VALUE AT RISK: ON THE STABILITY AND FORECASTING OF THE VARIANCE-COVARIANCE MATRIX

James Engel* and Marianne Gizycki**

Research Discussion Paper
1999-04

May 1999

*Research and Finance Studies, Australian Prudential Regulation Authority
**System Stability Department, Reserve Bank of Australia

The authors would like to thank Luci Ellis, Brian Gray, David Gruen, Philip Lowe and Geoff Shuetrim for helpful comments. We, nevertheless, retain responsibility for any remaining errors. The views expressed herein are those of the authors and do not necessarily reflect the views of the Reserve Bank of Australia and the Australian Prudential Regulation Authority.

Abstract

Over the past decade value at risk (VaR) has become the most widely used technique for the quantification of market-risk exposure. VaR is a measure of the potential loss that may occur from adverse moves in market prices (interest rates, exchange rates, equity prices and so forth). The capacity for a VaR measure to accurately predict future risk exposures depends upon the forecasts of the volatility of market rates and the correlations between the various market rates (that is, the variance-covariance matrix) incorporated into the VaR model. In this paper we first present the results of tests of the stability of the variances, covariances and correlations for exchange rates and Australian interest rates. Secondly, we assess the performance of several time-series models that may be used to forecast the variance-covariance matrix. In particular three models for the variance-covariance matrix are considered: equally weighted historical variances and covariances, exponentially weighted averages of historical variances and generalised autoregressive conditional heteroskedasticity (GARCH). We conclude that simple models perform as well as their more sophisticated GARCH counterparts.

JEL Classification Numbers: C22, E47, G21

Keywords: value at risk, market risk, volatility, correlation, GARCH

Table of Contents

VALUE AT RISK: ON THE STABILITY AND FORECASTING OF THE VARIANCE-COVARIANCE MATRIX

James Engel and Marianne Gizycki

1. Introduction

Over the past decade value at risk (VaR) has become the most widely used technique for the quantification of market-risk exposure. From January 1998 banks in Australia have been permitted (subject to a range of conditions) to use their VaR models as the basis for determining the capital that is required to cover market-risk exposure. VaR is a measure of the potential loss that may occur from adverse moves in market prices (interest rates, exchange rates, equity prices and so forth). Specifically it is the dollar amount that portfolio losses are not expected to exceed, with a specified degree of statistical confidence, over a pre-specified period of time. There are a number of different methodologies used to calculate VaR (Cassidy and Gizycki (1997) provide a discussion of these models). The most widely used method is that known as the variance-covariance approach. This approach is based on the simplifying assumption that financial-asset returns are normally distributed and hence, the statistical distribution of these returns can be completely described by the mean of the market returns, the variance of market returns and the correlations between the various market rates (that is, the variance-covariance matrix).

The capacity for a variance-covariance VaR measure to accurately predict future risk exposures depends upon the quality of forecasts of the variance-covariance matrix incorporated into the VaR model. In this paper we first present the results of tests of the stability of the variances, covariances and correlations for exchange rates and Australian interest rates. Next we assess the performance of several time-series models that may be used to forecast the variance-covariance matrix, in particular three models for the variance-covariance matrix are considered: simple historical variances, exponentially weighted averages of historical variances (where the weights are progressively smaller for observations further in the past) and generalised autoregressive conditional heteroskedasticity (GARCH) models. While most Australian banks use the simpler models (past, observed variances or exponentially weighted moving averages) there are some banks that use the more complex GARCH approaches within their risk measurement models. This paper investigates the benefits of the more sophisticated approaches.

2. Stability of the Variance-covariance Matrix

The standard approach, when using the variance-covariance method of estimating VaR, is to use historical variances and covariances as the forecasts of the future variance-covariance matrix. This approach assumes that the variances and covariances are constant over the period of estimation and forecast.

We test this assumption using daily data for working days from 15 December 1983 to 29 October 1997. Nine foreign exchange rate returns and eight interest rate return series are used.[1]

Figure 1 shows the correlation between selected pairs of exchange rate returns and interest rate changes. Correlations calculated using moving windows of 250 and 1 250 working days (that is, approximately one and five years) are shown. These graphs indicate, particularly when shorter window lengths are considered, that correlations move considerably through time.

To formally test this impression the global test for a constant unconditional correlation/covariance matrix as outlined in Jenrich (1970) is applied. This test quantifies the difference between two matrices via the trace of a relative difference matrix; the relative difference matrix being the difference between the two

1 The exchange rates are the Australian dollar against the US dollar, German mark, Canadian dollar, French franc, British pound, Japanese yen, New Zealand dollar, Swiss franc and Dutch guilder. The foreign exchange series are the indirect rates (that is, rates are expressed in terms of the foreign currency value of one Australian dollar). The interest rates are the overnight cash rate; the 30, 90 and 180-day bank accepted bill rates; and the 1, 2, 5 and 10-year Treasury bond yields. Returns are calculated as proportional changes in underlying rates.

Figure 1: Correlation Stability

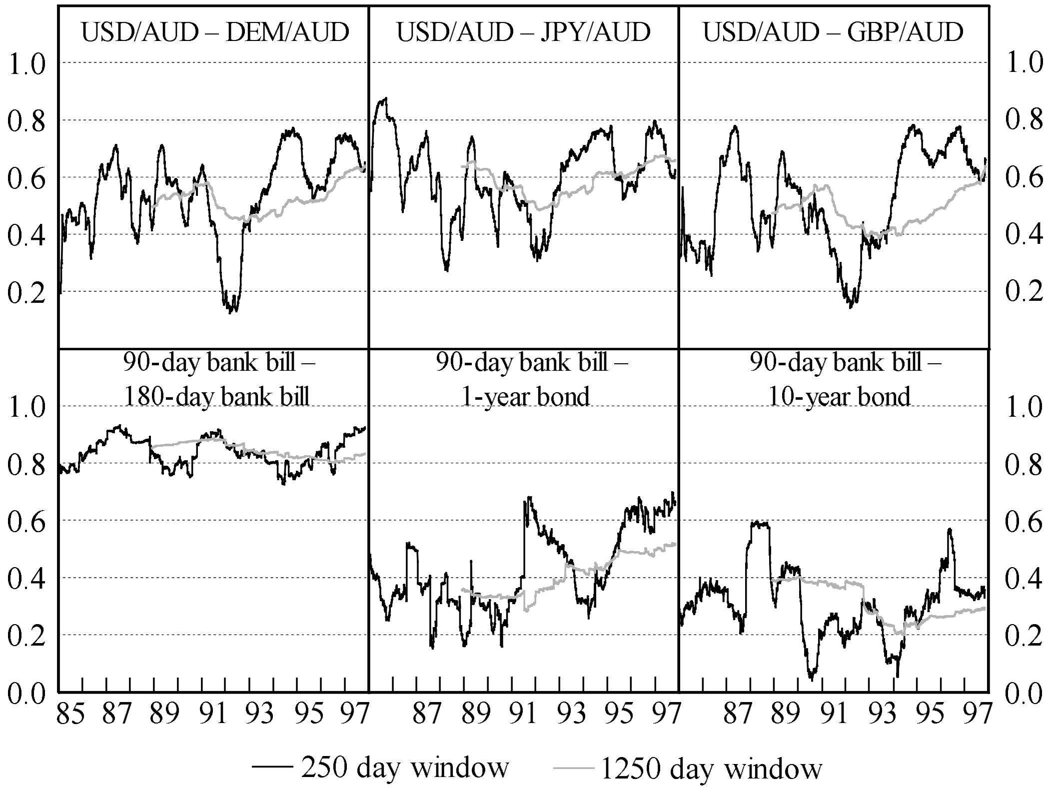

matrices divided by their sum.[2] Firstly, within-asset class covariances and correlations are considered. The unconditional covariance and correlation matrices for the nine foreign exchange return series and the eight interest rate return series

---

[2] A variance-covariance matrix with dimension p has $p(p\text{-}1)/2 + p$ independent elements, but the corresponding *correlation* matrix has only $p(p\text{-}1)/2$ independent elements since its diagonal elements are all unity. The Jenrich test has, in each case, an asymptotic chi-squared distribution with the number of degrees of freedom equal to the number of independent elements in the matrix. For testing the equality of the correlation matrix the test statistic has the form:

$$\chi^2 = \frac{1}{2} tr(c^{0.5}\overline{R}^{-1}(R_1 - R_2))^2 - dg'((c^{0.5}\overline{R}^{-1}(R_1 - R_2))^2)S^{-1}dg((c^{0.5}\overline{R}^{-1}(R_1 - R_2))^2)$$

$$c = n_1 n_2/(n_1 + n_2),\ \overline{R} = (n_1 R_1 + n_2 R_2)/(n_1 + n_2),\ S = (d_{ij} + \overline{r}_{ij}\overline{r}^{ij})$$

where the Rs denote sample correlations, dg the diagonal operator, n_1 and n_2 the length of data used to estimate matrices R_1 and R_2 respectively and δ_{ij} the kronecker delta (one if i=j and zero otherwise). If the sample correlations are replaced by sample covariances, the first term in the equation above becomes the test for stability of covariances. The second term is hence, a correction employed when testing correlation matrices.

are analysed separately. The full sample was broken into separate fixed-length sub-periods. The Jenrich test is then applied to test the equality of each pair of matrices calculated from adjacent sub-periods. The analysis was repeated for sub-period lengths ranging between 125 and 1 250 days.

The analysis was conducted for the full foreign exchange and interest rate covariance matrices and also for subsets of these matrices. The subsets considered were the Australian dollar against the US dollar, Japanese yen, German mark and British pound (major currencies); and the Australian dollar against the Canadian dollar, French franc, Dutch guilder and New Zealand dollar (other currencies) in the case of foreign exchange. For the interest rate series the subsets were the Treasury bond yields for one, two, five and ten-years (bond yields); and the yields on bank accepted bills of 30, 90 and 180-day maturities in conjunction with the overnight cash rate (discount securities). The numbers reported in Table 1 are the proportion of times a given covariance or correlation matrix was found to be equal to the matrix estimated from the following consecutive sub-period of data (at a 5 per cent level of significance).

Table 1: Proportion of Stable Comparisons

Per cent

Window length	Covariances	Correlations	Covariances	Correlations	Covariances	Correlations
Foreign exchange	All series		Major currencies		Other currencies	
125 day	0	0	19	56	0	15
250 day	0	0	8	46	0	8
500 day	0	0	17	67	0	0
750 day	0	0	0	0	0	0
1 250 day	0	0	0	100	0	0
Interest rates	All series		Bond yields		Discount securities	
125 day	0	4	0	26	0	19
250 day	0	0	0	15	0	0
500 day	0	0	0	0	0	0
750 day	0	0	0	0	0	0
1 250 day	0	0	0	0	0	0

It can be seen that the proportion of comparisons found to be equal is low. The results indicate that the correlation and covariance matrices are far from constant. In only one case was the correlation matrix stable over the full sample period. There is some tendency for correlations to be relatively more stable than covariances. To some extent this is to be expected since covariances reflect not just the relation between two series but the variances of the series as well. To the extent that variances are not stable this will be reflected in the covariances but not the correlations.

The statistical significance of the difference in the matrices appears to be an increasing function of the length of the sub-period compared. This could be indicative of a gradual change in the true covariance matrix over time. It can also be seen that instability is related to the dimension of the covariance matrix. Testing was carried out on various bivariate matrices to isolate any potential outlier within the full matrix. Appendix A reports the proportion of stable comparisons for the bivariate covariances and correlations. The USD/AUD and DEM/AUD covariance matrix and the 90-day bank accepted bill and overnight cash rate covariance matrix appear to be less stable than the other bivariate systems. However, removing these from the full analysis had no substantial effect. With the bivariate systems the proportion of stable matrices is dramatically increased. Thus it seems that the greater the number of financial-asset returns that a bank is exposed to the less stable will be the associated variance-covariance matrix.

