimates provide early warnings about this risk increase and therefore are a useful supplement to existing risk management systems. ❖

Giovanni Barone-Adesi [gbarone@gpu.srv.ualberta.ca] teaches at the Faculty of Business, University of Alberta (Canada), and City Business School (UK); and Kostas Giannopoulos [100331.154@compuserve.com] at the Faculty of Business, University of Westminster (UK). Free software for the calculations used in this paper is available from the authors at the e-mail addresses given

1 Giannopoulos and Eales, 1996

2 If W is known a priori, the portfolio's (unconditional) volatility can be computed easily, as in (3.2).

3 The downhill simplex algorithm was used to find the optimal smoothing coefficients for a variety of specifications. Then we selected the smoothing model that minimised the Schwarz criterion.

REFERENCES

- Giovanni Barone-Adesi (1994) 'ALM in Banks', working paper, University of Alberta.
- Bollerslev T (1986), 'Generalised Autoregressive Conditional Heteroskedasticity', Journal of Econometrics, 31, 307-28.
- Christie A (1982), 'The Stochastic Behaviour of Common Stock Variance: Value, Leverage and Interest Rate Effects', Journal of Financial Economics, 10, 407-432.
- Engle R (1982), 'Autoregressive Conditional Heteroskedasticity with Estimates of the Variance in the UK inflation', Econometrica, 50, 987-1008.
- Giannopoulos K and B Eales (1996), 'Educated Estimates', Futures and Options World, April, 45-47.
- Mandelbrot B (1963), 'The Variation of Certain Speculative Prices', Journal of Business, 36, 394-419.
- JP Morgan (1995) RiskMetrics