

# **LONG-RUN VOLATILITY FORECASTING: WHAT IS BEST FOR FAIR VALUATION OF ESOs?**

May 2009

## **Abstract**

Employee stock options (ESOs) are prominently used as part of the remuneration offered by many companies. For the sake of financial transparency, it is important that these ESOs be fairly valued. A very important, but certainly a vague or “rubbery” input for ESO valuation is the expected future volatility of the underlying stock price. This vagueness stems in large part from the fact that ESOs are generally long-lived. We address this problem by testing the predictive ability of several volatility estimation models, mostly from the ARCH family, to see which best forecasts future stock price volatility over a long term of five years. Our results indicate that the EGARCH model generally produces the best five year volatility forecasts, in terms of cross-sectional mean relative (forecast versus realised) volatility errors.

---

## 1. Introduction

Many companies in Australia and overseas grant employee stock options (ESOs) to their employees as part of their remuneration package. International Financial Reporting Standard 2, *Share-based Payment* (IFRS 2), issued by the International Accounting Standards Board (IASB), requires companies to recognise the “fair value” of ESOs granted to employees as an expense over the vesting period. Australia has adopted this standard, along with many other countries.<sup>1</sup> In the United States, Financial Accounting Standard 123 (revised) (FAS 123(R)) requires companies to meet similar fair value reporting standards for ESOs.

Determining the fair value of ESOs is problematic. The expected future volatility of the underlying stock price is an essential and very consequential input to pricing methods such as the Black-Scholes model and Cox-Ross-Rubinstein binomial model. Fair valuation at time of granting of ESOs requires future volatility to be estimated, but different estimates can produce substantially different values for otherwise identical ESOs. There is no universally accepted “best practice” method of estimating volatility, and IFRS 2 offers no recommendation or guidance.<sup>2</sup> Consequently ESO expenses cannot be easily compared across firms that employ different volatility estimation methods.

We address this problem by testing the predictive ability of several historical volatility estimation methods to see which “best” forecasts future long term volatility. The approach is different from most previous volatility forecasting research, which has generally focused purely on short or medium term predictive ability. We are concerned with predictive performance over a long horizon of five years, being a suitable life-span for ESOs. This paper is thus of practical importance to: accountants and human resource managers and their advisors who have to fairly value ESOs; employees who

---

<sup>1</sup> See <http://www.iasplus.com/country/useias.htm> for a list of countries that have adopted accounting standards issued by the IASB.

<sup>2</sup> Compounding the problem is that the long life of ESOs (typically five years in Australia and ten years in the United States) generally renders irrelevant any implied volatility reference taken from market valuations of usually far shorter-lived exchange traded options.

have to personally assess the fair value of their ESO packages; and investors who have to judge the corporate valuation impacts of ESOs.

The historical volatility estimation methods we test for forecast accuracy are primarily ARCH family models, in particular: generalised autoregressive conditional heteroskedasticity (GARCH) of Bollerslev (1986); exponential GARCH (EGARCH) of Nelson (1991); integrated GARCH (IGARCH) of Engle and Bollerslev (1986); and, fractionally integrated GARCH (FIGARCH) of Baillie, Bollerslev and Mikkelsen (1996). Two other methods are also tested: exponentially weighted moving average (EWMA) variance, and simple historical variance.

The long-run predictive performance of the volatility estimators is assessed using stock price data from the Australian Securities Exchange (ASX). For each company in our dataset and each volatility estimator, weekly stock returns for a specified estimation time-window are used to estimate future volatility for a subsequent adjacent forecast time-window. The forecast volatility can then be assessed against the realised volatility over the forecast window. We consider estimation windows of two and five years, and a forecast window of five years. The estimation plus forecast windows are repeatedly rolled forward by one week at a time resulting in a rolling time series of overlapping estimation plus forecast windows and their associated volatility forecasts and realisations. Rolling time series of forecast versus realised volatility error metrics are calculated and averaged in cross-section to indicate which volatility estimator provides the “best” volatility forecasts. Our results indicate that the EGARCH model generally produces the best five year volatility forecasts.

The rest of the paper is organised as follows. Section 2 outlines IFRS 2 and provides an analysis of relevant literature. Section 3 describes the technical specifications of the volatility estimators. Section 4 describes the data and method for the empirical analysis. Section 5 provides the results, and Section 6 concludes the paper.