Our results are consistent with Kaplanis (1988) in that the correlation matrix is relatively more stable than the covariance matrix. Kaplanis, however, found evidence of a constant correlation matrix over adjacent 46-month periods for ten stock markets' monthly returns over the period from 1967–1982. Longin and Solnik (1995) find the unconditional correlation matrix of monthly excess returns for seven countries' share price indices to be unstable over periods of five years. Sheedy (1997) applied the same test to equity index data for the US, UK, Japan, Germany and the World Index and foreign currency returns covering the US dollar against the British pound, Japanese yen, German mark; the British pound against the Japanese yen and German mark; and the Japanese yen against the German mark. Her results, with respect to the equity data, provide mixed evidence for the stability of the correlation matrix. She reports, however, that the foreign exchange data consistently reject the hypothesis of constant correlation. This is consistent with our findings.

When calculating the VaR measures it is common practice among banks to take full account of correlations within asset classes (for example, across a number of exchange rates or commodity prices) but to make more simplistic assumptions about correlations across asset classes (that is, between exchange rates, interest rates, equity returns and commodity returns). For example, the Basle Committee on Banking Supervision prohibited the use of empirical correlations when aggregating risk exposures across asset classes in their initial proposals to allow banks to use internally developed VaR models to determine required capital. One commonly cited rationale for this approach is that while reliance may be placed on within asset class correlations, correlations across asset classes are considerably more unstable.

To examine this proposition we examined the stability of a selection of correlations and covariances between exchange rates, interest rates and equity prices (the All Ordinaries Index). The results (shown in Table 2) are not significantly different from the within asset class bivariate analysis. Across-class correlations appear to be neither systematically more or less stable than correlations within asset classes (both in terms of the proportion of stable comparisons and the average magnitude of differences between matrices over time, as measured by the Jenrich statistic). This supports the Australian Prudential Regulation Authority's approach towards banks' internally-developed VaR models that may be used for capital-adequacy purposes which does not draw any distinction between across-class and within-class correlations.

Table 2: Proportion of Stable Comparisons – Across Asset Classes
Per cent

	Covariances					Correlations				
Window length:	125	250	500	750	1 250	125	250	500	750	1 250
USD versus										
BAB 90-day	22	23	0	0	0	85	85	50	33	100
Two-year bond	55	38	0	0	0	92	92	67	33	100
Five-year bond	55	38	0	0	0	96	92	67	33	100
Ten-year bond	62	38	0	0	0	100	92	67	33	100
DEM versus										
BAB 90-day	85	30	17	0	0	100	92	67	33	100
Two-year bond	85	69	17	0	0	100	92	67	67	100
Five-year bond	85	69	17	0	0	100	100	67	67	100
Ten-year bond	85	69	17	33	0	100	100	67	67	100
JPY versus										
BAB 90-day	85	69	17	33	0	100	100	67	67	100
Two-year bond	85	69	17	33	0	100	100	83	67	100
Five-year bond	85	69	17	33	0	100	100	83	67	100
Ten-year bond	85	69	17	33	0	100	100	83	67	100
All Ordinaries versus										
BAB 90-day	48	48	17	0	0	85	62	67	67	100
One-year bond	40	38	33	33	0	100	92	83	67	0
Ten-year bond	48	23	17	33	0	92	61	50	67	0
USD	37	38	33	33	0	92	92	50	0	100
DEM	33	30	17	33	100	92	92	50	67	100
JPY	33	31	0	33	0	89	84	67	33	100

3. Models for Forecasting the Variance-covariance Matrix

The apparent instability of the unconditional covariance matrix suggests that the historical covariance approach will be an inaccurate estimator of the true variance-covariance matrix. Therefore, more complex models of the evolution of the variance-covariance matrix may be required when forecasting risk exposures.

The three classes of models that we investigate are the equally weighted historical approach, the exponentially weighted moving average approach and the GARCH approach. There are many other types of models that may be used. We have, however, restricted ourselves to those models (and simple variations of those models) that are currently used by Australian banks.

The previous section's stability testing was based on standard covariance and correlation measures which take account of each series' sample mean over each sub-period. In the analysis that follows it is assumed that each financial return series has a zero mean. This assumption is a commonly used market practice when measuring market risk exposures. From a theoretical perspective the mean is both close to zero and prone to estimation error; thus, estimates of the variance-covariance matrix may be made worse by the inclusion of an inaccurate estimate of the mean (see Figlewski (1994)). If the standard formula for the variances and covariances is considered the squared return component is of the order 100 to 1 000 times greater than the mean component, hence, the inclusion of the estimate of the means will not make a significant difference.

While several banks, in their implementation of a VaR model, re-estimate the variance-covariance matrix daily, it is common practice in other banks to update the variance-covariance matrix only once a quarter. As a result we consider two sets of forecasts: the one-day-ahead forecasts and the forecast average variances and covariances over the quarter ahead.

3.1 Fixed-weight Historical

The fixed-weight approach assumes that return covariances and variances are constant over the sample period. Our finding of instability in the variance-covariance matrix indicates that this is not a good assumption. However, it is widely used on simplicity grounds. Using this approach, each element in the variance-covariance matrix can be represented by:

$$\sigma^2{}_{ij,t+1} = \frac{1}{N}\sum_{s=0}^{N-1} r_{i,t-s} r_{j,t-s} \tag{1}$$

where $r_{i,t\text{-}s}$ represents the market return for asset i between days t-s-1 and t-s.

3.2 Exponential Smoothing

Rather than placing equal weight on past observations, exponential smoothing places more weight on the most recent. This approach was popularised by JP Morgan in their RiskMetrics VaR model (JP Morgan and Reuters, 1996). The exponentially weighted moving average approach reacts faster to short-term movements in variances and covariances. If the underlying variances and covariances are not constant through time this faster reaction is an advantage. On the other hand, giving a greater weight to recent data effectively reduces the overall sample size, increasing the possibility of measurement error. Each element of the variance-covariance matrix is represented by:

$$\sigma^2_{ij,t+1} = \lambda\sigma^2_{ij,t} + (1-\lambda)r_{i,t}r_{j,t} \qquad (2)$$

where $0<\lambda<1$

An exponentially weighted average on any given day is a combination of two components: yesterday's weighted average, with weight λ, and yesterday's product of returns, which receives a weight of $(1-\lambda)$. This equation incorporates an autoregressive structure for the variance-covariance, thus reflecting the concept of volatility clustering. In the subsequent analysis two approaches are implemented. The first is to assume, consistent with the RiskMetrics specification, that λ is constant at 0.94. The second approach is to estimate λ over successive rolling windows using maximum likelihood techniques (we shall refer to this approach as the dynamic exponentially weighted moving average approach).

To gauge the accuracy of fixing λ at 0.94, the model is estimated using maximum likelihood methods over the full sample (λ was constrained to take the same value for all elements of the matrix). The value obtained from this analysis using the foreign exchange covariance matrix is 0.995.[3]

---

[3] This is not statistically significantly different from 0.94.

Like the equally weighted method the k-step ahead one-day forecasts are constant and the quarter-average forecast is equal to $\sigma^2_{ij,t+1}$. To see this:

$$\begin{aligned} E_t(\sigma^2_{ij,t+k}) &= E_t((1-\lambda)r_{i,t+k-1}r_{j,t+k-1} + \lambda\sigma^2_{ij,t+k-1}) \\ &= (1-\lambda)\sigma^2_{ij,t+k-1} + \lambda\sigma^2_{ij,t+k-1} \\ &= \sigma^2_{ij,t+k-1} \end{aligned} \tag{3}$$

3.3 Multivariate GARCH

The GARCH model of Bollerslev (1986) is a generalisation of the ARCH model introduced by Engle (1982). Variances and covariances are specified as stochastic processes that evolve over time. The intuition behind these models is similar to the exponentially weighted approach in that volatility clustering may be explicitly modelled. Less restrictions, however, are placed on the specification of the volatilities' behaviour. The previous two models are nested within the GARCH model. If α and β are zero in the below specification then the model collapses to the fixed-weight historical model. If ω is equal to zero, α = (1-λ) and β = λ then the model is equivalent to the exponentially weighted model. In a univariate setting the zero-mean GARCH(1,1)[4] model has the form:

$$\begin{aligned} &R_t = r_t \\ &r_t/I_{t-1} \sim N(0, H_t) \\ &H_t = \omega + \alpha\, r^2_{t-1} + \beta\, H_{t-1} \end{aligned} \tag{4}$$

The vector of innovations or unexpected returns is assumed to be conditionally normal with a conditional variance of H_t. The multivariate framework is analogous to the univariate in that the variance-covariance matrix is conditioned on past realisations of covariances of financial returns but the specification of the evolution of the covariances can become more complicated.

---

4 The (1,1) denotes one lagged variance term and one lagged squared return.

The more general multivariate models assume that variances and covariances rely on their own past values and innovations as well as other variables' past values and innovations. The number of parameters to be estimated in these general models is such that, as the number of variables increases, computation can become intractable. To illustrate, for our nine-by-nine foreign exchange matrix the number of parameters to be estimated by one of the more general models is 243. Given that our focus is on a model's forecasting performance, which requires repeated rolling estimation of the models, a more parsimonious parameterisation is needed.

To this end two models are used. The first model is the constant correlation multivariate GARCH model developed by Bollerslev (1990). The model has the advantage of reducing the number of parameters to be estimated to $3p + p(p\text{-}1)/2$ where p is the number of financial returns. Variances are estimated using a simple GARCH(1,1) formulation:

$$\sigma^2_{i,t+1} = \omega_i + \alpha_i r^2_{i,t} + \beta_i \sigma^2_{i,t} \tag{5}$$

The covariances are formulated as:

$$\sigma_{ij,t+1} = \rho_{ij} \sigma_{i,t+1} \sigma_{j,t+1} \tag{6}$$

Non-negativity constraints need to be imposed on the variance parameters to ensure that the conditional variance estimates are always positive. Given the recursive nature of the system, stationarity requires that $\alpha_i + \beta_i < 1$ for all i.

The parameters of the model are estimated by maximum likelihood techniques. Under standard regularity conditions the maximum likelihood estimator is asymptotically normal.[5] The log-likelihood is maximised using the Bernt, Hall, Hall and Hausman (1974) algorithm. Given the highly non-linear structure of the log-likelihood the iteration process is extremely time intensive. Even after the

---

5 Following Bollerslev (1986), if the model correctly specifies the first two conditional moments but the conditional normality assumption is violated, under suitable regularity conditions the quasi-maximum likelihood estimates will be consistent and asymptotically normal, but the usual standard errors have to be modified.

constant correlation assumption is imposed on the model the 63 parameters in the full system (for the nine-by-nine foreign exchange variance-covariance matrix) make rolling estimation computationally intractable. To facilitate rolling estimation the approach taken is to estimate separate bivariate systems for each pair of financial returns. Each of these 36 systems have seven parameters to be estimated. From these bivariate systems the full variance-covariance matrix can be constructed. To the extent that covariances are in fact jointly determined estimates produced from the pair-wise estimation may be biased and inefficient. Also, this approach can not guarantee positive definiteness of the covariance matrix but does enable forecasts to be constructed in a tractable fashion. This approach provides only one estimate of each covariance, but p-1 estimates for the ω, α and β parameters for each variance. The average of the p-1 forecasts of each variance is used.

