Testing value-at-risk approaches to capital adequacy

By Patricia Jackson and William Perraudin of the Bank's Regulatory Policy Division and David Maude of the Bank's Monetary Assessment and Strategy Division.

This article[(1)] looks at the nature of whole-book value-at-risk models, and describes how the Bank of England set out in 1995 to assess their performance in accurately predicting risk and in providing a basis for reliable trading-book capital calculations.

Introduction

In the past three years, a revision to the Basle Accord and new EU Directives have radically changed the method for calculating capital to back the trading books of banks. The 1988 Basle Accord applied a credit-risk capital treatment to both the banking and trading books of banks—in other words, not only to loans, but also to readily tradable items such as securities. This credit-risk approach had a number of drawbacks when applied to trading books, and in 1988 work was started by both the Basle Committee on Banking Supervision and the European Union to find an alternative approach.

Two new methodologies were developed. The first new 'standard' approach rested heavily on the risk-based capital weights already applied by some supervisors of securities firms in the United Kingdom and the United States. These fixed weights were based on calculations of price volatility for different types of security. In the case of bonds, for example, the weights varied substantially for different maturities. But this approach did not provide a way of reflecting diversification benefits across whole trading books, which was important for the largest firms.[(2)] The solution was to develop an alternative methodology, based on the internal value-at-risk (VaR) models that had been developed by the largest firms as a management tool to assess risk on whole trading-book portfolios.

The move towards a risk-based approach to calculating capital for the trading books of banks was clearly an important development. The credit-risk approach did not enable hedges within the trading book to be recognised, nor did it take into account short positions or positions in government bonds (although in the United Kingdom, a requirement was introduced for the latter). It also did not enable the influence of maturity on the price volatility of interest rate items to be recognised. However, an important question for the Bank of England was whether the new approaches (both the standard and the VaR) would actually deliver adequate capital, given potential trading-book losses.

For the standard approach, which was incorporated in the Capital Adequacy Directive (CAD),[(3)] the issue was particularly acute with regard to the UK gilt-edged market makers (GEMMs). These central players in the gilts market were required to have specialist books, limited to sterling bonds. The additive structure of the standard approach (according to which risk-based capital requirements are calculated market by market, by type of risk, before being summed) generates a sizable cushion of capital for diversified books, but the Bank wanted to know whether it would generate an adequate cushion of capital for specialist books such as those of the GEMMs.

In 1994, the Bank conducted studies in which the profits and losses on actual GEMM books were simulated over daily and weekly periods back to 1988. The results were compared with the capital that would have been required by the CAD. The conclusion reached was that, although the CAD did generate capital sufficient to cover 99% of weekly losses, in some periods this would not have provided a sufficient cushion. Because of this, when the CAD was introduced, the requirement for these firms was increased, to 125% of the CAD general market risk standard.

When VaR models were proposed as a way of capturing the effect of risk diversification in trading books, a similar question arose for the Bank, of whether this approach would deliver sufficient capital relative to the losses that might be experienced on actual books. In particular, it was important to know whether VaR models could predict losses accurately.

The Basle approach to models

Under both the 1997 Amendment to the Basle Accord and the Second Capital Adequacy Directive adopted by the European Union, banks can choose whether to use the standard approach to calculating capital requirements for trading books (equities, interest rate instruments, foreign exchange and commodities), or to seek supervisory approval to employ their own in-house VaR models as the basis for

(1) Based on the Bank's *Working Paper* No 79.
(2) The United Kingdom had developed a method of allowing for diversification of equity books for securities firms, but this was not adopted by the European Union or the Basle Committee on Banking Supervision.
(3) The European Commission's Capital Adequacy Directive, agreed in 1993 and introduced at the start of 1996, which established EU minimum capital requirements for the trading books of banks and securities firms.

the capital calculation. Even in the standard approach, models are employed to a limited extent to enable some positions to be correctly processed for inclusion in the standard methodology—this is particularly true of options positions. The alternative approach, however, relies entirely on internal models.

Exclusive reliance on models raises questions about necessary safeguards to ensure that the capital requirements generated are adequate. Basle addressed this in a number of ways. One was to lay down simple standards for the construction of the models. For example, models must be formulated to yield a value-at-risk estimate that will not be exceeded on more than 1% of occasions. The losses must be calculated for a ten-day holding period, and at least twelve months of returns data must be used. Basle does not, however, prescribe the type of model to be used. Basle also included a substantial additional buffer, by requiring banks to hold capital equivalent to the higher of (i) the VaR number yielded by the model or (ii) three times the 60-day moving average of the VaR numbers generated on the current and past books. In addition to these quantitative safeguards, Basle also included a number of qualitative safeguards, for example that the model had to be part of the bank's own risk-measurement system, and that stress testing had to be carried out on the portfolios to look at extreme losses.

As a check on the accuracy of the models, the supervisors carry out back-testing—a comparison of actual trading results with model-generated risk measures. This may pose problems, first because trading results are often affected by changes to portfolios in the period following the calculation of the VaR. Because of this, Basle has urged banks to develop their own capability to perform back-tests, using the losses that would have been made if the book had been held constant over a one-day period. Second, as Kupiec (1995) argues, back-testing requires a large number of observations in order to make a judgment about the accuracy of the model. Nevertheless, back-testing and some kind of penalty are essential to provide incentives for firms to increase the accuracy of the models. Firms that do not meet the back-testing criterion for accuracy suffer additional capital charges (see below).

Value-at-risk analysis

The typical VaR models developed by the firms for their internal risk-management purposes attempt to measure the loss on a portfolio over a specified period (often the next 24 hours) that will only be exceeded on a given fraction of occasions (usually 1% or 5%). Two broad types of VaR analysis are used:

(i) under *parametric* VaR analysis, the distribution of asset returns is estimated from historical data, under the assumption that this distribution is a member of a given parametric class. The commonest approach is to suppose that returns are stationary, joint normal and independent over time. Using estimates of the means and covariances of returns, it is then possible to calculate the loss in a one-day holding period that will be exceeded with a given probability; and

(ii) the *simulation* approach to VaR analysis consists of calculating the losses that would have been experienced on a particular portfolio in previous 24-hour periods (using a run of historical returns data) and finding the loss that is exceeded on a given percentage of days in the sample. As a non-parametric procedure, this approach imposes no assumptions about the distribution of returns, other than that they are independent over time.

Testing the VaR models

Before the amendment to the Basle Accord had been agreed, we tested what the VaR models delivered, by taking data on actual trading books from a bank with sizable trading exposure, covering equity, interest rate and foreign exchange risk (see the boxes on pages 258 and 259 for details). We examined the impact of window length (ie the length of the period from which returns data are taken for the models) and the effect of weighting returns data in the parametric VaR calculations. We also compared the empirical performance of parametric and simulation-based VaR models when used to calculate the possible losses on these books.(1)

A finding of considerable practical significance was that the various approaches to VaR modelling differ widely in the accuracy with which they predict the fraction of times a given loss will be exceeded. In this respect, simulation-based were better than parametric VaR techniques. This is clearly important when these models are used to generate capital requirements. On the other hand, parametric VaR analysis tracks the time-series behaviour of volatility better than simulation-based techniques, and appears to yield slightly superior volatility forecasts. However, with well-diversified fixed-income books, the gains in forecasting accuracy are relatively slight.