## 2. Accounting standards and background literature

In February 2004 the International Accounting Standards Board (IASB) issued International Financial Reporting Standard 2, *Share-based Payment* (IFRS 2). The effect of this standard is to require entities to recognise the fair value of employee stock options (ESOs) granted to employees as an expense over the vesting period. This requirement places a more onerous accounting burden on companies, who previously did not have to record the fair value of ESOs in their profit-and-loss statements. Australia adopted IFRS 2 in July 2004 through AASB 2, *Share-based Payment*,<sup>3</sup> which took effect for annual reporting periods beginning on or after 1 January 2005.

The United States retains the Financial Accounting Standards Board (FASB) as its primary standard-setter. The FASB's FAS 123, *Accounting for Stock-Based Compensation*, issued in 1995, previously allowed companies to report the "intrinsic value" of ESOs rather than the fair value.<sup>4</sup> However, in December 2004 the FASB issued FAS 123(R), which removed the option of reporting the intrinsic value of ESOs and required companies to recognise the fair value. FAS 123(R) was designed to harmonise US reporting requirements with those of the IASB, and the provisions of the standard are similar to those of IFRS 2.

IFRS 2 states that the fair value of an ESO is to be determined using an option pricing model (Para 17).<sup>5</sup> The standard does not specify which option pricing model is to be used, other than it "be consistent with generally accepted valuation methodologies for pricing financial instruments" (Para 17). The standard also requires the model to take into account six factors: the exercise price of the option, the life of the option, the cur-

---

<sup>3</sup> AASB 2 contains the same provisions as, and is equivalent to, IFRS 2. This paper will refer to IFRS 2 rather than AASB 2.

<sup>4</sup> The intrinsic value of a call option is simply the current stock price minus the exercise price, assuming that the stock price is greater than or equal to the exercise price. It is almost always lower than the fair value, because it does not include the time value of the option.

<sup>5</sup> This assumes that the price cannot be gauged from an active options market. In practice, most ESOs have long lives and vesting conditions that preclude inferring the price from publicly traded options.

rent stock price, the expected volatility of the stock, the expected dividends on the stock, and the risk-free interest rate (Para B6). “Volatility” is defined as the annualised standard deviation of the continuously compounded rate of return of the stock over the life of the option (Para B22). Though not explicit, these factors anticipate the use of a risk-neutral pricing method such as the Black-Scholes model or Cox-Ross-Rubinstein binomial model (the models seemingly most popular in practice).<sup>6</sup>

IFRS 2 gives some guidance on estimating future volatility, without providing a “standard” or explicit formula. Some factors that may be taken into account include the implied volatility from actively traded options, the historical volatility of the stock and the possibly mean-reverting behaviour of volatility (Para 25). The standard also allows some historical periods to be excluded from the historical volatility calculation if they were “extraordinarily volatile” periods (Para 25d). Ultimately it is left to reporting entities to decide over what intervals returns will be measured, what periods of historical data will be disregarded, and how much historical data will be used in estimation.

## **2.1. Conditional variance models**

Conditional variance models predict future volatility by modelling the conditional variance as a function of time, past volatility and past errors in estimation; they use accrued “memory” of previous returns or volatility to make predictions about future volatility. The dominant class of conditional variance models, in both theory and practice, is the ARCH (autoregressive conditional heteroskedasticity) family of models. The first ARCH model was proposed by Engle (1982), who modelled conditional stock return variance (volatility squared) as a function of previous stock return innovations (ARCH terms) and a constant level of volatility. The number of lagged return innovations and the relative weight that they were given determined the memory of the process.

---

<sup>6</sup> For example, see Deloitte (2004) *A Guide to IFRS 2*.

The generalised ARCH (GARCH) of Bollerslev (1986) is the best-known ARCH-class model. Volatility is modelled as a function of past volatility, some underlying “long-run” level of volatility and the latest return innovation. Compared to traditional ARCH models, GARCH models allow for persistence of shocks to the conditional variance where the shocks decay at a geometric rate. Exponential GARCH (EGARCH) was developed by Nelson (1991) to allow for asymmetric responses of the conditional variance to positive and negative return shocks; this was prompted by the observation that negative return shocks tend to increase conditional variance more than positive shocks (see Black (1976)). Integrated GARCH (IGARCH), proposed by Engle and Bollerslev (1986), is similar to GARCH except that shocks to the conditional variance persist indefinitely. The fractionally integrated GARCH (FIGARCH) of Baillie et al. (1996) incorporates fractional long memory into the ARCH type model as a “middle ground” between non-integrated ARCH models (e.g. GARCH and EGARCH) and integrated ARCH models (e.g. IGARCH); shocks to the conditional variance decay at a hyperbolic rate, which is slower than the geometric decay of GARCH but faster than the indefinite persistence of IGARCH.