The one-day-ahead GARCH variance forecast for the constant correlation GARCH model is given by:

$$\begin{aligned} \sigma_{i,t+1}^2 &= \omega_i + \alpha_i r_{i,t}^2 + \beta_i \sigma_{i,t}^2 \\ &= \frac{\omega_i}{1 - \beta_i} + \alpha_i \sum_{j=0}^{T} \beta_i{}^j r_{i,t-j}^2 \end{aligned} \tag{7}$$

T represents the length of data used in the estimation. It follows that the k-step ahead forecast has the form:

$$\sigma_{i,t+k/t}^2 = \varpi_i + (\alpha_i + \beta_i)^{k-1}(\sigma_{i,t+1}^2 - \varpi_i) \text{ where } \varpi_i = \frac{\omega_i}{1 - \alpha_i - \beta_i} \tag{8}$$

and hence, are not constant in k. Given these variance forecast functions the average one-day forecast over a quarter (containing N days) is:

$$\sigma_{Ai}^2 = \varpi_i + \frac{1}{\mathrm{N}}(\sigma_{i,t+1}^2 - \varpi_i)\left(\frac{1 - (\alpha_i + \beta_i)^N}{1 - \alpha_i - \beta_i}\right) \quad \text{if } \alpha_i + \beta_i \neq 1 \tag{9}$$

The covariance forecasts are simple functions of these variance forecasts and the parameters ρ_{ij} given the constant correlation assumption.

This constant conditional correlation specification has been used widely in the literature to facilitate estimation given the difficulty in estimating multivariate GARCH models, but its validity is open to debate. The assumption could be justified if the conditional correlation remains constant over time but the market expected returns and variances vary over time. There is evidence of predictable time variations in the equity return distributions; the variance of returns has been shown to be heteroscedastic and univariate GARCH models have had success in modelling returns' variances (see Alexander and Leigh (1997) and Figlewski (1994)). Factors which weigh against the constant correlation assumption include the increased interdependence of international markets (growing integration could lead to increasing correlations through time); the fact that markets tend to be more strongly correlated in times of high volatility than in times of low volatility and the results of our own stability testing discussed previously. Appendix B analyses deviations from the simple constant correlation model. The results imply that this assumption of a constant conditional correlation is questionable; threshold and time-trend models are able to explain movements in conditional correlations. In the following forecasting exercise the results from this analysis should be taken into account. Poor forecasting performance of the constant conditional GARCH model may be due to mis-specification of the model.

The second multivariate GARCH model that we use for forecasting is the Babba, Engle, Kraft and Kroner (BEKK) parameterisation. Engle and Kroner (1995) introduced this model because its quadratic form guarantees that the conditional covariance matrix will be positive definite. The model has the form:

$$H_{t+1} = C'C + B'H_t B + A'R_t R_t' A \tag{10}$$

Matrices A, B and C are the parameter matrices that need to be estimated. R_t is a vector of returns for time t. H_t is the estimated variance-covariance matrix at time t. Again the unconstrained BEKK model is too computationally time consuming for use in this forecasting exercise. To facilitate tractability, a diagonal structure is imposed on the parameter matrices, which removes cross market influences. The model automatically imposes the necessary non-negativity constraints. Rather than producing estimates pair-by-pair, the full model is estimated. Variance forecasts from the BEKK model have the same form as those from the GARCH constant correlation model, with the GARCH parameters being replaced by squared parameters. The covariance forecasts in the BEKK model relax the constant correlation assumption and have the same specification as the constant correlation GARCH variance equations.

4. Forecast Performance

There is a wide literature on modelling financial return variability. To date there has been little agreement in the findings of this literature. West and Cho (1994) compared the out-of-sample forecasting performance of univariate homoskedastic, GARCH, autoregressive and non-parametric models of exchange rate volatility to find that, over a one-week horizon, GARCH models tend to be slightly more accurate. However, for longer forecast horizons West and Cho found that there was little difference in the forecast performance of the various models. Similarly, Brailsford and Faff's (1996) analysis of Australian stock market variability provides some support for the use of GARCH modelling. However, the rankings of the various model forecasts are sensitive to the choice of performance criteria. In contrast Boudoukh, Richardson and Whitelaw (1997), in forecasting the volatility of US interest rates, found that a non-parametric approach outperformed the GARCH model. Campa and Chang (1997), also using foreign exchange data, found that, for shorter time horizons, exponentially weighted moving average models outperform both the fixed-weight historical and GARCH models. However, for longer forecast horizons, fixed-weight models are found to be superior.

More recently the literature has considered forecasts of covariances and correlations. Alexander and Leigh (1997), using equity and foreign exchange data, found that exponentially weighted moving average methods outperform fixed-weight and GARCH methods. It was noted in this study that GARCH models do not perform well when judged by statistical criteria that measure the centre of the distribution. Sheedy (1997) noted that, when comparing various GARCH-type models, the parsimonious models, such as the constant correlation model, perform as well as the more complicated specifications.

One-day-ahead forecasts and quarterly average forecasts are computed by moving the window lengths through the sample and re-estimating the models at each point (we consider moving windows of length 125 days, 250 days, 500 days, 750 days and 1 250 days). Due to the assumption of a zero mean, the one-day-ahead realised matrix will have the form:

$$\hat{\sigma}^2_{ij,t+1} = r_{i,t+1} r_{j,t+1} \tag{11}$$

and the elements of the quarterly (63 days) average realised matrix are calculated as:

$$\hat{\sigma}^2_{ij,AV} = \frac{1}{63} \sum_{t=1}^{63} \sigma^2_{T+t} \tag{12}$$

Since past work has shown that model choice is sensitive to the performance criteria, when comparing the forecasts from the different models a number of performance measures are employed. If we assume conditional normality and zero mean, forecasting the variances and covariances is equivalent to forecasting the probability density function of returns. We can evaluate their accuracy by measuring how well the forecast distribution fits the actual data over the forecast horizon. The greater the log-likelihood function for a given sample the better is the fit of the estimated distribution.

In all cases, the exponentially weighted average log-likelihood falls significantly below the others. This is not unexpected since the parameter in this model is imposed and not estimated. When a window length of 250 days is used, the fixed-weight model outperforms the GARCH models, but when a window length of 1 250 days is used these rankings are reversed. Due to the large number of parameters that the GARCH models need to estimate, the length of data needed to obtain accurate parameter estimates is large. This is illustrated by the increased performance of the GARCH models when 1 250 days of data are used. The BEKK model outperforms the constant correlation based on its log-likelihood. This follows from the increased freedom of the parameterisation of the BEKK model.

Model comparisons based on the log-likelihood are conditional on the assumption of normality. As normality does not hold for many financial-return series, seven distribution-free performance measures are analysed.

Four symmetric performance criteria are considered:

Mean error $$\frac{1}{N}\sum_{t=1}^{N}(\hat{\sigma}_t^2 - \sigma_t^2)$$

Mean absolute error $$\frac{1}{N}\sum_{t=1}^{N}\left|\hat{\sigma}_t^2 - \sigma_t^2\right|$$

Root mean squared error $$\sqrt{\frac{1}{N}\sum_{t=1}^{N}(\hat{\sigma}_t^2 - \sigma_t^2)^2}$$

Mean absolute percentage error $$\frac{1}{N}\sum_{t=1}^{N}\left|\frac{\hat{\sigma}_t^2 - \sigma_t^2}{\sigma_t^2}\right|$$

The mean error offsets the effect of errors of different signs, however, the mean error can be used as a general guide as to the direction of over- or under-prediction. Both the mean absolute error and the root mean squared error (RMSE) focus on the magnitude of errors without taking into account the direction of error, with the RMSE placing greater weight on larger errors. The mean absolute percentage error gives a relative indication of overall forecasting performance.

To account for asymmetry in the loss function we use two error statistics developed by Brailsford and Faff (1996):

Mean mixed error (under) $$\frac{1}{N}\left[\sum_{t=1}^{o}\left|\hat{\sigma}_t^2 - \sigma_t^2\right| + \sum_{t=1}^{u}\sqrt{\hat{\sigma}_t^2 - \sigma_t^2}\right]$$

where o refers to number of over predictions and u to the number of under predictions

Mean mixed error (over) $$\frac{1}{N}\left[\sum_{t=1}^{o}\sqrt{\hat{\sigma}_t^2 - \sigma_t^2} + \sum_{t=1}^{u}\left|\hat{\sigma}_t^2 - \sigma_t^2\right|\right]$$

The mean mixed error (under) penalises under-predictions more heavily while the mean mixed error (over) places greater weight on over-predictions. Finally, to test the efficiency of each model's forecasts we consider the regression R^2 from $\hat{\sigma}_t^2 = \phi + \delta\sigma_t^2$. If the model were fully efficient ϕ would not be significantly different from zero, and δ and the R^2 would both be close to one.

At each point in time the rolling-window estimation results in a forecast variance-covariance matrix to be compared with the actual realisation. Given that the foreign exchange matrix contains 45 elements and the interest rate matrix contains 36 elements, rather than assess each model's forecast performance for each individual variance and covariance a more parsimonious approach was adopted. At each point in time each element of the variance-covariance matrix is treated as a separate observation. The forecast performance measures were then averaged across all observations. The results for the daily forecasts are summarised in Figures 2 and 3. For all criteria, the smaller the number the better (except the R-squared measure). Full details are reported in Appendix C.

Clearly, model choice depends crucially on the metric used. Across criteria, no one model consistently outperforms any other. Given this variation, previous work that relies on one metric should be viewed with caution. In terms of mean error, mean absolute error, root mean squared error, mean under-prediction, and R-squared the simpler models (the fixed-weight and the static exponential models) are preferred. However, the GARCH models tend to do better when the models are assessed against the mean-absolute percentage error. Also the BEKK formulation of the GARCH model tends to produce the lowest average over-prediction (particularly for the foreign exchange data).