Finally, we investigated the size of buffer that would come out of the Basle requirement that capital must exceed the higher of the current VaR or three times the average VaR of the previous 60 days.

Parametric VaR analysis

The first question that we addressed is how sensitive parametric VaR analyses are to the way in which the volatilities are estimated. The approach to volatility estimation typically used in VaR applications is to take a weighted average of the squared deviations of each return from an estimate of the mean return, using a window of past data. So if r_t is the holding return at t, a typical estimator for $\sigma^2 = VaR(r_t)$ would be:

(1) A significant omission in our study was the treatment of derivatives in VaR models.

Portfolio data

The main advantage of using actual books for the predominant bank trading risks is that it ensures that the pattern of risk exposures along the yield curve and between markets is realistic. The amount of exposure taken at different points on the yield curve and between markets clearly reflects a bank's investment decisions. Randomly generated portfolios are unlikely to be representative, and it would be difficult to build stylised books that were representative without basing them on actual books.

The table shows the breakdown of the four different books that we employed in our statistical analysis. The first three portfolios were those held by a bank with a sizable trading operation in three consecutive months. In the table, the foreign exchange exposure for a particular currency represents the total net sterling value of assets denominated in that currency. So for example, if the bank acquires a ten-year Deutsche Mark-denominated bond, both the foreign exchange exposure and the six to ten-year bond categories in the Deutsche Mark column of the table increase.(1)

Two features of the data stand out. First, the degree to which the bank's fixed-income exposure fluctuates over relatively short periods of time is quite striking. This fact underlines the importance of banks satisfying capital requirements for market risk almost continuously. VaR models need to be run daily. Second, the bank's net foreign exchange exposure is relatively small, except for the large short dollar position in portfolio 4. This suggests that the bank is systematically hedging the net foreign exchange risk in its trading book.(2) Other data that we saw suggest that the months we chose were fairly typical of the bank's general behaviour, in that foreign exchange risk is systematically hedged, whereas other exposures fluctuate considerably.

Portfolio amounts

£ millions

Portfolio 1

	FFr	£	$	Yen	DM
FX	-10.89	n/a	-46.02	4.31	40.95
3–12 months	24.04	56.82	-191.56	-590.78	462.35
2–5 years	-11.45	-336.42	83.13	1,247.51	-139.10
6–10 years	-3.52	-14.62	69.96	-65.45	-144.32
11+ years	0.00	0.00	-3.19	5.52	-41.66

Portfolio 2

	FFr	£	$	Yen	DM
FX	-5.95	n/a	5.72	-22.23	10.20
3–12 months	64.96	40.01	-135.10	-529.87	629.00
2–5 years	-130.29	-268.84	-33.18	1194.70	-178.89
6–10 years	19.39	11.17	0.93	-58.66	-107.47
11+ years	0.00	0.00	-2.17	5.20	-8.76

Portfolio 3

	FFr	£	$	Yen	DM
FX	-9.86	n/a	33.50	-5.59	22.48
3–12 months	-237.72	105.39	4.56	-1314.62	11.69
2–5 years	43.46	-245.85	11.11	346.49	89.64
6–10 years	39.53	22.44	0.26	-58.31	-69.96
11+ years	0.00	-26.70	-2.72	-4.75	-8.81

Portfolio 4

	FFr	£	$	Yen	DM
FX	28.51	n/a	-132.10	11.84	-26.08
3–12 months	-11.00	2.22	-153.15	-341.36	-327.05
2–5 years	-160.38	13.88	24.53	357.72	559.87
6–10 years	179.83	-53.34	53.92	40.87	-298.86
11+ years	43.13	39.72	29.99	0.00	0.00
Equities	1.50	2.81	-37.69	6.06	8.24

n/a = not applicable.

Most of our data on the bank's portfolio consisted of interest rate exposures in different currencies. But it is also important to examine whether VaR analysis performs differently when applied to portfolios containing equities, rather than only fixed-income and foreign exchange positions. The bank also provided us with data on a single additional portfolio, here labelled portfolio 4, which contained equity exposures. The relatively small size of this equity book is typical of what most banks hold.

(1) The practice of considering the exchange rate and foreign currency price risks separately is common among practitioners.
(2) The exposures were the consolidated exposures for the bank and its securities companies, and therefore this did not simply reflect the effect of the Bank of England's guideline on overnight FX exposures that applies to the bank.

$$\hat{\sigma}_t^2 = \frac{1}{T-1}\sum_{i=0}^{T-1}\lambda_i (r_{t-T+i} - \bar{r}_t)^2 \quad \textbf{(1)}$$

where $\lambda_i \in [0,1]$, $\sum_{i=0}^{T-1}\lambda_i / T = 1$, and $\bar{r}_t \equiv \sum_{j=0}^{T-1} r_{t-T+j} / T$

In implementing the VaR models, we worked out the returns for one-day or rolling ten-day holding periods on a given portfolio, and then calculated volatilities, tail probabilities etc, using that single series.(1)

Three choices must be made in implementing the parametric VaR model described above, namely (a) an appropriate length for the data 'window', (T); (b) the weighting scheme to be adopted, $(\lambda_0, \lambda_1, \ldots \lambda_{T-1})$; and (c) whether the mean should be estimated using the sample mean, $\sum_{j=0}^{T-1} r_{t-T+j} / T$, or set to zero as some empirical researchers have advocated.(2)

(a) Window length

Table A shows two ways of assessing the sensitivity of the VaR results to the choice of T. In the upper block of the table, we show the mean absolute forecast error, where we define the forecast error at period t as:

$$\left| |r_t - \bar{r}_t| - \hat{\sigma}_t \right| \quad \textbf{(2)}$$

(1) This approach yields results that are arithmetically identical to those one would obtain if one estimated a full covariance matrix for n individual asset return series (Σ), and then estimated the volatility of a portfolio with portfolio holdings, $a \equiv (a_1, a_2, \ldots, a_n)'$, by calculating the quadratic form, $a'\Sigma a$. The latter approach is taken by practitioners, including JP Morgan in their *RiskMetrics* system, and is more efficient if one has many portfolios for which one wants the value at risk on a single date. When a large number of VaR calculations are required for a small number of portfolios on different dates, our approach is quicker.
(2) See, for example, J P Morgan (1995), page 66.