## **2.2. Empirical evidence on the performance of ARCH models**

There is no shortage of literature about the performance of ARCH models in predicting the future volatility of stock returns. For example, Poon and Granger (2003) identified 39 major studies of the predictive ability of the simple GARCH model. Most studies focus on the short-run performance of these models, often for forecast horizons of only a few days or weeks. In contrast, there is little evidence on the long-run predictive ability of ARCH models, a matter relevant to the pricing of long-lived ESOs.

Some studies that do look at the long-run performance of ARCH volatility estimators are those of Baillie et al. (1996) and Bollerslev and Mikkelsen (1996, 1999). Baillie et al. found that their newly proposed FIGARCH model outperforms other ARCH models in specifying the volatility process of the US\$/Deutschmark exchange rate. Bollerslev and Mikkelsen (1996) found that fractionally integrated ARCH models outperform GARCH and EGARCH models for pricing options on the S&P 500 index. And

Bollerslev and Mikkelsen (1999) showed that fractionally integrated ARCH models outperform simpler ARCH models in pricing three year LEAPS options on the S&P 500 index.

With a view to the requirements of ESO pricing, there are three primary aspects that distinguish our study from the majority of previous long-run volatility forecasting studies. Firstly, whereas previous studies have generally considered long-run forecasting horizons of three years or less, our forecasting horizon of five years is a better match to the (initial) maturities of ESOs in general. In this regard it is reasonable to suppose that there may be a difference between the three year (or less) and five year forecast performance of the various volatility estimators, just as there is evidence of a difference between short-run and three year forecast performance. Secondly, previous long-run volatility forecasting studies have generally been undertaken with foreign exchange or share price index data, rather than individual stock return data relevant to ESO pricing. Thirdly, we use comparatively low frequency (weekly) stock return data to estimate and forecast volatility so as to assuage market illiquidity issues associated with higher frequency returns.

Our paper examines the long-run forecasting ability of several volatility estimation models, mostly from the ARCH family, using a wide-ranging dataset of individual stock returns. Specifically we consider Australian equities, which have been subject to relatively few conditional variance studies.<sup>7</sup>

### **3. Volatility models**

The following conditional variance models from the ARCH family are examined: GARCH(1,1); EGARCH(1,1); IGARCH(1,1); and FIGARCH(1,d,0). Furthermore an exponentially weighted moving average (EWMA) estimator and an estimator based purely on the simple historical variance are considered.

---

<sup>7</sup> A notable contribution to the Australian evidence is Brailsford and Faff (1996), who used the Statex-Actuaries Accumulation Index to test conditional variance models in the Australian market.

### 3.1. ARCH family volatility models

Specify the information set of the market at time  $t$  to be given by  $\mathfrak{S}_t$ . Define  $R_t$  to be the continuously compounding total return rate for a stock for the period  $t-1$  to  $t$ . Now suppose the stock return follows an AR(1) process:<sup>8</sup>

$$R_t = \mu_0 + \mu_1 R_{t-1} + \varepsilon_t , \quad (1)$$

where  $\mu_0$  and  $\mu_1$  are constants and  $\varepsilon_t$  is a random error term. The error term  $\varepsilon_t$  determines return innovations according to a stochastic process:

$$\varepsilon_t = \sqrt{h_t} \eta_t ,$$

where  $\eta_t$  is an independent and identically distributed (i.i.d.) white noise series with zero expectation and unit variance (i.e.  $E[\eta_t]=0$  and  $\text{VAR}[\eta_t]=1$ ), and  $h_t$  is a  $\mathfrak{S}_{t-1}$ -measurable process. By conditioning on the previous information set  $\mathfrak{S}_{t-1}$ , the conditional variance of the stock return can be inferred:

$$\text{VAR}[R_t | \mathfrak{S}_{t-1}] = \text{VAR}[\varepsilon_t | \mathfrak{S}_{t-1}] = h_t .$$

The generalised autoregressive conditional heteroskedasticity (GARCh) model of Bollerslev (1986) defines the conditional variance of the returns as a function of some constant level of variance, lagged return innovations (up to  $p$  lags), and lagged conditional variances (up to  $q$  lags). For the GARCh(1,1) model only one innovation lag ( $p=1$ ) and one conditional variance lag ( $q=1$ ) are used:

$$h_t = \omega + \alpha \varepsilon_{t-1}^2 + \beta h_{t-1} , \quad (2)$$

where  $\omega$  is the constant level of variance,  $\alpha$  determines the impact of the lagged innovation and  $\beta$  determines the persistence or “inertia” of the lagged conditional vari-

---

<sup>8</sup> An AR(1) process was chosen to describe the return process, as it provided a good “fit” to data in preliminary testing, and is in line with other research. The AR(1) specification does not materially reduce the generality of this section.

ance. To ensure the process is stationary (i.e. does not explode as time passes) it is sufficient to have  $\alpha + \beta < 1$ . Having  $\alpha$ ,  $\beta$  and  $\omega > 0$  is sufficient to ensure that the variance is strictly positive. Under this model, shocks to the conditional variance decay at a geometric rate.