Figure 2: Foreign Exchange Series – Daily Forecasts

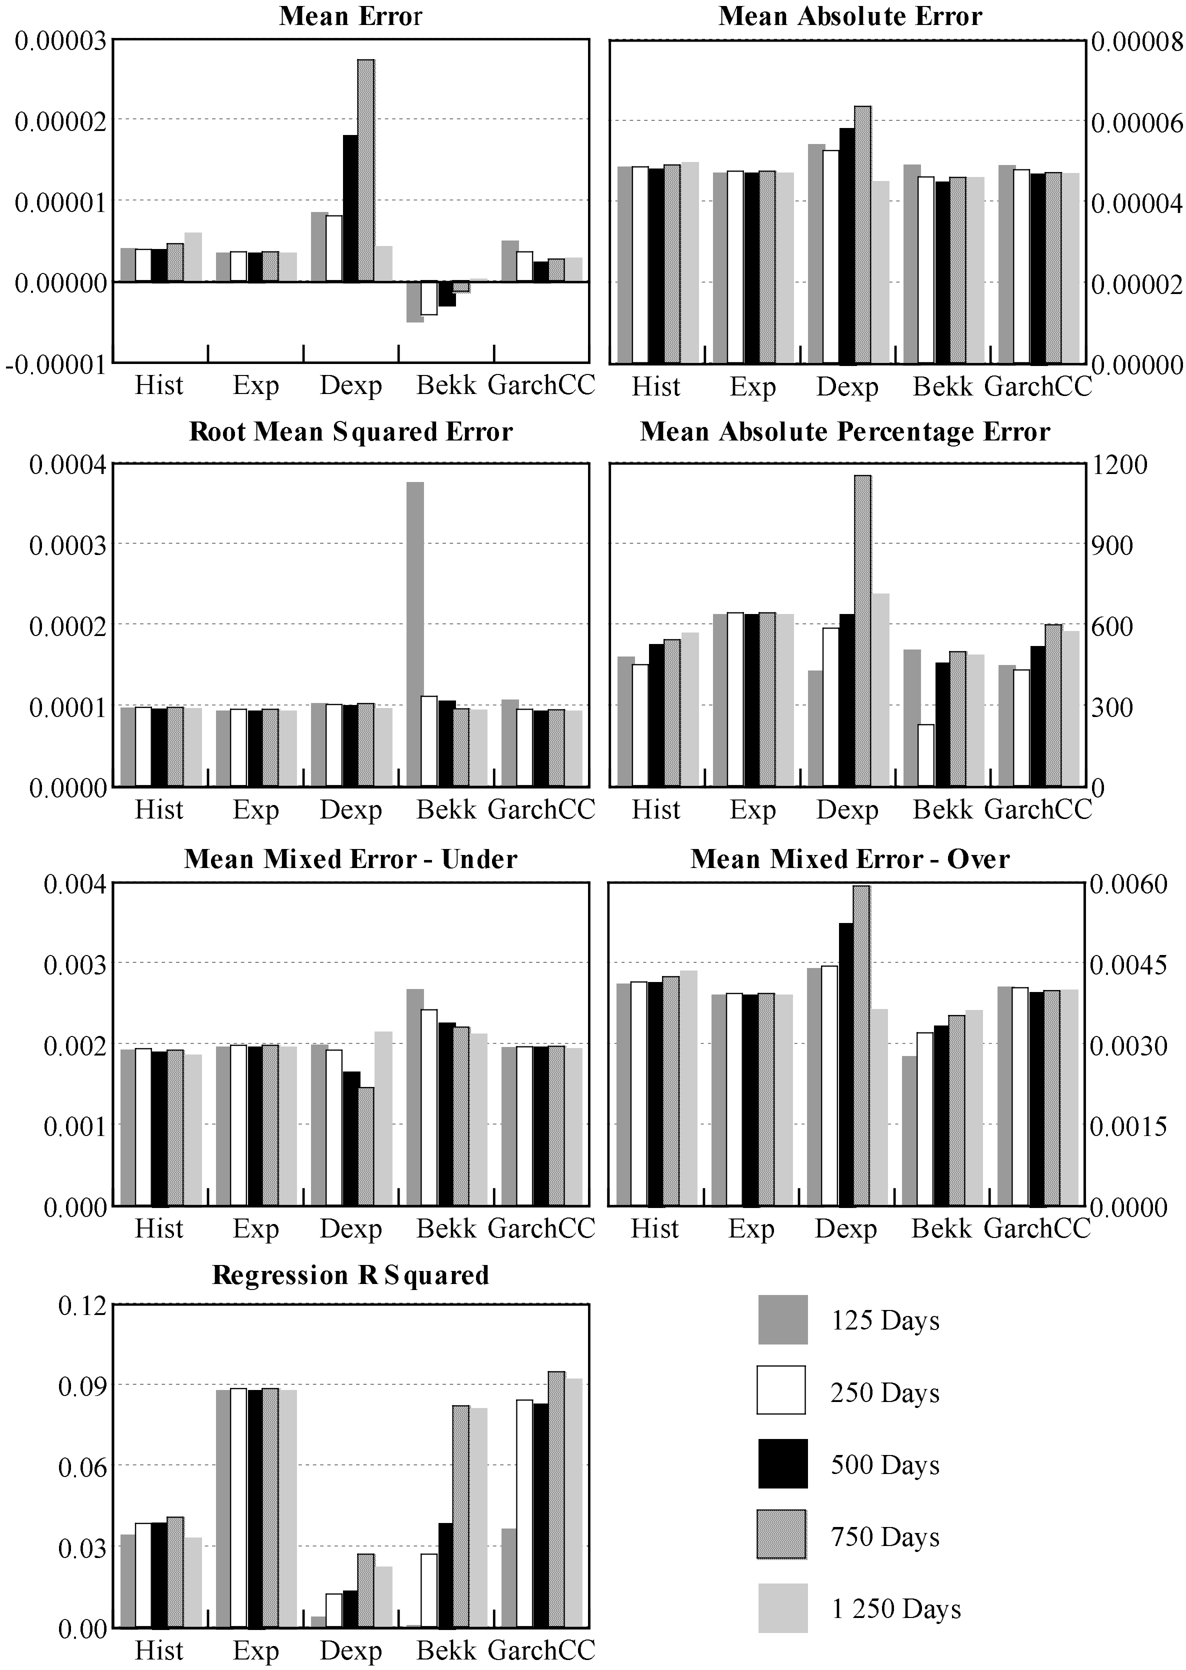

Figure 3: Interest Rate Series – Daily Forecasts

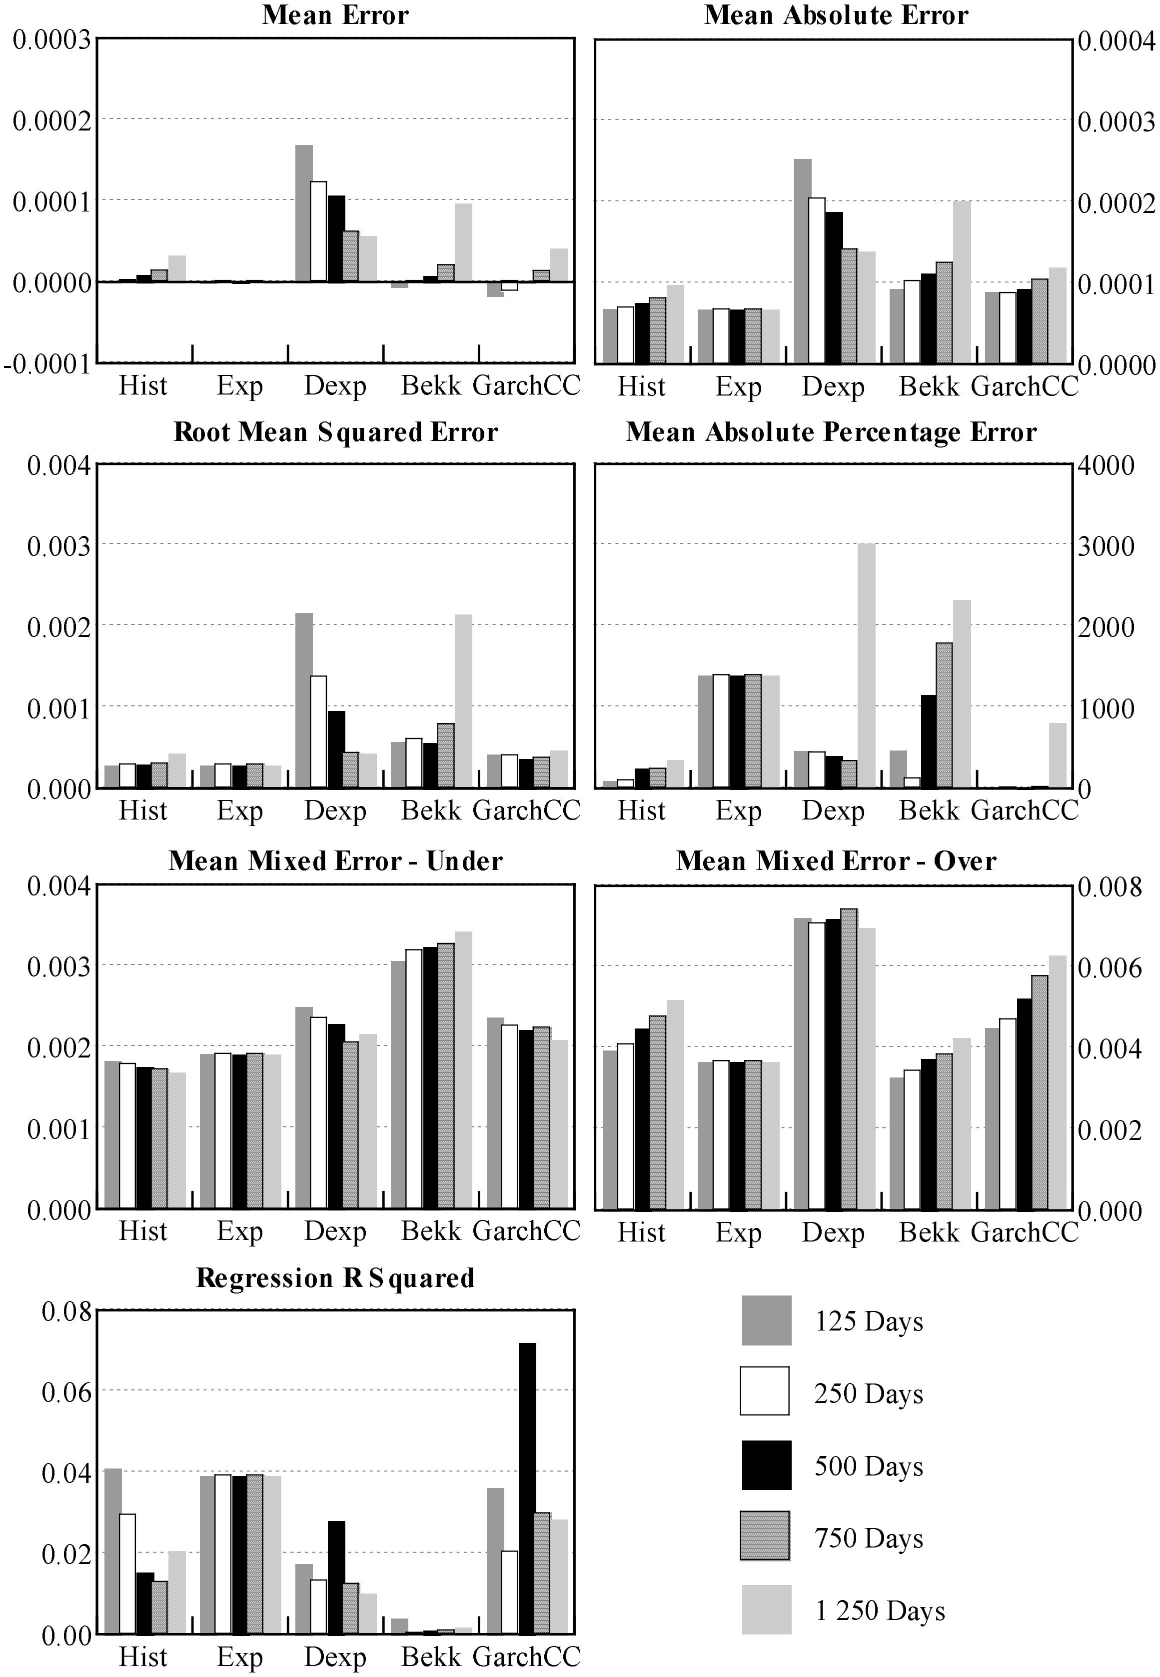

When the different models are compared across all metrics the more complicated GARCH models do not, in general, out-perform the simpler fixed-weight and exponentially weighted moving average approaches. For interest rates, the static exponentially weighted moving average model usually dominates all other models. For all criteria except the MAPE and MMEU the exponentially-weighted model performs best, both when forecasting daily and quarter-average variances and covariances. The fact that the fixed-parameter exponential model, in general, outperforms the dynamic exponential model provides quite strong evidence for the use of the constant parameter simplifying assumption used by RiskMetrics. The sharp decay in the weights of the static exponential model is such that the effective data window is quite short – lengthening the data window has little impact on the forecasted variances and covariances (this can be seen in Figure 3: the performance of the static model is invariant to the window length). This indicates that, for the interest rates' variance-covariance, shorter window lengths provide more efficient forecasts. This finding is consistent with the results of the diagnostic testing of the constant correlation model, which suggest that correlations tend to evolve gradually over time.

The relative performance of GARCH models is strongest in the case of the daily foreign-exchange forecasts. However, the differences in performance across models are not large and for shorter window lengths the simpler models tend to be favoured. The constant-parameter exponential moving average forecasts one-day-ahead variances and covariances performs well while the equally-weighted historical average performs relatively strongly in forecasting quarter-average variances and covariances. Although the simpler models' advantage dissipates as the window length is increased, the more complicated models do not then dominate. In conclusion, the simpler models apparently do not consistently under-perform their more complicated counterparts – in fact, there is some support for the contrary.

To test whether performance levels differ significantly across the five models we use the test of equality of the mean squared error values presented by West and Cho (1994). Under the null hypothesis of equality of mean squared errors, the test statistic has a $\chi^2(4)$ distribution. The tests of equality are carried out on the daily and quarter-average forecasts for each of the different window lengths. Table 3 contains these results.