Table A
Parametric VaR models: window length

		3 months' data	6 months' data	12 months' data	24 months' data
		Mean absolute forecast error			
Portfolio 1	Mean	26.71*	26.79	27.02	27.12
	Standard error	(0.85)	(0.73)	(0.64)	(0.60)
	t-statistic	*n/a*	*0.20*	*0.57*	*0.79*
Portfolio 2	Mean	17.26*	17.32	17.40	17.29
	Standard error	(0.55)	(0.47)	(0.42)	(0.41)
	t-statistic	*n/a*	*0.21*	*0.39*	*0.08*
Portfolio 3	Mean	5.43	5.42	5.44	5.40*
	Standard error	(0.21)	(0.17)	(0.15)	(0.14)
	t-statistic	*0.23*	*0.18*	*0.72*	*n/a*
Portfolio 4	Mean	77.12*	78.11	78.10	78.60
	Standard error	(2.10)	(1.85)	(1.78)	(1.72)
	t-statistic	*n/a*	*0.89*	*0.68*	*0.99*

Tail probabilities

	3 months' data	6 months' data	12 months' data	24 months' data
Portfolio 1	1.71	1.38	1.32*	1.32*
Portfolio 2	2.11	1.91	1.58	1.51*
Portfolio 3	1.58	1.32	1.45	1.25*
Portfolio 4	1.71	1.65	1.71	1.28*

n/a = not applicable.

Notes: Calculations use equal weights ($\lambda_i = 1 \; \forall \; i$), zero means and daily returns.
Forecast errors are scaled up by 10,000.
Asterisks indicate lowest in the row.
Newey-West standard errors are in parentheses.
T-statistics are given for the difference from the lowest mean absolute error in the same row.

Averaging the absolute forecast errors over the entire sample period yields a measure of the accuracy of the volatility estimates. Standard errors are reported in parentheses under each mean. These are calculated using the technique of Newey and West (1987), and so are robust to complex patterns of time dependence. The standard errors give a very conservative impression of the statistical significance of differences in mean forecast errors, since means calculated under different assumptions are highly positively correlated, reducing the variability of the average difference. So we also give the t-statistics for the difference between each mean absolute forecast error and the lowest mean in the same row of the table. The t-statistics are also calculated using the Newey-West technique.

Note that we tried working with various other measures of forecast accuracy. First, one may define the forecast error as $\left|(r_t - \bar{r}_t)^2 - \hat{\sigma}_t^2\right|$, and then employ the sample mean of these absolute differences. In this case, one is evaluating forecasts of the instantaneous variance rather than the instantaneous standard deviation. Since VaR calculations employ the latter, this is probably not appropriate. Second, we experimented by using root mean squares of the forecast errors instead of simply means. The problem with this approach is that it attributes more weight to outliers. We thought it better, therefore, to use means. In the lower block of Table A, we provide measures of the degree to which capital requirements based on different VaR models cover losses that occur with a given probability. Assuming normally distributed returns, one may deduce from the time series of estimated volatilities a corresponding series for what we shall call '1% cut-off points', meaning the loss that, according to the model, will be exceeded on average

Returns data

The bond returns employed in our study were based on a time series of zero-coupon yield curves calculated by an investment bank (not the one that supplied us with portfolio data). From this, we calculated holding returns for the maturity categories on which we had portfolio data. For equities, we employed the returns on the French CAC 40, the UK FT All-Share, the German DAX, the US S&P Composite and the Japanese Nikkei 225. Including equities and foreign exchange positions meant that in total we were dealing with 79 different sources of risk. All returns were calculated as changes in log prices.

Throughout the analysis, we took sterling as the base currency and employed data from July 1987 to April 1995. The table below shows the annualised sample standard deviations of the daily returns on our 79 asset categories. The figures in the table suggest that returns on fixed-income books are much less volatile than returns on equities, unless the fixed-income portfolio includes very long-dated securities. Even holdings heavily weighted towards long-dated bonds will have relatively low average durations, and so are likely to exhibit lower volatilities than portfolios that include equities or foreign exchange. Although the returns data covered the period July 1987 to April 1995, estimates of the VaRs were made only for the period from June 1989. Data from the earlier period were used in whole or in part (depending on the length of the data window) to construct the first VaR estimate. This meant that it was not possible to compute a VaR estimate for the 1987 equity market crash, although the crash did appear in the past data when VaR estimates were calculated using a 24-month window.(1)

Standard deviations of daily returns

	FFr	£	$	Yen	DM
FX	6.32	n/a	10.74	10.00	6.63
<3 months	0.90	0.48	0.31	0.22	0.25
3–6 months	1.09	0.86	0.53	0.34	0.45
6–9 months	1.31	1.32	0.83	0.53	0.67
9–12 months	1.49	1.76	1.16	0.70	0.88
1–2 years	2.63	3.33	2.09	1.30	1.72
2–3 years	3.62	4.42	3.10	1.95	2.27
3–4 years	4.59	5.53	4.13	2.67	2.93
4–5 years	5.58	6.57	5.15	3.43	3.50
5–6 years	6.65	7.55	6.14	4.36	4.06
6–7 years	7.99	8.55	7.13	5.62	4.97
7–8 years	9.36	9.80	8.13	6.73	6.19
8–9 years	10.15	10.97	9.08	7.66	7.34
9–10 years	10.40	12.05	9.94	8.43	8.53
11+ years	11.45	13.66	11.63	10.09	10.50
Equities	19.48	14.24	16.51	22.43	20.02

n/a = not applicable.

Note: Standard deviations are annualised (multiplied by $\sqrt{250}$) and in per cent.

(1) This explains the high estimates for portfolio 4 at the very start of the estimation period, shown in Chart 1.