The exponential GARCH (EGARCH) model was introduced by Nelson (1991) to accommodate asymmetric response of the conditional variance to positive and negative return innovations. Assuming that innovations are normally distributed,<sup>9</sup> the conditional variance of the EGARCH(1,1) model is given by:

$$\ln(h_t) = \omega + \theta \frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + \gamma \left[ \left| \frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} \right| - \sqrt{\frac{2}{\pi}} \right] + \beta \ln(h_{t-1}) . \quad (3)$$

Since the left hand side of the equation takes the logarithm of the conditional variance, the right hand side can become negative without the conditional variance becoming negative. The parameters  $\omega$ ,  $\theta$ ,  $\gamma$  and  $\beta$  are free to take positive or negative values. If  $\theta < 0$ , a negative shock will increase the conditional variance more than a positive shock of the same magnitude, and vice versa for  $\theta > 0$ . Having  $|\beta| < 1$  is sufficient to ensure the process is stationary (He, Terasvirta and Malmsten (2002)). Like the GARCH model, shocks to the conditional variance decay at a geometric rate.

The integrated GARCH (IGARCH) model of Engle and Bollerslev (1986) is a “tweaked” version of the GARCH model where the condition  $\alpha + \beta = 1$  is imposed. This is to reflect that, in practice, the observed  $\alpha$  and  $\beta$  GARCH parameters (for various lags) often sum approximately to one (see Chou (1988) and Mikosch and Stărică (2004)). Thus, for the IGARCH(1,1) model, the conditional variance is specified by equation (2) with the constraint that  $\alpha + \beta = 1$ . Like the GARCH model, having  $\alpha$ ,  $\beta$  and  $\omega > 0$  is sufficient to ensure the process is always non-negative. Unlike the GARCH(1,1) model, under the IGARCH(1,1) model volatility shocks per-

---

<sup>9</sup> This assumption is discussed in Section 4.1. It does not substantially affect the formulation of the conditional variance.

sist indefinitely, and the conditional variance exhibits infinite dependence on initial conditions (Nelson (1990)).

Fractionally integrated GARCH (FIGARCH) reflects the fractional long memory documented in stock returns (Ding, Granger and Engle (1993)). Empirical research indicates that shocks to volatility decay more slowly than commensurate with GARCH or EGARCH specifications, but more quickly than commensurate with the infinite persistence of IGARCH. Following Bollerslev and Mikkelsen (1996), the FIGARCH(1,d,0) specification is adopted:

$$h_t = \omega - \beta \varepsilon_{t-1}^2 + \left[ d \varepsilon_{t-1}^2 - \frac{d(d-1)}{2!} \varepsilon_{t-2}^2 + \frac{d(d-1)(d-2)}{3!} \varepsilon_{t-3}^2 - \dots \right] + \beta h_{t-1} . \quad (4)$$

When  $0 < d < 0.5$ , the process has long memory and the stock return process follows fractional Brownian motion. To ensure that the process is stationary, it is sufficient to have  $\beta < d$  and  $-0.5 < d < 0.5$  (Chung (2001)). To ensure that the conditional variance is strictly positive, it is sufficient to have  $d$ ,  $\beta$  and  $\omega > 0$ .

### 3.2. Other volatility models

The exponentially weighted moving average (EWMA) estimator is often used in risk management to forecast volatility in value-at-risk calculations (Jorion (1997)).<sup>10</sup> In calculating the mean of the squares of the historic return variations from mean return, this estimator assigns maximum weight to the most recent observation, and exponentially decreasing weights to progressively older observations. Thus:

$$h_t = \left( \sum_{i=1}^K \lambda^{i-1} \right)^{-1} \sum_{i=1}^K \lambda^{i-1} (R_{t+1-i} - \bar{R})^2 , \quad 0 < \lambda < 1 ,$$

where  $K$  is the number of observations and  $\bar{R}$  is the mean return of the sample. The parameter  $\lambda$  determines the persistence of volatility; the higher is  $\lambda$ , the longer

---

<sup>10</sup> One of the better known EWMA estimators is JP Morgan's RiskMetrics ([www.jporganchase.com](http://www.jporganchase.com)).

shocks to the conditional variance take to die away. If  $K$  is large, we can conveniently approximate the conditional variance as:

$$h_t \approx \lambda h_{t-1} + (1-\lambda)(R_t - \bar{R})^2 . \quad (5)$$

To ensure stationarity it is sufficient to have  $0 < \lambda < 1$ .