Table 3: Equality of Mean Squared Errors Across Models

Window length	125	250	500	750	1 250
Foreign exchange					
Daily forecast	257.6*	193.3*	687.5*	1075.4*	50.5*
Quarter average forecast	461.7*	959.0*	2009.2*	4220.3*	1028.0*
Interest rates					
Daily forecast	11797.8*	3304.9*	81.6*	186.2*	201.7*
Quarter average forecast	728.5*	305.1*	130.1*	526.8*	383.7*

Note: * denotes significance at the 1 per cent level

The null of equality is rejected in all cases. To further investigate the differences amongst models we applied West and Cho's test to other groupings of models. In most instances we found that each model produced a mean squared error that differed significantly from that of all other models. When considering daily foreign exchange variances and covariances based on a 125 day window, and daily interest rate forecasts using 125, 250 and 1 250 day calibration windows all models produced significantly different mean squared errors. Similarly, in the case of the quarter-average forecasts only one instance of equality was identified (the dynamic exponential, BEKK and constant-correlation models when applied to interest rates and calibrated on a 500 day window).

For the daily foreign exchange forecasts the model groupings based on mean squared errors vary across the different window lengths. When 250 or 500 day window lengths are used the West and Cho test groups together the dynamic exponential and BEKK models, and the fixed-weight, fixed-parameter exponential and constant-correlation models. Mean squared errors produced by the dynamic exponential, BEKK and fixed-weight models, and the fixed-parameter exponential and constant correlation models do not differ significantly when the 750 day window is used. For the 1 250 day window the fixed-weight and dynamic exponential models, and the fixed-parameter exponential, BEKK and constant-correlation models may be grouped together. In the case of the daily interest rate forecasts (500 and 750 day windows) the dynamic exponential and BEKK models, and the three other models may be grouped together.

In addition to testing across models we used the West and Cho test to test whether, for a given model, the mean squared errors differed significantly across the various data window lengths. The results of this testing are presented in Table 4.

Table 4: Equality of Mean Squared Errors Across Window Lengths

Model	Hist	Exp	Dexp	Bekk	GarchCC
Foreign exchange					
Daily forecast	4.3	4.0	421.2*	11.8**	6.4
Quarter average forecast	1189.9*	4.0	833.2*	5.1	7.0
Interest rates					
Daily forecast	115.2*	4.0	60.8*	9.3	292.3*
Quarter average forecast	217.5*	4.0	89.9*	20.9*	22.4*

Notes: * and ** denotes significance at the 1 per cent and 5 per cent level respectively

The effect of increasing the data window length varies across the different models. For the historical approach, the shorter the length of data used, the better the model (at least down to our smallest window of 125 days). For the quarter-average foreign exchange, daily interest rate and quarter-average interest rate forecasts the differences in mean squared errors are significant. The fixed-parameter exponentially-weighted moving average approach gains no benefit from increasing the data window, since little weight falls on data more than a quarter ago. The GARCH models do not consistently favour longer window lengths. For instance, in the case of the BEKK model, increasing the window length significantly reduces mean squared errors for daily foreign exchange forecasts, but significantly increases mean squared errors for quarter-average interest rate forecasts.

The fact that the forecast performance of the dynamic exponentially moving average and the GARCH models do not systematically improve as the length of data used for model estimation increases is a little surprising. Increased data length should provide more accurate parameter estimates. The fact that more precise parameter estimates are not resulting in more precise forecasts suggests that these may not be appropriate models for this purpose and that other classes of models should be considered.

Further analysis of forecast errors for the individual variances and covariances shows that much of the relatively poor forecasting performance of the GARCH models can be attributed to extremely poor prediction of a small number of elements within the variance-covariance matrix (for example, the variance of the Australian dollar – New Zealand dollar exchange rate). When these elements are removed, however, the GARCH models still do not outperform the simpler models. The simpler models exhibit fairly constant behaviour across the elements of the variance-covariance matrix. This is consistent across all forecast-error metrics. As the sample size is increased these simple models perform better on average across all matrix elements, but the dispersion around this average increases. This behaviour is found in both the daily and quarter-average results. For the GARCH models there does not seem to be a consistent relationship between the length of data window and the dispersion of forecasting accuracy across the elements of the variance-covariance matrix. At times increasing the data window increases the spread of forecast metrics across the matrix elements, while at other times the spread is reduced.

The users of a VaR model are not likely to view over- and under-prediction of variances and covariances equally. A model that consistently over-predicts volatility will overstate a portfolio's risk. This may be attractive to supervisors who may prefer models to err on the conservative side. However, individual traders within a firm may prefer models which under-predict risk and thus overstate risk-adjusted returns. It is not clear whether the banking firm as a whole would prefer a model that over- or under-predicts risk. The capital allocation flowing from an overly conservative model will be more expensive and hurdle rates of return unnecessarily high. Against this, a model that consistently under-predicts will expose the bank to an unexpectedly high probability of bankruptcy. We need to consider prediction of variances and correlations separately. Over-estimation of variances unambiguously over-predicts true risk. The effect of over-prediction of correlation depends upon the composition of the portfolio subject to the VaR model. Hence, we separately consider three measures of forecast bias in variance prediction: the mean error, and the mean mixed errors, both over and under. These are shown in Figures 4 and 5. Full details are given in Appendix C.