Chart 1
Plots of forecasts

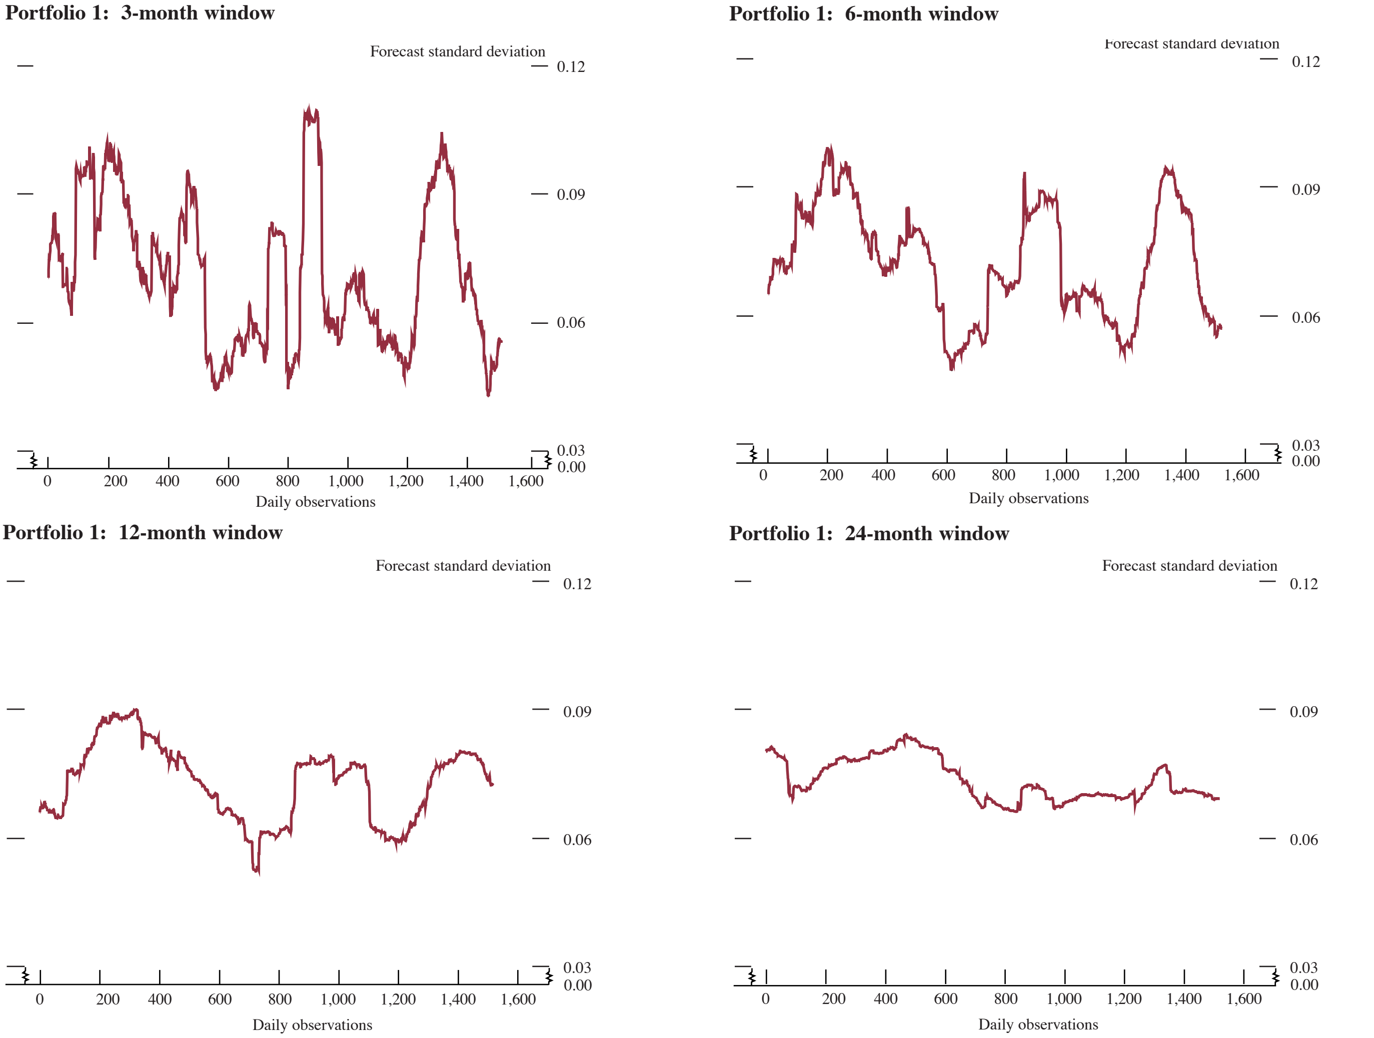

1% of the time.(1) As a measure of the performance of different VaR models, the lower panel in Table A shows the proportion of actual portfolio returns that fall below the 1% cut-off points.

As the upper panel of Table A shows, the mean absolute forecast errors are relatively insensitive to the length of the data window, though in most cases a short window yields slightly more accurate forecasts. On the face of it, the insensitivity is surprising, since plots of the forecasts based on long or short windows look quite different (see Chart 1). Furthermore, comparisons of the forecasting accuracy of different VaR techniques applied to individual exchange rate returns included in J P Morgan (1995) suggest that different window lengths *do* make a difference (although not a large one). In fact, the accuracy of forecasts of volatilities and the sensitivity of the forecasts to different techniques depend very much on the return series in question. When we repeated the analyses reported in Table A using the return on a single exchange rate, as in J P Morgan (1995), we found distinctly greater differences between the forecasting performances of different VaR techniques. However, it is important to note that using a different window size significantly affects the tail probabilities shown in the lower part of the table. In general, the figures in the table show that losses exceed the 1% cut-off points much more than 1% of the time, demonstrating the inaccuracy of the measures of tail probability implied by parametric VaR models based on normal distributions. Hendricks (1996) reaches a similar conclusion in his study of VaR models applied to foreign exchange portfolio returns. This is not surprising given the widely documented leptokurtosis (ie

(1) More precisely, the cut-off points may be obtained by inverting the equation:

$$\text{Prob}\left[\sum_{n=1}^{n} r_{nt} a_n < -\gamma \,\middle|\, \sigma^2, \mu\right] = 0.01$$

for γ on a period-by-period basis. (In the equation above, a_n is the holding of the nth asset. Throughout our analysis, we shall normalise initial wealth to unity, so that $\sum_{n=1}^{N} a_n = 1$.) Inverting the equation yields:

$$\gamma \equiv -\mu - \Phi^{-1}(0.01)\sigma$$

where $\Phi(\cdot)$ is the cumulative distribution function for a standard normal random variable.

fat-tailed distribution) of interest rates and stock returns. But the results in Table A suggest that a longer data window helps to reduce the tail probability bias.

(b) Weighting schemes

As mentioned before, a common procedure is to calculate variance estimates for VaR-type analyses using *weighted* squared deviations from an estimate of the mean. Rapidly declining weights mean that variance estimates are largely based on the last few observations, though information contained in more lagged observations is not totally ignored. The motivation for this approach is the widely recognised fact that financial market returns are conditionally heteroskedastic.

A range of more or less complicated techniques has been developed to model this feature of financial returns. In particular, Generalised Autoregressive and Conditionally Heteroskedastic (GARCH) models are specifically designed for this purpose.(1) Most implementations of VaR analysis have taken the simpler approach of estimating variances using the weighted average of squared deviations from the mean described above, with weights that decline exponentially as the lag length increases. The weights are thus of the form:

$$\lambda_i \equiv T\frac{1-\lambda}{1-\lambda^{T-1}}\lambda^i \qquad i = 0, 1, 2, \ldots, T-1 \qquad \textbf{(3)}$$

for a constant $\lambda \in [0,1]$, where $\sum_0^{T-1}\lambda_i = T$.

The upper panel of Table B shows mean absolute volatility forecast errors obtained using different weighting schemes. The calculations are carried out using daily returns with 24-month windows of lagged data, and means fixed at zero. Once again, the volatility forecasts for the fixed income and foreign exchange books are quite insensitive to the precise approach followed, although rapidly declining weights (λ = 0.94) perform somewhat better for all four portfolios, and yield a statistically significant improvement in forecast accuracy for portfolio 4. The lower panel of Table B shows the tail probabilities for different weighting schemes. It is apparent that using weighting schemes with rapidly declining weights increases the upward bias in the tail probabilities. As with window length, there appears to be a trade-off, in that weighting schemes may improve the degree to which the VaR calculations track time-varying volatilities (ie the mean absolute forecast errors may be reduced to some small degree), but worsen the bias in the tail probabilities.