Finally we also consider the simple historical variance model, also known as the “random walk” model. This method simply uses the past sample variance of the continuously compounded returns over a certain period to forecast future variance.

#### **4. Data and Method**

Time series of weekly (Wednesday-to-Wednesday) stock returns are obtained from the SIRCA daily database of Australian Securities Exchange (ASX) stock prices and used with our chosen estimators to forecast stock return variance, and to provide realised variance by which the forecasting performance of the estimators can be assessed. With weekly steps, we apply rolling estimation plus forecast windows of two plus five years and five plus five years over the period January 1980 to August 2007; thus we consider all companies in the database with a continuous weekly returns history of at least seven years and at least ten years during January 1980 to August 2007. For the two plus five year (estimation plus forecast) rolling window, the number of eligible companies varies from 21 to 645; and for the five plus five year rolling window, the number of eligible companies varies from 16 to 491. The number of eligible companies generally increases steadily and moderately as the rolling windows move forward through the January 1980 to August 2007 time-frame. However the database presents a very large jump upwards in eligible companies for the two plus five years and five plus five years commencing January 1991 (ending January 1998 and January 2001 respectively).

#### 4.1. Estimating the ARCH family volatility models

For each of our chosen ARCH family estimators, parameters are obtained using maximum likelihood estimation.

Given a time series of  $K$  weekly stock returns ( $R_{t-K+1}$  to  $R_t$ ) for either a two year or five year estimation window (i.e.  $K \approx 104$  or 261 weeks),  $K-1$  return innovations ( $\varepsilon_{t-K+2}$  to  $\varepsilon_t$ ) are computed from the return equation (1):  $\varepsilon_t = R_t - \mu_0 - \mu_1 R_{t-1}$ . Weiss (1986) and Bollerslev and Wooldridge (1992) showed that, asymptotically, maximum likelihood estimation is insensitive to the underlying distribution of innovations. Thus we follow Baillie et al. (1996) and Bollerslev and Mikkelsen (1996, 1999) and assume that innovations are normally distributed such that  $\varepsilon_t \sim N(0, h_t)$ .

Equations (2) and (3) specify the conditional variance time series for our GARCH(1,1), EGARCH(1,1) and IGARCH(1,1) estimators. To initiate each conditional variance time series,  $h_{t-K+2}$  is set equal to the sample return variance over the estimation window. Thus  $K-2$  conditional variances are computed ( $h_{t-K+3}$  to  $h_t$ ).

For the GARCH(1,1) model, equation (2) is reparameterised to ensure the estimated conditional variance process is stationary and strictly positive:

$$\begin{aligned} h_t &= \omega + \alpha \varepsilon_{t-1}^2 + \beta h_{t-1} \\ &= \omega + \alpha (R_{t-1} - \mu_0 - \mu_1 R_{t-2})^2 + \beta h_{t-1} \\ &= \exp(\ln_{\omega}) \\ &\quad + [1 - \text{logistic}(\text{invlogistic}_{\beta})] \text{logistic}(\text{invlogistic}_{\alpha} \ln_{\omega} \text{take}_{\beta}) (R_{t-1} - \mu_0 - \mu_1 R_{t-2})^2 \\ &\quad + \text{logistic}(\text{invlogistic}_{\beta}) h_{t-1}, \end{aligned}$$

where  $\ln_{\omega}$ ,  $\text{invlogistic}_{\alpha} \ln_{\omega} \text{take}_{\beta}$  and  $\text{invlogistic}_{\beta}$  are the new parameters, which together with the exponential and logistic functions put into effect the constraints  $\alpha$ ,  $\beta$  and  $\omega > 0$ , and  $\alpha + \beta < 1$ .