Figure 4: Foreign Exchange Variance

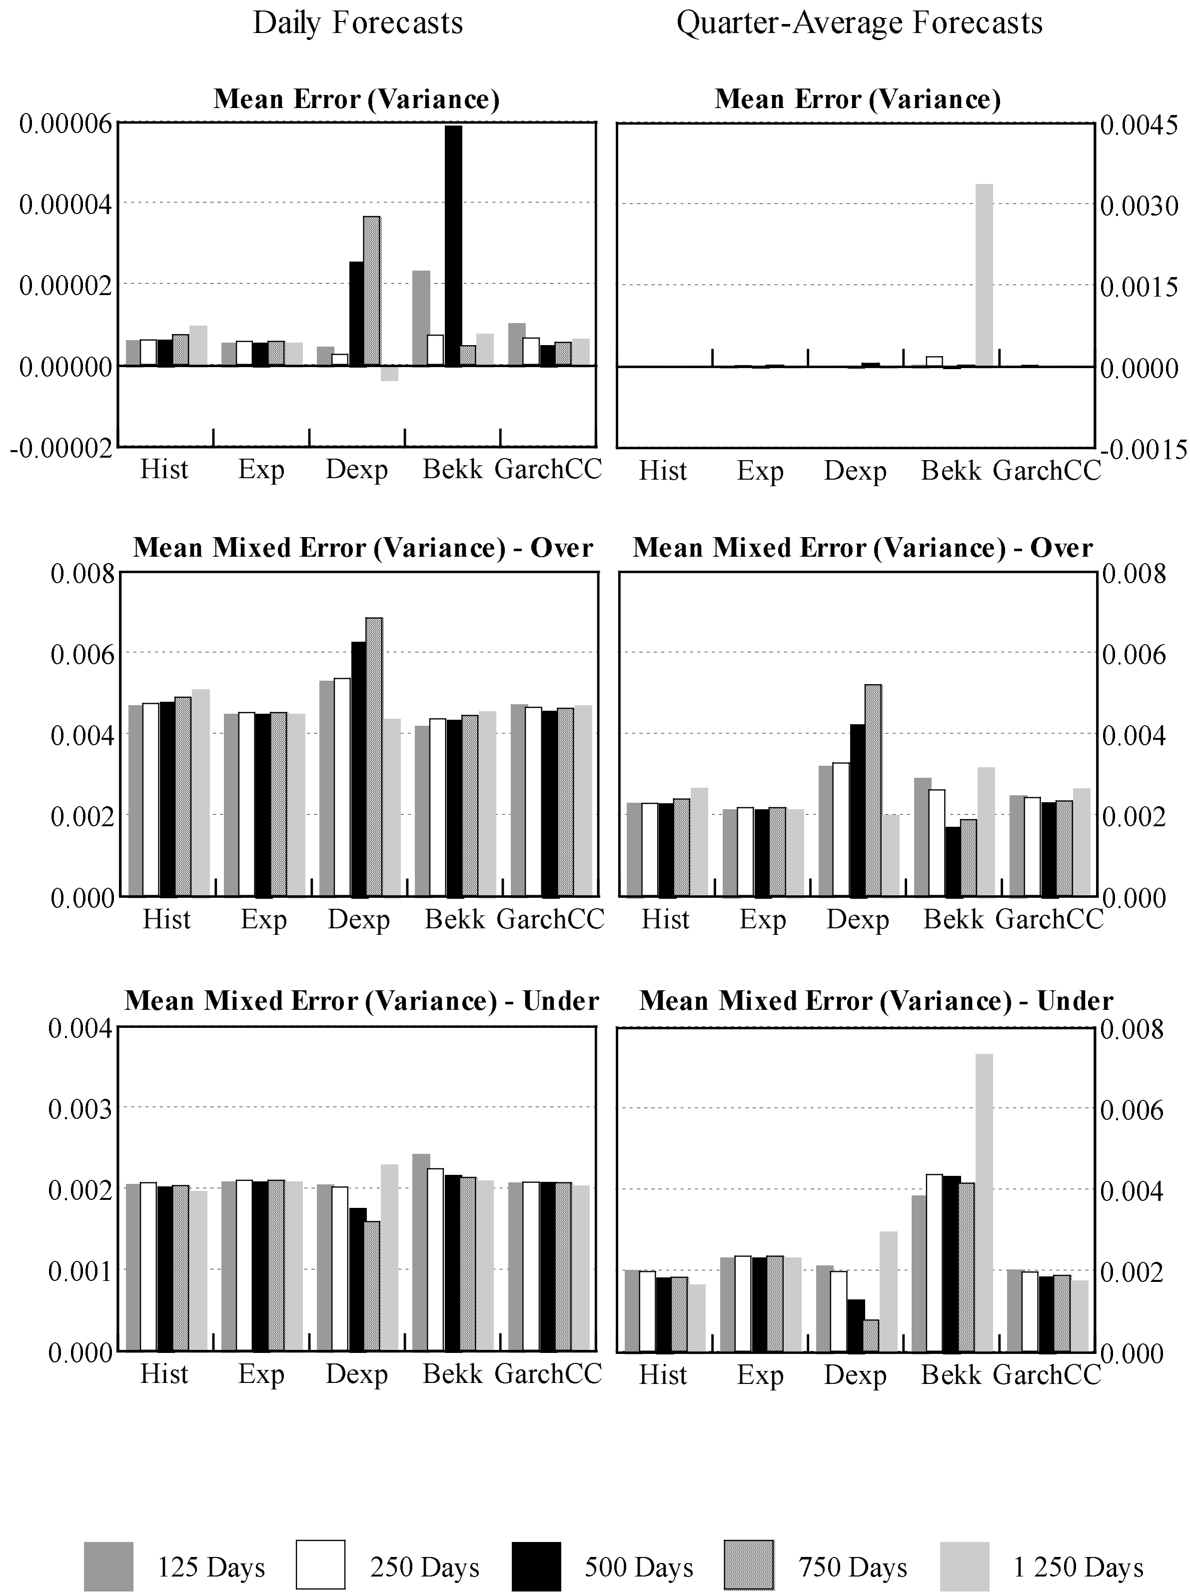

Figure 5: Interest Rates Variances

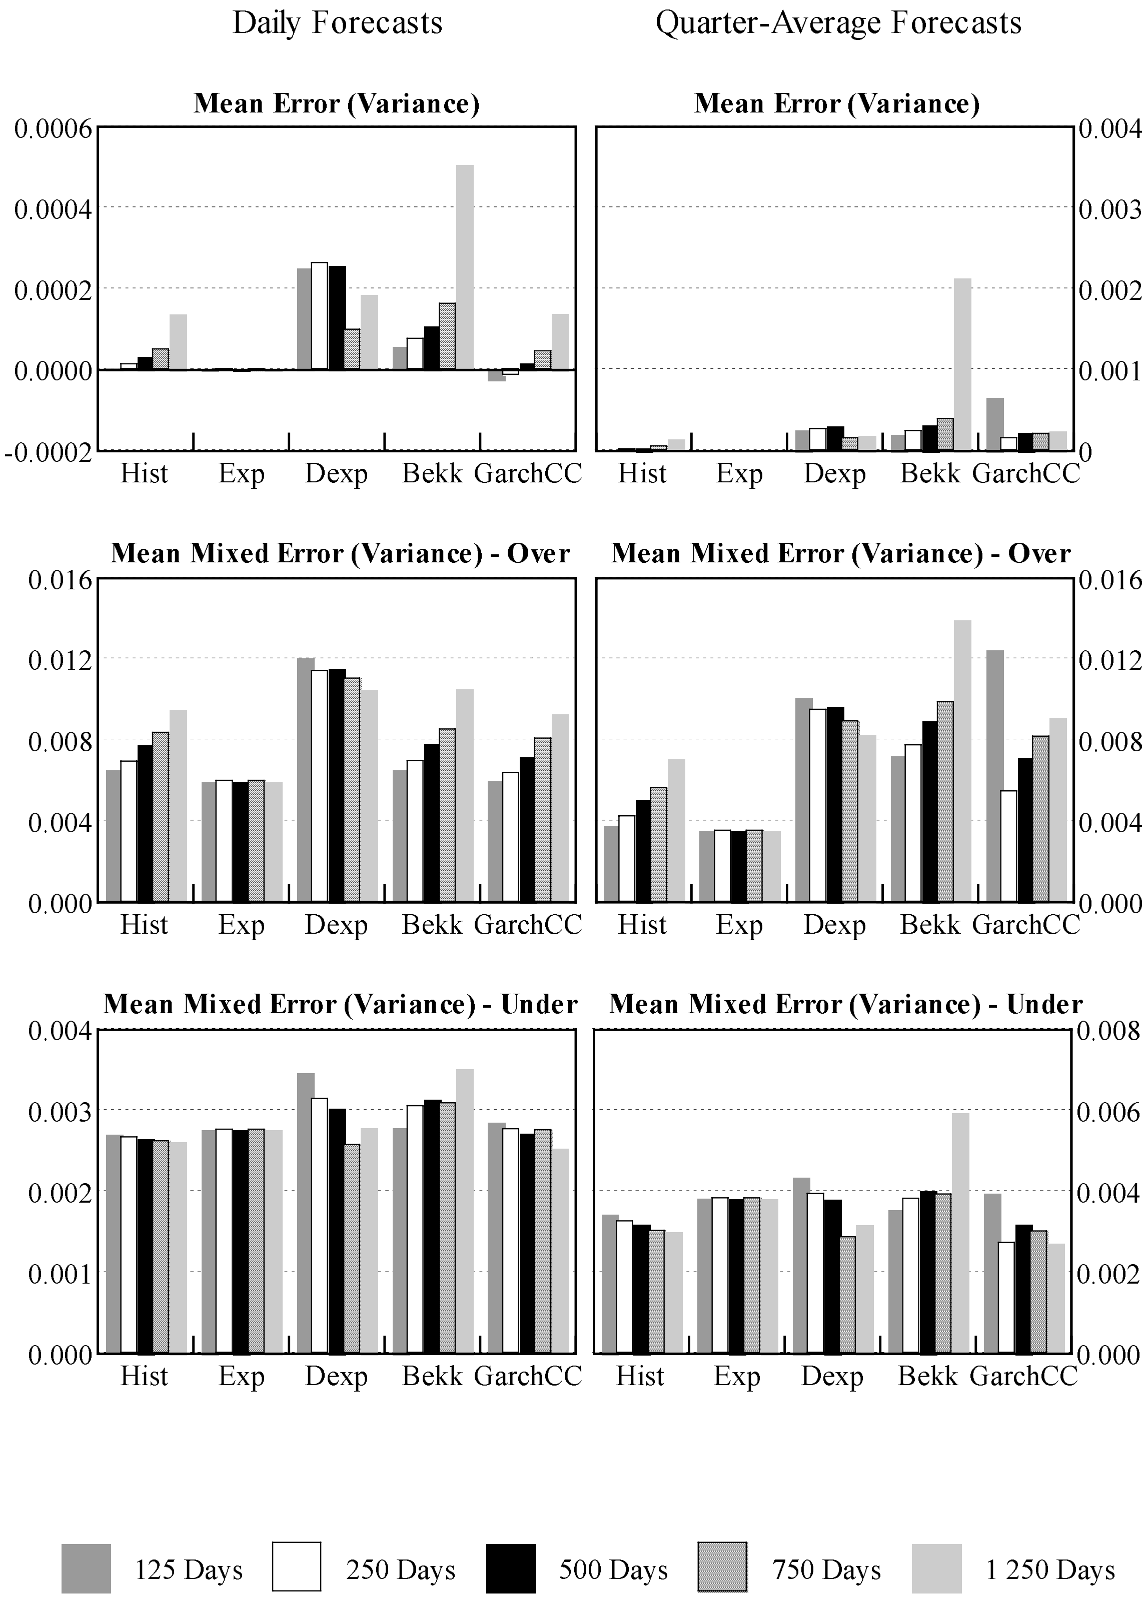

Consistent with the previous results, the fixed-parameter exponentially-weighted moving average tends to outperform the other models when forecasting interest rates, producing low mean-error and low over-prediction results. For daily and quarterly interest-rate variance forecasts the historical and constant-correlation GARCH models respectively provide the least under-prediction and thus can be taken to be the most conservative models. The results are more mixed for foreign exchange with no model consistently outperforming the others.

As we noted earlier, the frequency with which banks re-estimate their variance-covariance matrix varies. While banks are required to re-estimate the matrix at least quarterly for regulatory purposes, many banks update their matrix each day. To gauge the impact of less frequent re-estimation on forecast performance we compare the root mean squared error of one-day-ahead forecasts with that of the forecast for the last day of the quarter. The percentage increase in root mean squared error is shown in Table 5.

The first three models predict that future variances and covariances will remain constant. The forecast accuracy of these simple models declines by something in the range of 2 to 10 per cent, which in comparison with the other measurement errors embodied in VaR models is probably not large (see, for example, Gizycki and Hereford (1998)). In contrast, the GARCH models (in which the future path of variances and covariances follows a smooth decay function) perform much more poorly over the longer forecast horizon. The poor performance of the models calibrated on the shorter data periods may be attributed to imprecision in estimation of the model parameters. However, the forecast accuracy of the more robust long-window estimates declines by as much as a half over the quarter. In several instances we obtain the unusual result that better forecast accuracy is obtained for the longer horizon than the day-ahead forecasts.

Table 5: Increase in RMSE over the Quarter

Per cent

	Foreign exchange	Interest rates
Fixed-weight historical		
125	2.21	2.12
250	-1.42	-1.14
500	0.44	0.25
750	-2.67	-2.92
1 250	2.32	2.18
Exponential smoothing – fixed-parameter		
125	8.03	8.00
250	4.76	5.05
500	5.21	5.15
750	1.17	1.07
1 250	6.23	6.27
Dynamic exponential smoothing		
125	11.19	11.32
250	7.39	7.84
500	6.12	6.17
750	2.46	2.44
1 250	6.73	6.80
Constant correlation GARCH		
125	2 712.01	3 360.55
250	21.58	29.51
500	-4.74	-4.65
750	25.39	25.08
1 250	39.38	39.38
BEKK GARCH		
125	33.21	38.53
250	238.23	293.77
500	-1.53	-1.39
750	50.69	50.41
1 250	50.50	51.52

Note: For each model and each element of the variance-covariance matrix the root mean squared error was calculated for forecasts of the variance or covariance to be observed on the next day and the last day of the quarter. This table shows the difference between the two root mean squared errors expressed as a percentage of the one-day-ahead root mean squared error. This ratio has been averaged across all elements of the variance-covariance matrix.

5. Conclusion

There are two principal conclusions to be drawn from this paper. Firstly that there is considerable variation in variances and, to a lesser extent, correlations over time. This provides support for APRA's approach to market-risk measurement, which requires that banks update the parameters of the variance-covariance matrix at least quarterly. That said, moving from quarterly to daily updating of parameters only slightly improves forecast accuracy. Secondly, the cost, in terms of forecasting accuracy, of using simple models of the time-series evolution of variances and covariances does not appear to be high. Simple models such as equally-weighted moving averages and fixed-parameter exponentially-weighted moving averages appear to perform as well or often better than more complex GARCH models.

These findings provide some support for the practices prevalent in the Australian banking industry where comparatively simple models of variances and covariances are employed in formulating VaR models and the financial return data underlying the VaR models are periodically updated, often daily but at least quarterly.

There are a number of caveats to the application of the results of this forecasting exercise to the practical assessment of VaR models. Modelling variances and covariances is just one component of a VaR model, which must also address the overall distribution of financial returns, portfolio composition and measurement of the sensitivity of financial instruments to movements in underlying prices.

If it is assumed that financial returns follow a normal distribution the mean returns and the variance-covariance matrix are sufficient to describe the full distribution of financial returns. There is strong evidence to suggest, however, that financial returns are not normally distributed. In such a case a model that forecasts covariances well, will accurately forecast behaviour around the centre of the distribution but not necessarily in the distribution's tails. The focus of VaR models in measuring market-risk is on the extreme tails of the distribution (typically the first or fifth percentiles). Going beyond the variance-covariance VaR model, banks have developed a range of VaR models (such as historical simulation and Monte Carlo simulation) which place stronger emphasis on modelling the extremes of financial return distributions. Clearly more work remains to be done comparing the tail forecasting performance of various variance-covariance VaR formulations with that of other VaR models.

In this paper, equal weight has been given to each variable within the variance-covariance matrix. In practice, banks' portfolios tend to be concentrated in a small number of assets. For example, the bulk of banks foreign exchange exposure may derive from trading in major currencies such as the US dollar, German mark and Japanese yen. Accurate forecasting of the variability of these rates will be much more important than for other less actively traded currencies. The assessment of the forecasting performance of VaR models needs to be calibrated against the composition of banks' portfolios.

Accurate prediction of the variability of the value of a portfolio requires an accurate forecast of the probability of larger moves in market prices and precise measurement of the sensitivity of the value of various instruments to those larger price moves. For most simple instruments, such as spot and forward foreign exchange and bonds, measurement of price-sensitivity is a straightforward matter. In the case of complex instruments, such as options, however, there remains wide variation in the practices adopted to incorporate them into a VaR framework. Further research on these issues remains to be done.

Appendix A: Stability Analysis

Table A1: Foreign Exchange

The proportion of stable matrices; per cent

	Covariances					Correlation				
Window length	125	250	500	750	1 250	125	250	500	750	1 250
USD versus										
DEM	26	0	17	33	0	70	38	50	33	0
JPY	30	0	17	33	0	74	38	50	33	0
GBP	33	15	17	33	0	89	54	50	33	0
CHF	33	15	17	33	0	89	54	50	33	0
NLG	44	23	50	33	0	100	69	83	67	0
NZD	44	23	67	33	0	100	69	83	67	0
FRF	44	23	67	33	0	100	69	83	67	0
CAD	48	31	67	33	0	100	100	83	67	0
DEM versus										
JPY	62	38	67	33	0	100	100	83	67	100
GBP	67	53	83	67	0	100	100	100	100	100
CHF	74	53	83	67	0	100	100	100	100	100
NLG	74	53	83	67	0	100	100	100	100	100
NZD	74	53	83	100	0	100	100	100	100	100
FRF	74	53	83	100	0	100	100	100	100	100
CAD	78	53	83	100	100	100	100	100	100	100
JPY versus										
GBP	78	62	83	100	100	100	100	100	100	100
CHF	78	62	83	100	100	100	100	100	100	100
NLG	78	62	83	100	100	100	100	100	100	100
NZD	78	62	83	100	100	100	100	100	100	100
FRF	78	62	83	100	100	100	100	100	100	100
CAD	78	62	83	100	100	100	100	100	100	100

Table A2: Interest Rates

The proportion of stable matrices; per cent

	Covariances					Correlations				
Window length	125	250	500	750	1 250	125	250	500	750	1 250
BAB 90-day versus										
Overnight cash	22	7	0	0	0	70	84	5	67	100
BAB 30-day	59	30	0	0	0	100	100	100	100	100
BAB 180-day	63	61	0	0	0	100	100	100	100	100
One-year bond	71	69	0	0	0	100	100	100	100	100
Two-year bond	70	69	17	0	0	100	100	100	100	100
Five-year bond	70	69	17	0	0	100	100	100	100	100
Ten-year bond	70	69	17	0	0	100	100	100	100	100

Appendix B: Conditional Correlation Analysis

Given the questions surrounding the constant correlation assumption various departures from the model are tested. Three alternative models are estimated along the lines of Longin and Solnik (1995). The constant correlation model is augmented with a time trend, threshold variables and asymmetry variables. The models are estimated over the full sample; over two halves of the sample; over four quarters of the full sample; and over consecutive 500 day periods. To simplify the analysis only bivariate systems involving the AUD\USD paired with the eight other exchange rates and 90-day bank bill rate paired with the other interest rate series are used.