(c) The inclusion of estimated means

The last exercise we perform to assess the sensitivity of VaR analyses to different assumptions is to calculate mean absolute forecast errors for parametric VaR models (i) with means estimated from lagged returns, and (ii) with the means set to zero. Fixing the means at zero might seem an unconventional statistical procedure, but the estimation error associated with badly determined mean estimates in relatively small samples may reduce the efficiency of the estimated volatilities. (Figlewski (1994) makes a similar point in the context of return variance estimation.) If the true mean returns are, as seems likely, very close to zero, fixing them at this level could enhance the forecasts. In fact, the results in Table C show that, for the particular books and return data we employ, the findings are mixed. The mean absolute forecast errors with means set to zero are in some cases lower and in some higher than when the means are freely estimated. With one-day returns, the

Table B
Parametric VaR models: exponential weights

		Equal weights	λ = 0.97	λ = 0.94
		Mean absolute errors		
Portfolio 1	Mean	27.17	26.37	26.11*
	Standard error	(0.60)	(0.84)	(0.94)
	t-statistic	*1.67*	*1.33*	*n/a*
Portfolio 2	Mean	17.29	17.05	16.86*
	Standard error	(0.41)	(0.53)	(0.60)
	t-statistic	*1.08*	*1.26*	*n/a*
Portfolio 3	Mean	5.40	5.36	5.30*
	Standard error	(0.14)	(0.19)	(0.22)
	t-statistic	*0.71*	*1.03*	*n/a*
Portfolio 4	Mean	78.60	76.49	75.62*
	Standard error	(1.72)	(1.98)	(2.15)
	t-statistic	*2.18*	*1.61*	*n/a*
		Tail probabilities		
Portfolio 1		1.32*	1.32*	1.72
Portfolio 2		1.51*	1.71	1.91
Portfolio 3		1.25*	1.45	1.45
Portfolio 4		1.38*	1.65	1.65

n/a = not applicable.

Notes: Calculations use zero means, daily returns, and a 24-month window.
Forecast errors are scaled up by 10,000.
Asterisks indicate lowest in the row.
Newey-West standard errors are in parentheses.
T-statistics are given for the difference from the lowest mean absolute error in the same row.

Table C
Parametric VaR models: sample mean inclusion

Mean absolute forecast errors

			Sample mean	Zero mean
Portfolio 1	one-day return	Mean	27.30	27.17
		Standard error	(0.61)	(0.60)
		t-statistic	*2.01*	*n/a*
	ten-day return (a)	Mean	82.54	81.58
		Standard error	(2.44)	(2.46)
		t-statistic	*0.95*	*n/a*
Portfolio 2	one-day return	Mean	17.31	17.29
		Standard error	(0.41)	(0.41)
		t-statistic	*0.56*	*n/a*
	ten-day return (a)	Mean	51.27	50.67
		Standard error	(1.34)	(1.38)
		t-statistic	*0.86*	*n/a*
Portfolio 3	one-day return	Mean	5.39	5.40
		Standard error	(0.14)	(0.14)
		t-statistic	*n/a*	*1.14*
	ten-day return (a)	Mean	16.34	16.38
		Standard error	(0.45)	(0.49)
		t-statistic	*n/a*	*0.23*
Portfolio 4	one-day return	Mean	78.53	78.60
		Standard error	(1.73)	(1.72)
		t-statistic	*n/a*	*0.34*
	ten-day return (a)	Mean	237.69	232.23
		Standard error	(7.23)	(7.65)
		t-statistic	*1.68*	*n/a*

n/a = not applicable.

Notes: Calculations use equal weights, one-day returns, and a 24-month window.
Forecast errors are scaled up by 10,000.
Asterisks indicate lowest in the row.
Newey-West standard errors are in parentheses.
T-statistics are given for the difference from the lowest mean absolute error in the same row.

(a) Calculated by multiplying one-day returns by $\sqrt{10}$.

(1) See the August 1997 *Quarterly Bulletin*, page 288, for more details on GARCH models.

differences are very small. With portfolio 1, one-day return forecast accuracy is improved in a statistically significant way, but the gain appears economically insignificant.

Parametric versus non-parametric VaR models

In this section, we compare the performance of parametric and non-parametric VaR models. Since non-parametric VaR models do not yield a time series of volatility forecast errors, we restrict our comparison to the tail probabilities that the two kinds of model produce. Table D shows the results for data window lengths ranging from 3 to 24 months. For the parametric approach, ten-day return tail probabilities were calculated by scaling up the one-day VaR estimates by $\sqrt{10}$, and then taking the fraction of observations for which the ten-day loss outturns exceed the implied cut-off level. The one-day tail probabilities are calculated as in previous sections. For the non-parametric approach, ten-day return tail probabilities were calculated using ten-day portfolio losses to compute the VaR, and then taking the fraction of observations for which the ten-day loss outturns exceed the implied cut-off level. For the one-day tail probabilities, the VaR was computed using one-day portfolio losses, and the result was compared with the one-day outturns. For both the parametric and the non-parametric approaches, the ten-day return outturns were computed on a rolling basis by summing the log daily returns.

Table D
Parametric and simulation VaRs: tail probabilities

	3 months' data	6 months' data	12 months' data	24 months' data
Portfolio 1				
One-day return parametric	1.71	1.38	1.32	1.32
Ten-day return parametric (a)	1.78	1.05	1.32	1.05
One-day return simulation	1.71	0.79	1.38	0.92
Ten-day simulation (b)	3.69	1.97	2.30	1.78
Portfolio 2				
One-day return parametric	2.11	1.91	1.58	1.51
Ten-day return parametric (a)	0.79	0.72	0.99	0.92
One-day return simulation	1.78	0.99	1.18	1.18
Ten-day return simulation (b)	2.63	1.32	1.45	1.65
Portfolio 3				
One-day return parametric	1.58	1.32	1.45	1.25
Ten-day return parametric (a)	1.58	1.12	1.05	1.05
One-day return simulation	1.51	0.86	1.18	0.86
Ten-day return simulation (b)	3.09	1.32	1.58	1.18
Portfolio 4				
One-day return parametric	1.71	1.65	1.71	1.38
Ten-day return parametric (a)	1.12	1.12	1.18	0.92
One-day return simulation	1.38	0.72	1.38	0.92
Ten-day return simulation (b)	3.09	1.58	1.38	1.25

(a) Calculated by multiplying the one-day VaR estimate by $\sqrt{10}$ and comparing this with the subsequent realised ten-day log returns.
(b) Calculated by estimating the VaR from the portfolio losses over ten-day periods and comparing these with the subsequent realised ten-day log returns.