For the EGARCH(1,1) model, to ensure the estimated conditional variance process is stationary, equation (3) is reparameterised to put into effect the constraint  $|\beta| < 1$ :

$$\ln(h_t) = \omega + \theta \frac{(R_{t-1} - \mu_0 - \mu_1 R_{t-2})}{\sqrt{h_{t-1}}} + \gamma \left[ \left| \frac{R_{t-1} - \mu_0 - \mu_1 R_{t-2}}{\sqrt{h_{t-1}}} \right| - \sqrt{\frac{2}{\pi}} \right] + [2 \text{logistic}(\text{invlogistic\_}\beta \text{ plus } 1 \text{ on } 2) - 1] \ln(h_{t-1}) .$$

For the IGARCH(1,1) model, equation (2) is reparameterised to put into effect the constraints  $\alpha$ ,  $\beta$  and  $\omega > 0$ , and  $\alpha + \beta = 1$ :

$$h_t = \exp(\ln\_ \omega) + [1 - \text{logistic}(\text{invlogistic\_}\beta)] (R_{t-1} - \mu_0 - \mu_1 R_{t-2})^2 + \text{logistic}(\text{invlogistic\_}\beta) h_{t-1} .$$

For the FIGARCH(1,d,0) model, equation (4) is reparameterised to put into effect the constraints  $d$ ,  $\beta$  and  $\omega > 0$  and  $\beta < d < 0.5$ ; and the infinite fractional difference series applied to the innovation lags is truncated at 20 lags:

$$\begin{aligned} h_t \approx & \omega - \beta (R_{t-1} - \mu_0 - \mu_1 R_{t-2})^2 + \beta h_{t-1} + \left[ d (R_{t-1} - \mu_0 - \mu_1 R_{t-2})^2 \right. \\ & - \frac{d(d-1)}{2!} (R_{t-2} - \mu_0 - \mu_1 R_{t-3})^2 + \frac{d(d-1)(d-2)}{3!} (R_{t-3} - \mu_0 - \mu_1 R_{t-4})^2 - \dots \\ & \left. + \frac{d(d-1)(d-2)\dots(d-19)}{20!} (R_{t-20} - \mu_0 - \mu_1 R_{t-21})^2 \right] \\ = & \exp(\ln\_ \omega) \\ & - 0.5 \text{logistic}(\text{invlogistic\_}2d) \text{logistic}(\text{invlogistic\_}\beta \text{ on } d) [(R_{t-1} - \mu_0 - \mu_1 R_{t-2})^2 - h_{t-1}] \\ & + \left[ 0.5 \text{logistic}(\text{invlogistic\_}2d) (R_{t-1} - \mu_0 - \mu_1 R_{t-2})^2 - \dots \right. \\ & \left. + \frac{\prod_{i=0}^{19} [0.5 \text{logistic}(\text{invlogistic\_}2d) - i]}{20!} (R_{t-20} - \mu_0 - \mu_1 R_{t-21})^2 \right] . \end{aligned}$$

To initiate the FIGARCH(1,d,0) conditional variance time series,  $h_{t-K+2}$  to  $h_{t-K+21}$  are set equal to the sample return variance over the estimation window. Thus  $K - 21$  conditional variances are computed ( $h_{t-K+22}$  to  $h_t$ ).

To initiate the GARCH(1,1), EGARCH(1,1) and FIGARCH(1,d,0) maximum likelihood estimations, the starting value of  $\mu_0$  is set to the sample mean return over the estimation window, and  $\mu_1$  is set to zero. The starting values of the  $\ln\_ \omega$  parameter for GARCH(1,1) and FIGARCH(1,d,0), and the  $\omega$  parameter for EGARCH(1,1), are

set to the natural log of the sample return variance over the estimation window. For GARCH(1,1), the starting values of  $invlogistic\_beta$  and  $invlogistic\_alpha1takebeta$  are both set to -100 (effectively setting  $\alpha$  and  $\beta=0+$ ). For EGARCH(1,1), the starting values of  $\theta$ ,  $\gamma$  and  $invlogistic\_betaplus1on2$  are all set to zero (effectively setting  $\beta=0$ ). For FIGARCH(1,d,0), the starting values of  $invlogistic\_2d$  and  $invlogistic\_bond$  are respectively set to zero and -100 (effectively setting  $d=0.25$  and  $\beta=0+$ ).

The IGARCH(1,1) maximum likelihood estimation is given a “helping hand” by setting the starting values of  $\mu_0$ ,  $\mu_1$ ,  $ln\_omega$  and  $invlogistic\_beta$  equal to the estimates obtained from the GARCH(1,1) maximum likelihood estimation.<sup>11</sup>

## 4.2. Estimating the EWMA volatility model

The EWMA parameter is also obtained via maximum likelihood estimation.