The covariance equation in the model augmented with a time trend (chosen arbitrarily as linear) has the form:

$$\sigma_{ij,t+1} = (\rho_{ij0} + \rho_{ij1}t)\sigma_{i,t+1}\sigma_{j,t+1} \quad \text{(B1)}$$

If the coefficient on the time trend is significantly different from zero then this provides evidence that the conditional correlation is not constant across time. Table B1 summarises the results from this exercise.

Table B1: Testing Constant Correlation GARCH Functional Form

	Full sample	Half sample	Quarter sample	500 days
Time trend				
Foreign exchange	13.56*	6.23*	4.88*	2.75*
Interest rates	4.83*	4.03*	3.73*	2.43*
Threshold				
Foreign exchange	6.28*	4.74*	3.14*	2.14
Interest rates	4.31*	5.02*	3.47*	2.03
Asymmetry				
Foreign exchange ($\phi_{ij1} \neq \phi_{ij2}$)	2.48*	1.58	1.83	1.43
Foreign exchange ($\phi_{ij0} \neq \phi_{ij3}$)	2.24	0.97	1.47	0.94
Interest rates ($\phi_{ij1} \neq \phi_{ij2}$)	2.38	1.99	1.53	1.39
Interest rates ($\phi_{ij0} \neq \phi_{ij3}$)	2.29	1.44	1.45	1.13

Notes: The results reported are the average standard t-statistic on the ρ_{ij} coefficient across time periods and all elements in the variance-covariance matrix.
* denotes significance at the 5 per cent level.

The time trend coefficient is significant across all samples. Over the full period, therefore, covariances have significantly changed. It is worth noting that as the sample size decreases the time trend becomes less significant. To this extent smaller data windows, when no time trend is included, would be expected to give a more accurate representation of the evolution of the underlying process.

To test the hypothesis that correlations increase during periods of high volatility, a threshold effect is introduced into the bivariate system. With this threshold on correlation, the covariance term of the GARCH specification can be written as:

$$\sigma_{ij,t+1} = (\rho_{ij0} + \xi_{ij1} S_t)\sigma_{i,t+1}\sigma_{j,t+1} \qquad \text{(B2)}$$

where S_t is a dummy variable that takes the value 1 if the estimated conditional variance of the USD/AUD exchange rate return is greater than its unconditional value and 0 otherwise. The threshold for the interest rate series is the variance of the change in the 90-day bank bill rate. This choice of threshold is arbitrary. The unconditional variance of innovations from the base model is taken as the exogenous threshold. The coefficient ξ_{ij1} will be positive if the correlation increases when the conditional variance is high. The average t-statistic on the dummy variable coefficient is shown in the table above.

The threshold coefficient is significantly different from zero in all cases except the 500 day window, hence, for the larger sized samples a threshold effect is present. This suggests that the use of shorter data windows may compensate for the failure to explicitly model a threshold effect. The threshold coefficients are significantly positive for all pairs of foreign exchange rates and longer-term interest rates indicating that periods of high volatility are associated with increased correlations.[6] This is supportive of the stress-testing approach set out in APRA's market-risk reporting requirements which assumes that 'worst case' price movements occur simultaneously across a range of markets.

---

6 The threshold coefficients are significantly negative for short-term interest rates. Hence, increased variation is associated with decreased correlation. This may reflect the interaction of the operation of monetary policy with market expectations.

The third model looks at the issue of asymmetry. In the previous models a negative or positive shock is assumed to have the same impact on correlation. Here we test whether negative and positive shocks have a different impact on the conditional correlation. This augmented model conditions the correlation on both the sign and magnitude of past shocks to the USD/AUD exchange rate and for the interest rate matrix the 90-day bank bill rate. This asymmetric correlation GARCH has the following form:

$$\sigma_{ij,t+1} = (\phi_{ij0}S_{1,t} + \phi_{ij1}S_{2,t} + \phi_{ij2}S_{3,t} + \phi_{ij3}S_{4,\text{t}})\sigma_{i,t+1}\sigma_{j,t+1} \quad \text{(B3)}$$

where $S_{k,\,t}$ are the dummy variables that take the values:

$S_{1,\,t} = 1$ if $\varepsilon_{i,\,t}$ is less than $-\sigma_i$ and zero otherwise
$S_{2,\,t} = 1$ if $\varepsilon_{i,\,t}$ is less than 0 and zero otherwise
$S_{3,\,t} = 1$ if $\varepsilon_{i,\,t}$ is greater than 0 and zero otherwise
$S_{4,\,t} = 1$ if $\varepsilon_{i,\,t}$ is greater than σ_i and zero otherwise

Asymmetry is captured when $\phi_{ij1} \neq \phi_{ij2}$ *and* $\phi_{ij0} \neq \phi_{ij\,3}$. As can be seen from the results asymmetry does not appear to be present. This is consistent with the literature regarding foreign exchange series (Sheedy (1997)).

Appendix C: Forecast Performance

An asterisk indicates best performance. The numbers presented are multiplied by 1000 except for the MAPE and R squared criteria.

Table C1: Foreign Exchange

Daily

	Hist	Exp	Dexp	BEKK	GARCHCC
125 day window					
ME	0.004*	0.004	0.009	-0.005	0.005
MAE	0.049	0.047*	0.054	0.049	0.049
RMSE	0.098	0.094*	0.103	0.377	0.108
MAPE	482.510	640.474	431.204	508.773	452.449*
MMEU	1.939*	1.975	1.994	2.683	1.968
MMEO	4.122	3.926	4.421	2.778*	4.069
R^2	0.035	0.088*	0.004	0.001	0.037
250 day window					
ME	0.004	0.004*	0.008	-0.004	0.004
MAE	0.048	0.047	0.053	0.046*	0.048
RMSE	0.097	0.094*	0.100	0.111	0.094
MAPE	448.560	640.634	584.847	227.159*	429.216
MMEU	1.934	1.974	1.917*	2.415	1.957
MMEO	4.133	3.928	4.427	3.196*	4.029
R^2	0.038	0.088*	0.012	0.027	0.084
500 day window					
ME	0.004	0.004	0.018	-0.003*	0.003
MAE	0.048	0.047	0.058	0.045*	0.047
RMSE	0.097	0.094*	0.101	0.107	0.094
MAPE	530.773	640.634	639.489	461.536*	521.527
MMEU	1.915	1.974	1.666*	2.269	1.973
MMEO	4.163	3.928	5.256	3.346*	3.971
R^2	0.039	0.088*	0.014	0.039	0.083
750 day window					
ME	0.005	0.004	0.027	-0.001*	0.003
MAE	0.049	0.047	0.063	0.046*	0.047
RMSE	0.097	0.094*	0.101	0.095	0.094
MAPE	541.055	640.634	1 151.385	496.691*	597.006
MMEU	1.914	1.974	1.453*	2.197	1.965
MMEO	4.233	3.928	5.915	3.513*	3.981
R^2	0.041	0.088	0.027	0.082	0.095*
1 250 day window					
ME	0.006	0.004	-0.004	0.000*	0.003
MAE	0.050	0.047	0.045*	0.046	0.047
RMSE	0.097	0.094*	0.097	0.095	0.094
MAPE	570.695	640.634	717.589	490.637*	578.418
MMEU	1.877*	1.974	2.156	2.133	1.954
MMEO	4.372	3.928	3.655	3.634*	4.014
R^2	0.033	0.088	0.023	0.081	0.092*

Table C2: Foreign Exchange

Quarter-average

	Hist	Exp	Dexp	BEKK	GARCHCC
125 day window					
ME	0.000*	-0.001	0.004	-0.010	0.002
MAE	0.019*	0.021	0.027	0.043	0.022
RMSE	0.030*	0.033	0.042	0.441	0.125
MAPE	1.370	1.693	5.663	1.209*	1.312
MMEU	1.802*	2.099	2.041	4.368	1.847
MMEO	2.039	1.930	2.554	0.984*	2.099
R^2	0.335*	0.327	0.104	0.002	0.025
250 day window					
ME	0.000*	-0.001	0.004	0.018	0.001
MAE	0.017*	0.021	0.024	0.066	0.019
RMSE	0.027*	0.033	0.035	3.388	0.062
MAPE	2.671	1.693*	1.873	2.188	2.072
MMEU	1.698*	2.097	1.809	3.960	1.714
MMEO	1.960	1.932	2.539	1.054*	2.008
R^2	0.389*	0.327	0.176	0.000	0.096
500 day window					
ME	0.000*	-0.001	0.014	-0.015	0.000
MAE	0.017*	0.021	0.026	0.029	0.017
RMSE	0.025*	0.033	0.035	0.106	0.027
MAPE	2.887	1.693	4.096	0.997*	2.688
MMEU	1.637	2.097	1.147*	3.713	1.663
MMEO	1.966	1.932	3.498	0.952*	1.960
R^2	0.442*	0.327	0.237	0.023	0.410
750 day window					
ME	0.000*	-0.001	0.023	-0.011	0.000
MAE	0.017*	0.021	0.030	0.028	0.017
RMSE	0.025*	0.033	0.036	0.108	0.027
MAPE	3.360	1.693	7.563	1.379*	3.344
MMEU	1.643	2.097	0.637*	3.392	1.685
MMEO	2.044	1.932	4.542	1.133*	1.970
R^2	0.448*	0.327	0.357	0.026	0.401
1 250 day window					
ME	0.002	-0.001	-0.009	0.668	0.001*
MAE	0.018*	0.021	0.019	0.703	0.018
RMSE	0.026*	0.033	0.030	96.804	0.029
MAPE	3.473	1.693*	3.527	21.811	3.319
MMEU	1.522*	2.097	2.499	3.766	1.614
MMEO	2.221	1.932	1.390*	1.523	2.147
R^2	0.401*	0.327	0.284	0.000	0.309