The results in the table suggest that calculating the one-day and ten-day VaR cut-off points from short data windows is inadvisable, in that the small-sample biases are substantial. For longer data windows, the non-parametric approach for the one-day returns consistently outperforms the parametric VaR model, in that the tail probabilities are matched more accurately. For the parametric approach, the tail probabilities computed using the different lag lengths consistently exceed the 1% level, reflecting the well-known non-normality of financial returns. Looking at the ten-day returns, the non-parametric approach appears to perform worse than the parametric VaR estimates for some portfolios. In general, the tail probability figures for ten-day returns underline the statistical problems involved in attempting to deduce ten-day volatilities directly from estimates of one-day volatilities.

'Spike' loss periods

An important question is whether the ability of parametric VaR analysis to 'track' the time-series behaviour of volatility enables it to outperform simulation-based VaR models in predictions of large 'spike' losses in portfolio values. It is possible that even if parametric VaR models do not yield lower mean absolute forecast errors, as we saw above, they are better at picking out large market movements. This issue is particularly important if VaR analysis is to be used for regulatory purposes, since the primary concern of regulators regarding trading-book risks is that banks will be wiped out by sudden large losses that occur before action can be taken to reduce the riskiness of the bank's portfolio. To examine this issue, we split our sample period into six-month intervals and identify, for each of our portfolios, the day within each period on which the largest loss occurred.

Before comparing the performance of the parametric and simulation-based VaR models, let us examine the composition of the spike portfolio losses. Table E provides detailed breakdowns of the constituent parts of each of these large-value declines for portfolio 4, which contains equity as well as interest rate and foreign exchange risk. As is apparent from Table E, bond risk was the most important factor in generating large losses, acting as the dominant factor in eight out of twelve cases. Foreign exchange risk was the most important factor in the remaining four cases. The table in the box on portfolios data shows that portfolio 4 contains greater foreign exchange exposure than the other portfolios (in particular, a relatively large net US dollar position).

It is surprising that the equity exposure created no spike losses in the period of our sample. We were concerned that this result reflects the fact that large changes in equity values tend to be negative, and the largest equity exposure in portfolio 4 is a short position in US equities. As an experiment, we re-ran the VaR calculations assuming that the equity exposures (and the corresponding components of the foreign exchange exposures) were of opposite sign. Even with this change, none of the spike losses were attributable mainly to equity losses. One may, therefore, conclude that the relatively small size of the equity exposure is enough to make equity risk minimal, even though equity returns themselves are much more volatile than those on bond portfolios.(1)

(1) The more 'spiky' and volatile nature of equities has been recognised by regulators, eg in the CAD building-block approach. Under the CAD, a single position in a ten-year government bond would carry a capital requirement of 2.4%, whereas a single position in an equity index would carry a charge of 8%. For a single equity, the charge would be 12%.

Table E
‘Spike losses’ (daily returns in per cent)—portfolio 4

Date		France	United Kingdom	United States	Japan	Germany	Total
03/07/89	FX	0.13	n/a	-2.03	-0.01	-0.11	-2.02
	Bond	0.26	-0.09	-0.05	-1.61	-1.12	-2.61
	Equities	0.01	0.02	-0.12	0.04	0.09	0.03
	Total	**0.39**	**-0.07**	**-2.20**	**-1.58**	**-1.14**	**-4.60**
21/02/90	FX	0.01	n/a	-0.72	0.06	-0.02	-0.67
	Bond	1.35	0.02	0.04	0.46	-4.22	-2.36
	Equities	-0.01	-0.02	0.03	-0.16	-0.06	-0.23
	Total	**1.34**	**0.00**	**-0.65**	**0.35**	**-4.30**	**-3.26**
06/08/90	FX	-0.04	n/a	-0.87	0.05	0.04	-0.82
	Bond	-3.18	-0.32	-2.41	-1.47	2.99	-4.38
	Equities	-0.07	-0.07	0.98	-0.16	-0.39	0.29
	Total	**-3.28**	**-0.38**	**-2.30**	**-1.58**	**2.64**	**-4.90**
11/02/91	FX	-0.04	n/a	-0.56	0.04	0 .06	-0.50
	Bond	0.75	-0.04	-0.13	-1.65	-1.38	-2.45
	Equities	0.01	0.04	-0.81	0.00	0.10	-0.66
	Total	**0.73**	**-0.00**	**-1.50**	**-1.61**	**-1.23**	**-3.61**
01/09/91	FX	-0.03	n/a	-2.08	0.11	0.06	-1.95
	Bond	0.35	-0.06	0.03	-1.09	-1.10	-1.88
	Equities	-0.00	-0.01	0.04	-0.03	-0.06	-0.05
	Total	**0.32**	**-0.07**	**-2.01**	**-1.01**	**-1.10**	**-3.87**
18/11/91	FX	-0.18	n/a	-1.35	0.09	0.15	-1.28
	Bond	-0.50	0.07	-0.04	-0.04	-0.14	-0.67
	Equities	-0.04	-0.04	-0.22	-0.15	-0.07	-0.52
	Total	**-0.72**	**0.03**	**-1.60**	**-0.11**	**-0.06**	**-2.47**
23/09/92	FX	0.09	n/a	0.03	-0.08	-0.17	-0.13
	Bond	-3.25	-0.05	-0.34	-0.06	-2.33	-6.02
	Equities	-0.00	-0.00	-0.02	0.00	0.03	0.01
	Total	**-3.16**	**-0.05**	**-0.33**	**-0.14**	**-2.46**	**-6.15**
05/01/93	FX	0.47	n/a	-3.14	0.26	-0.46	-2.87
	Bond	-0.30	-0.24	-0.13	0.06	-0.54	-1.15
	Equities	0.01	-0.01	0.08	-0.05	0.11	0.13
	Total	**0.18**	**-0.25**	**-3.19**	**0.27**	**-0.89**	**-3.89**
13/04/93	FX	0.09	n/a	-2.46	0.20	-0.11	-2.27
	Bond	0.31	0.06	-0.23	-0.81	-0.20	-0.88
	Equities	0.02	0.02	-0.06	0.22	0.06	0.26
	Total	**0.42**	**0.08**	**-2.75**	**-0.40**	**-0.25**	**-2.89**
01/03/94	FX	0.05	n/a	0.01	0.04	-0.03	0.07
	Bond	-1.51	-0.17	-1.07	-1.79	0.86	-3.68
	Equities	-0.03	-0.03	0.18	0.06	-0.08	0.09
	Total	**-1.50**	**-0.20**	**-0.88**	**-1.69**	**0.75**	**-3.52**
28/06/94	FX	0.00	n/a	0.58	-0.02	-0.01	0.55
	Bond	-0.23	-0.08	-0.78	-1.44	-3.15	-5.67
	Equities	0.01	0.01	0.09	0.08	0.10	0.29
	Total	**-0.22**	**-0.07**	**-0.11**	**-1.37**	**-3.06**	**-4.82**
03/10/94	FX	0.10	n/a	-0.07	0.09	-0.10	0.02
	Bond	-1.64	-0.06	-0.49	-1.19	-0.03	-3.42
	Equities	-0.02	-0.03	0.07	0.02	0.00	0.04
	Total	**-1.57**	**-0.09**	**-0.49**	**-1.08**	**-0.13**	**-3.36**

n/a = not applicable.