From the time series of  $K$  weekly stock returns ( $R_{t-K+1}$  to  $R_t$ ),  $K$  return innovations ( $\varepsilon_{t-K+1}$  to  $\varepsilon_t$ ) are computed from the random walk return equation:  $\varepsilon_t = R_t - \mu_0$ ,  $\varepsilon_t \sim N(0, h_t)$ .

Equation (5) specifies the conditional variance time series. To initiate the conditional variance time series,  $h_{t-K+1}$  is set equal to the sample return variance over the estimation window. Thus  $K-1$  conditional variances are computed ( $h_{t-K+2}$  to  $h_t$ ). Equation (5) is reparameterised to ensure stationarity:

$$\begin{aligned} h_t &\approx \lambda h_{t-1} + (1-\lambda)(R_t - \bar{R})^2 \\ &= \lambda h_{t-1} + (1-\lambda)\varepsilon_t^2 \\ &= \text{logistic}(invlogistic\_lambda)h_{t-1} + [1 - \text{logistic}(invlogistic\_lambda)]\varepsilon_t^2 . \end{aligned}$$

---

<sup>11</sup> Regardless of parameter starting values, the IGARCH(1,1) maximum likelihood estimation has a very high rate of failure.

To initiate the EWMA maximum likelihood estimation, the starting value of *invlogistic* $_{\lambda}$  is set to 2.7515, which effectively sets  $\lambda = 0.94$  (being the value used in the RiskMetrics database, created by JP Morgan (1995) for value-at-risk and volatility consulting).

### 4.3. Forecasting conditional variance

Having obtained the maximum likelihood parameters of our various volatility models for a given estimation window, for each model a time series of  $N$  weekly conditional variance forecasts ( $\hat{h}_{t+1}$  to  $\hat{h}_{t+N}$ ) is projected for the adjacent five year forecast window (i.e.  $N \approx 260$  weeks).

Having reversed any reparameterisation of parameter estimates, and knowing  $\varepsilon_t$  and  $h_t$ , the GARCH(1,1) and IGARCH(1,1) conditional variance forecasts are given by:

$$\hat{h}_{t+1} = \hat{\omega} + \hat{\alpha}\varepsilon_t^2 + \hat{\beta}h_t ,$$

$$\hat{h}_{t+s} = \hat{\omega} + (\hat{\alpha} + \hat{\beta})\hat{h}_{t+s-1} , \quad s \geq 2 ,$$

noting that  $\hat{\alpha} + \hat{\beta} = 1$  for the IGARCH(1,1) model.

For the EGARCH(1,1) model the conditional variance forecasts are given by:

$$\ln(\hat{h}_{t+1}) = \hat{\omega} + \hat{\theta} \frac{\varepsilon_t}{\sqrt{h_t}} + \hat{\gamma} \left[ \left| \frac{\varepsilon_t}{\sqrt{h_t}} \right| - \sqrt{\frac{2}{\pi}} \right] + \hat{\beta} \ln(h_t) ,$$

$$\ln(\hat{h}_{t+s}) = \hat{\omega} + \hat{\beta} \ln(\hat{h}_{t+s-1}) , \quad s \geq 2 .$$

For the FIGARCH(1,d,0) model, 20 lagged innovations are used to make the  $t+1$  conditional variance forecast:

$$\hat{h}_{t+1} = \hat{\omega} - \hat{\beta}\varepsilon_t^2 + \hat{\beta}h_t$$

$$+ \left[ \hat{d}\varepsilon_t^2 + \frac{\hat{d}(1-\hat{d})}{2!}\varepsilon_{t-1}^2 + \frac{\hat{d}(1-\hat{d})(2-\hat{d})}{3!}\varepsilon_{t-2}^2 + \dots + \frac{\hat{d}(1-\hat{d})(2-\hat{d})\dots(19-\hat{d})}{20!}\varepsilon_{t-19}^2 \right] .$$

As the FIGARCH(1,d,0) conditional variance forecast is stepped forward, future squared innovations ( $\varepsilon_{t+1}^2$ ,  $\varepsilon_{t+2}^2$ , etc.) are represented by their estimated expectations (i.e.  $\hat{h}_{t+1}$ ,  $\hat{h}_{t+2}$ , etc.):

$$\hat{h}_{t+2} = \hat{\omega} + \left[ \hat{d}\hat{h}_{t+1} + \frac{\hat{d}(1-\hat{d})}{2!} \varepsilon_t^2 + \dots + \frac{\hat{d}(1-\hat{d})(2-\hat{d})\dots(19-\hat{d})}{20!} \varepsilon_{t-18}^2 \right], \text{ etc. ,}$$

$$\hat{h}_{t+s} = \hat{\omega} + \left[ \hat{d}\hat{h}_{t+s-1} + \frac{\hat{d}(1-\hat{d})}{2!} \hat{h}_{t+s-2} + \dots + \frac{\hat{d}(1-\hat{d})(2-\hat{d})\dots(19-\hat{d})}{20!} \hat{h}_{t+s-20} \right], \quad s \geq 21 .$$