Table C3: Interest Rates

Daily

	Hist	Exp	Dexp	BEKK	GARCHCC
125 day window					
ME	0.000*	-0.001	0.168	-0.006	-0.017
MAE	0.068	0.067*	0.253	0.093	0.089
RMSE	0.283*	0.284	2.161	0.570	0.415
MAPE	90.571	1 388.708	464.170	469.233	5.818*
MMEU	1.824*	1.905	2.484	3.055	2.361
MMEO	3.932	3.654	7.205	3.274*	4.498
R^2	0.041*	0.039	0.017	0.004	0.036
250 day window					
ME	0.001*	-0.001	0.122	-0.001	-0.012
MAE	0.069	0.067*	0.204	0.102	0.087
RMSE	0.287	0.284*	1.367	0.600	0.399
MAPE	90.406	1 388.757	436.538	113.641	5.566*
MMEU	1.779*	1.904	2.346	3.180	2.254
MMEO	4.078	3.656	7.054	3.419*	4.693
R^2	0.029	0.039*	0.013	0.000	0.020
500 day window					
ME	0.008	-0.001*	0.106	0.007	0.001
MAE	0.076	0.067*	0.187	0.111	0.093
RMSE	0.296	0.284*	0.951	0.553	0.361
MAPE	243.407	1 388.757	398.508	1 147.206	6.641*
MMEU	1.746*	1.904	2.277	3.225	2.201
MMEO	4.488	3.656*	7.180	3.722	5.215
R^2	0.015	0.039	0.028	0.001	0.072*
750 day window					
ME	0.013	-0.001*	0.061	0.020	0.013
MAE	0.080	0.067*	0.141	0.124	0.104
RMSE	0.301	0.284*	0.429	0.785	0.367
MAPE	231.027	1 388.757	326.254	1 778.201	9.810*
MMEU	1.715*	1.904	2.048	3.256	2.228
MMEO	4.756	3.656*	7.401	3.824	5.750
R^2	0.013	0.039*	0.012	0.001	0.030
1 250 day window					
ME	0.032	-0.001*	0.056	0.096	0.041
MAE	0.097	0.067*	0.138	0.200	0.119
RMSE	0.424	0.284*	0.428	2.140	0.469
MAPE	346.281*	1 388.757	3 019.276	2 321.331	806.829
MMEU	1.684*	1.904	2.156	3.422	2.081
MMEO	5.177	3.656*	6.957	4.248	6.279
R^2	0.020	0.039*	0.010	0.002	0.028

Table C4: Interest Rates

Quarter-average

	Hist	Exp	Dexp	BEKK	GARCHCC
125 day window					
ME	0.000*	0.000	0.169	0.028	0.175
MAE	0.033*	0.037	0.219	0.092	0.225
RMSE	0.081*	0.082	2.166	0.685	2.348
MAPE	4.626	6.410	17.833	2.361*	39.055
MMEU	2.114*	2.374	2.971	4.113	3.060
MMEO	2.225	2.161*	5.573	2.433	5.959
R^2	0.191	0.209*	0.036	0.017	0.000
250 day window					
ME	0.002	0.000*	0.123	0.036	0.040
MAE	0.034*	0.037	0.170	0.103	0.068
RMSE	0.090	0.082*	1.351	0.756	1.083
MAPE	5.597	6.412	14.340	3.256*	12.369
MMEU	2.002	2.373	2.766	4.217	1.991*
MMEO	2.343	2.162*	5.352	2.559	3.417
R^2	0.167	0.209*	0.049	0.017	0.006
500 day window					
ME	0.009	0.000*	0.107	0.053	0.062
MAE	0.041	0.037*	0.154	0.123	0.098
RMSE	0.110	0.082*	0.935	0.884	0.880
MAPE	7.082	6.412	15.057	3.701*	11.379
MMEU	2.010	2.373	2.700	4.298	2.282*
MMEO	2.774	2.162*	5.524	2.869	4.365
R^2	0.126	0.209*	0.054	0.032	0.013
750 day window					
ME	0.014	0.000*	0.062	0.069	0.062
MAE	0.044	0.037*	0.098	0.139	0.098
RMSE	0.122	0.082*	0.313	1.058	0.726
MAPE	8.467	6.412*	14.021	14.118	12.450
MMEU	1.888*	2.373	2.120	4.321	2.187
MMEO	3.057	2.162*	5.631	3.047	4.983
R^2	0.107	0.209*	0.052	0.009	0.013
1 250 day window					
ME	0.033	0.000*	0.056	0.453	0.073
MAE	0.063	0.037*	0.096	0.527	0.106
RMSE	0.345	0.082*	0.308	10.772	0.565
MAPE	10.097	6.412*	17.106	8.469	12.484
MMEU	1.814*	2.373	2.284	4.826	2.050
MMEO	3.572	2.162*	5.099	3.876	5.394
R^2	0.053	0.209*	0.054	0.008	0.051

Table C5: Foreign Exchange Variances

Variance – daily

	Hist	Exp	Dexp	BEKK	GARCHCC
125 day window					
ME	0.00001	0.00001	0.00000	0.00002	0.00001
MMEU	0.00210	0.00210	0.00210	0.00243	0.00208
MMEO	0.00470	0.00451	0.00530	0.00421	0.00473
250 day window					
ME	0.00001	0.00001	0.00000	0.00001	0.00001
MMEU	0.00212	0.00210	0.00203	0.00224	0.00207
MMEO	0.00473	0.00451	0.00532	0.00435	0.00464
500 day window					
ME	0.00001	0.00001	0.00003	0.00006	0.00001
MMEU	0.00201	0.00210	0.00184	0.00217	0.00209
MMEO	0.00482	0.00451	0.00631	0.00437	0.00459
750 day window					
ME	0.00001	0.00001	0.00004	0.00000	0.00001
MMEU	0.00203	0.00210	0.00163	0.00213	0.00206
MMEO	0.00494	0.00451	0.00682	0.00444	0.00461
1 250 day window					
ME	0.00001	0.00001	0.00000	0.00001	0.00001
MMEU	0.00201	0.00210	0.00232	0.00211	0.00204
MMEO	0.00513	0.00451	0.00444	0.00458	0.00472

Table C6: Foreign Exchange Variances

Variance – quarter-average

	Hist	Exp	Dexp	BEKK	GARCHCC
125 day window					
ME	0.00000	0.00000	0.00000	0.00004	0.00001
MMEU	0.00201	0.00235	0.00212	0.00388	0.00205
MMEO	0.00232	0.00216	0.00322	0.00293	0.00251
250 day window					
ME	0.00000	0.00000	0.00000	0.00017	0.00000
MMEU	0.00202	0.00235	0.00201	0.00435	0.00196
MMEO	0.00233	0.00216	0.00332	0.00260	0.00242
500 day window					
ME	0.00000	0.00000	0.00002	-0.00001	0.00000
MMEU	0.00192	0.00235	0.00133	0.00436	0.00188
MMEO	0.00234	0.00216	0.00434	0.00172	0.00232
750 day window					
ME	0.00000	0.00000	0.00003	-0.00001	0.00000
MMEU	0.00182	0.00235	0.00082	0.00415	0.00187
MMEO	0.00243	0.00216	0.00521	0.00188	0.00234
1 250 day window					
ME	0.00000	0.00000	-0.00001	0.00338	0.00000
MMEU	0.00172	0.00235	0.00303	0.00735	0.00177
MMEO	0.00273	0.00216	0.00202	0.00320	0.00267

Table C7: Interest Rate Variances

Variance – daily

	Hist	Exp	Dexp	BEKK	GARCHCC
125 day window					
ME	0.00000	0.00000	0.00025	0.00006	-0.00002
MMEU	0.00274	0.00276	0.00353	0.00278	0.00285
MMEO	0.00651	0.00595	0.01202	0.00652	0.00599
250 day window					
ME	0.00001	0.00000	0.00026	0.00008	-0.00001
MMEU	0.00272	0.00276	0.00311	0.00305	0.00276
MMEO	0.00693	0.00595	0.01145	0.00694	0.00636
500 day window					
ME	0.00003	0.00000	0.00026	0.00011	0.00002
MMEU	0.00262	0.00276	0.00307	0.00313	0.00271
MMEO	0.00774	0.00595	0.01155	0.00783	0.00715
750 day window					
ME	0.00005	0.00000	0.00010	0.00016	0.00004
MMEU	0.00262	0.00276	0.00268	0.00308	0.00275
MMEO	0.00833	0.00595	0.01105	0.00851	0.00804
1 250 day window					
ME	0.00014	0.00000	0.00018	0.00051	0.00014
MMEU	0.00262	0.00276	0.00282	0.00351	0.00253
MMEO	0.00951	0.00595	0.01054	0.01049	0.00927

Table C8: Interest Rate Variances

Variance – quarter-average

	Hist	Exp	Dexp	BEKK	GARCHCC
125 day window					
ME	0.00001	0.00000	0.00026	0.00020	0.00066
MMEU	0.00342	0.00382	0.00436	0.00354	0.00390
MMEO	0.00378	0.00351	0.01016	0.00720	0.01240
250 day window					
ME	0.00001	0.00000	0.00026	0.00024	0.00015
MMEU	0.00325	0.00382	0.00396	0.00381	0.00272
MMEO	0.00422	0.00351	0.00954	0.00773	0.00546
500 day window					
ME	0.00003	0.00000	0.00030	0.00032	0.00022
MMEU	0.00325	0.00382	0.00382	0.00402	0.00319
MMEO	0.00516	0.00351	0.00975	0.00894	0.00715
750 day window					
ME	0.00005	0.00000	0.00015	0.00038	0.00020
MMEU	0.00305	0.00382	0.00296	0.00391	0.00301
MMEO	0.00562	0.00351	0.00894	0.00984	0.00815
1 250 day window					
ME	0.00014	0.00000	0.00019	0.00213	0.00024
MMEU	0.00305	0.00382	0.00324	0.00594	0.00271
MMEO	0.00704	0.00351	0.00832	0.01391	0.00910

References

Alexander, C.O. and C.T. Leigh (1997), 'On the Covariance Matrices used in Value at Risk Models', *The Journal of Derivatives*, 4(3), pp. 50–62.

Berndt, K., B. Hall, R. Hall and J. Hausman (1974), 'Estimation and Inference in Nonlinear Structual Model', *Annals of Economic and Social Measurement*, 3(4), pp. 653–665.

Bollerslev, T. (1986), 'Generalized Autoregressive Conditional Heteroskedasticity', *Journal of Econometrics*, 31(3), pp. 307–327.

Bollerslev, T. (1990), 'Modelling the Coherence in Short-run Nominal Exchange Rates: A Multivariate Generalised ARCH Model', *The Review of Economics and Statistics*, 72(3), 498–505.

Boudoukh, J., M. Richardson and R.F. Whitelaw (1997), 'Investigation of a Class of Volatility Estimators', *The Journal of Derivatives*, 4(3), pp. 63–71.

Brailsford, T. And R. Faff (1996), 'An Evolution of Volatility Forecasting Techniques', *Journal of Banking and Finance*, 20(3), pp. 419–438.

Campa, J.M. and P.H.K. Chang (1997), 'The Forecasting Ability of Correlations Implied in Foreign Exchange Options', NBER Working Paper No. 5974.

Cassidy, C. and M. Gizycki (1997), 'Measuring Traded Market Risk: Value-at-Risk and Backtesting Techniques', Reserve Bank of Australia Research Discussion Paper No. 9708.

Engle, R. (1982), 'Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation', *Econometrica*, 50(4), pp. 987–1007.

Engle, R. and K. Kroner (1995), 'Multivariate Simultaneous Generalized ARCH', *Econometric Theory*, 11(1), pp. 122–150.

Figlewski, S. (1994), 'Forecasting Volatility Using Historical Data', New York University Salomon Brothers Working Paper Series. No. S-94-13.

Gizycki, M. and N. Hereford (1998), 'Assessing the Dispersion in Banks' Estimates of Market Risk: The Results of a Value-at-Risk Survey', Australian Prudential Regulation Authority Discussion Paper No. 1.

Jenrich, R. (1970), 'An Asymptotic χ^2 Test for the Equality of Two Correlation Matrices', *Journal of the American Statistical Association*, 65(330), pp. 904–912.

JP Morgan and Reuters (1996), 'RiskMetrics™ – Technical Document', Fourth Edition, New York.

Kaplanis, E.C. (1988), 'Stability and Forecasting of the Comovement Measures of International Stock Market Returns', *Journal of International Money and Finance*, 7(1), pp. 63–75.

Longin, F. and B. Solnik (1995), 'Is the Correlation in International Equity Returns Constant: 1960–1990?', *Journal of International Money and Finance*, 14(1), pp. 3–26.

Sheedy, E. (1997), 'Correlation in International Equity and Currency Markets: A Risk Adjusted Perspective', Macquarie University Centre for Studies in Money Banking and Finance Paper No. 17.

West, K.D. and D. Cho (1994), 'The Predictive Ability of Several Models of Exchange Rate Volatility', NBER Working Paper No. 152.