Note: Components may not sum to totals because of rounding.

Table F shows the capital requirement implied by the VaR estimates minus the actual loss sustained.(1) We term this quantity the capital surplus (+) or capital shortfall (–). As one may see, parametric and simulation-based VaR models perform somewhat differently. When capital is based on the simulation-based VaR model, the bank has a capital surplus on 16 of the 48 spike loss dates. When the parametric VaR model is used, the bank has a surplus on nine occasions. Whether the capital surplus is positive or negative, on most spike loss dates, the simulation-based VaR model implies a larger capital surplus than the parametric VaR model. The implication is that, though it does not exploit the conditional structure of volatility, the simulation-based VaR model seems to do a somewhat better job of establishing appropriate capital requirements. Chart 2 illustrates this, using a 24-month window, for each of the portfolios.

Basle alternative approach capital calculations

A final important question is how much of a capital cushion the proposed Basle alternative approach would deliver for actual books, given not only the 99% confidence level, but also the multiplier of three. We look at this issue for our portfolios, by comparing the capital requirement that would be generated by one part of the proposed two-stage test, namely three times the 60-day average of the VaR estimates calculated to cover a ten-day holding period, using the parameters laid down by Basle. A bank would be required to hold capital equivalent to the greater of (i) this amount and (ii) the VaR for the current book. With a multiplier of three, the first of these tests will ‘bite’, unless the bank’s current book is abnormally risky.

Table F
Model performance on ‘spike’ loss dates

Model	Portfolio 1 Simulation	Portfolio 1 Parametric	Portfolio 2 Simulation	Portfolio 2 Parametric
Period 1	-1.63	-1.51	-0.49	-0.47
Period 2	-0.56	-0.64	-0.42	-0.43
Period 3	-0.75	-0.89	-0.48	-0.54
Period 4	-0.03	-0.08	-0.29	-0.39
Period 5	0.28	0.11	0.15	0.02
Period 6	-1.08	-1.34	-1.05	-1.22
Period 7	-1.81	-2.09	-1.39	-1.51
Period 8	0.04	-0.24	-0.31	-0.35
Period 9	0.40	0.15	-0.08	-0.10
Period 10	0.11	-0.08	0.06	0.00
Period 11	-0.07	-0.10	-0.04	-0.04
Period 12	-0.16	-0.08	0.18	0.12

Model	Portfolio 3 Simulation	Portfolio 3 Parametric	Portfolio 4 Simulation	Portfolio 4 Parametric
Period 1	-0.06	-0.05	-0.81	-0.58
Period 2	-0.08	-0.10	0.05	-0.15
Period 3	-0.11	-0.13	-1.62	-1.95
Period 4	-0.10	-0.12	-0.32	-0.53
Period 5	-0.09	-0.12	-0.62	-0.79
Period 6	-0.08	-0.16	0.79	0.58
Period 7	-0.75	-0.80	-3.19	-3.29
Period 8	-0.01	-0.10	-0.34	-0.79
Period 9	0.16	0.06	0.66	0.13
Period 10	0.04	-0.03	-0.54	-0.47
Period 11	0.03	0.01	-1.28	-1.40
Period 12	0.04	0.04	0.29	-0.09

Notes: The table shows the capital shortfall (-) or surplus (+) for the largest loss in each six-month period. Parametric approach uses zero mean. Figures are expressed as daily returns in per cent. Equal weights, daily returns, 24-month window.

We compared the ten-day returns that would have been secured on our four portfolios during the period July 1989 to April 1995 with the capital requirement based on three times the 60-day average of the ten-day VaR estimates calculated by multiplying the daily VaR estimates by $\sqrt{10}$. (The Basle requirement would usually be calculated using the 60-day average for VaR estimates for different books held on different days.) In performing the calculations, we used the parametric approach, with a 24-month window of past returns data, equal weights, and a zero mean. We calculated the capital requirement implied by multipliers of two and two and a half, as well as three. None of the portfolios had a single loss outlier (losses that exceeded the capital requirement) when the multiplier was either two and a half or three. Three of the portfolios had a single (marginal) loss outlier for a multiplier of two.

The Basle approach to back-testing

The alternative Basle approach includes a requirement that banks would suffer increases in their capital requirements if, over a twelve-month period (250 trading days), their VaR models under-predict the number of losses exceeding the 1% cut-off point. Such losses are termed ‘exceptions’. If a bank’s VaR model has generated zero to four exceptions, it

(1) The capital ‘requirement’ is the VaR for the whole book produced using a 99% confidence level. We do not incorporate in this calculation any other aspects of the Basle proposals, such as the three-times multiplier.

Chart 2
Comparison of simulation and parametric-based VaR models

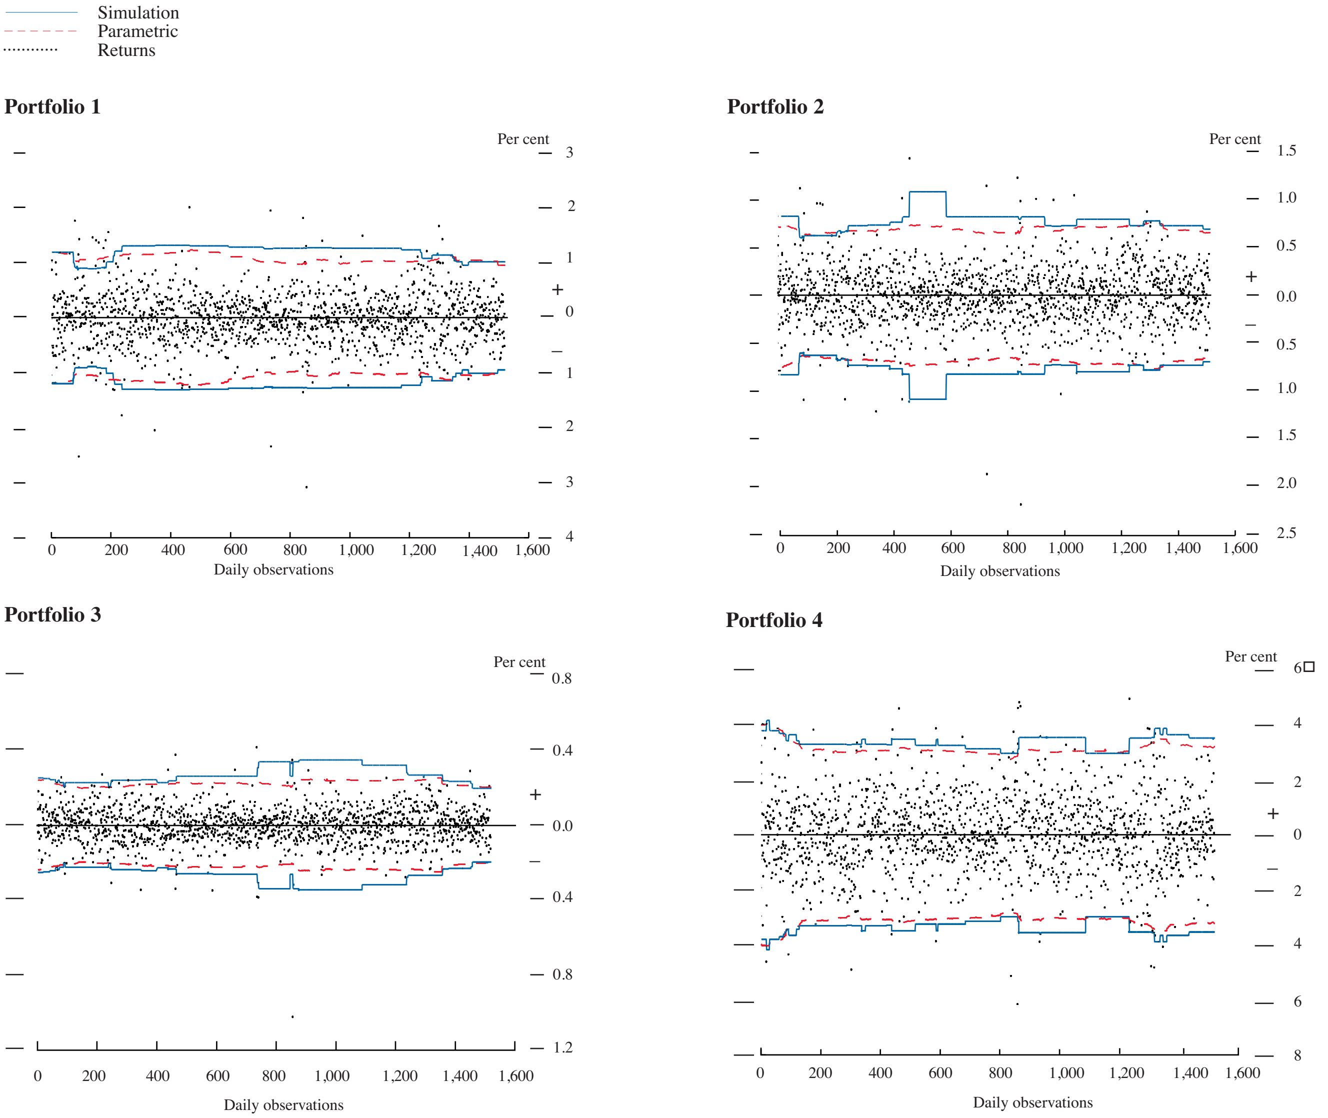

is said to be in the Green Zone; if five to nine, it is in the Yellow Zone; and if there are more than ten exceptions, it is in the Red Zone. The capital requirement for banks whose models are in the Yellow Zone may be increased by regulators; if they are in the Red Zone, the requirement would almost certainly be increased.

We ran back-tests for all four of our portfolios, comparing the VaR figures calculated for one-day holding periods (again, using the parametric approach) with the actual return on each book. The number of exceptions for each portfolio over the different twelve-month periods are set out in Table G. The results vary for different portfolios. For three of the six periods, if portfolio 2 were held, the model would generate more than four exceptions. The highest number of exceptions was seven, which occurred twice for portfolio 2 and once for portfolio 4. According to the Basle guidelines, this would normally lead to an increase of 0.65 in the multiplier, unless the supervisor could be persuaded that special factors had affected outcomes.(1) The fact that the model moves from the Green to the Yellow Zone so much from period to period underlines the difficulty of distinguishing between good and bad models using samples of only 250 observations. However, our results suggest that a grossly inaccurate model would be picked up by such back-testing.

Table G
Back-testing results—number of exceptions in each twelve-month period

Portfolio	1	2	3	4
Period 1	6	7	4	3
Period 2	4	7	5	3
Period 3	3	2	4	1
Period 4	4	5	4	4
Period 5	1	1	2	3
Period 6	2	1	0	7

Green Zone = 0–4 exceptions.
Yellow Zone = 5–9 exceptions.
Red Zone = 10+ exceptions.

(1) Supervisors can disregard the Yellow Zone if they believe that there is a good reason for the poor performance, unrelated to the model. However, the Red Zone can only be disregarded in extraordinary circumstances.

Summary

This article has set out the results of the tests carried out by the Bank to assess the accuracy of the risk-measurement models used by firms to evaluate risk on their trading-book portfolios. The main conclusions from these tests were as follows:

- Different VaR models performed more or less well in supplying unbiased measures of the value at risk. (For some VaR models built with a 99% confidence level, significantly more than 1% of losses exceeded the value-at-risk estimate.)

- Simulation-based VaR models met this test better than parametric VaR models based on normal distributions, because of the severely fat-tailed nature of reasonably diversified fixed-income exposures. Most banks' trading books are made up largely of such exposures.

- Use of short data samples (or a weighting scheme that places heavy weight on recent data) worsened the biases in the VaR estimates for parametric models.

- The extra safeguards around the use of the VaR models (the requirement that a firm must meet the higher of the estimated VaR, or three times the 60-day moving average of the current and past VaRs) would probably mean, for market-risk models of the kind tested, that only extremely risky portfolios would fail to be covered by sufficient capital.

- The back-testing requirements incorporated in the Basle approach are likely to lead to some banks holding higher capital. A bank holding the portfolios employed in the study could find its capital requirements adjusted upwards from time to time if it used the parametric approach.

References

Basle Committee on Banking Supervision (1994), *Risk Management Guidelines for Derivatives*, Basle.

Basle Committee on Banking Supervision (1995a), *Planned Supplement to the Capital Accord to Incorporate Market Risks*, Basle.

Basle Committee on Banking Supervision (1995b), *An Internal Model-Based Approach to Market Risk Capital Requirements*, Basle.

Basle Committee on Banking Supervision (1996), *Overview of the Amendment to the Capital Accord to Incorporate Market Risk*, Basle.

Figlewski, S (1994), 'Forecasting Volatility Using Historical Data', *New York University Working Paper*, S-94-13.

Hendricks, D (1996), 'Evaluation of Value-at-Risk Models Using Historical Data', *Federal Reserve Bank of New York Economic Policy Review*, Vol 2, No 1, April, pages 39–69.

J P Morgan (1995), *RiskMetrics Technical Document*, Third edition, JP Morgan: New York.

Kupiec, P H (1995), 'Techniques for Verifying the Accuracy of Risk Measurement Models', *Journal of Derivatives*, Winter.

Newey, W K, and West, K D (1987), 'A Simple Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix', *Econometrica*, Vol 55, pages 703–708.