For the EWMA model, the conditional variance forecasts all equal the current variance:  $\hat{h}_{t+s} = h_t$ ,  $s \geq 1$ . For the simple historical variance estimator, the conditional variance forecasts all equal the sample weekly return variance over the estimation window.

For each volatility model, the  $N$  weekly conditional variance forecasts over the forecast window are converted to an annualised “integrated” volatility:

$$\hat{S} = \sqrt{52.1 \frac{1}{N} \sum_{i=1}^N \hat{h}_{t+i}} .$$

This integrated volatility represents our volatility forecast. It can be considered analogous to the input for the Black-Scholes formula (Andersen, Bollerslev and Lange (1999)).<sup>12</sup>

#### 4.4. Assessing forecasting performance

For each stock return time series (i.e. each individual cross-sectional company), and each estimation plus forecast window, we seek to obtain six different volatility fore-

---

<sup>12</sup> Merton (1973) showed that integrated volatility,  $\sqrt{\int_t^T \sigma_u^2 du / (T-t)}$ , can be used in the Black-Scholes formula, where  $\sigma_t$  is a deterministic spot volatility function.

casts. However, for the ARCH family and EWMA models, if maximum likelihood estimation fails, or if the volatility forecast is more than twice the simple volatility forecast (also considered failure), the simple volatility forecast is used as a substitute. Therefore the volatility forecasts are termed: GARCH/simple, EGARCH/simple, IGARCH/simple, FIGARCH/simple, EWMA/simple, and simple. These six forecasts are assessed against the realised simple volatility over the forecast window,  $\sigma$  (i.e. the annualised sample standard deviation of weekly stock returns over the forecast window).

Following Bollerslev and Mikkelsen (1999), three cross-sectional volatility forecast error metrics are computed: mean relative error (MRE), mean relative absolute error (MRAE), and mean relative square error (MRSE):

$$\text{MRE} = \frac{1}{M} \sum_{i=1}^M \left( \frac{\hat{S}_i - \sigma_i}{\sigma_i} \right), \quad \text{MRAE} = \frac{1}{M} \sum_{i=1}^M \left( \frac{|\hat{S}_i - \sigma_i|}{\sigma_i} \right), \quad \text{MRSE} = \frac{1}{M} \sum_{i=1}^M \left( \frac{(\hat{S}_i - \sigma_i)^2}{\sigma_i^2} \right),$$

where  $M$  is the number of companies in our database with a weekly stock return time series for the full period of the given estimation plus forecast window.

A fourth metric, the substitution rate, indicates the proportion of  $M$  cross-sectional volatility forecasts that have the simple forecast substituted for the failed ARCH family or EWMA forecast.

By rolling the two plus five year and five plus five year estimation plus forecast windows forward, a time series of volatility forecast errors for each volatility model and each error metric is obtained.

## 5. Results

To provide a pattern of reference for the volatility forecasting error results, Figure 1 shows the time series of five years backward sample volatility and five years forward sample volatility for the Australian All Ordinaries Accumulation Index, and the relative error that results from using historic volatility to forecast future volatility. Figure 1 starkly illustrates how a short and sharp volatility shock, the October 1987 stock

market crash, leads to a distinct step up in the volatility forecast error. After October 1987, the five years forward volatility realisation is much lower than the five years backward volatility estimation, leading to a large positive relative forecast error; this situation remains until the five years backward estimation window rolls past October 1987, at which point the relative forecast error steps back down. Another short and sharp volatility shock (albeit not nearly as sharp as October 1987) is prominently observable as a step up in the volatility forecast error in late 1997, attributable to the Asian financial crisis.

The volatility forecast error time series in Figure 1 finishes with the five years forward volatility realisation extending into the current period of market turmoil (stemming from the US sub-prime mortgage crisis and the ensuing global financial crisis). Consequently, in late 2003, the five years backward volatility underestimates future volatility, leading to a negative relative forecast error.

As Figure 1 demonstrates, the variable severity and intervals of volatility shocks means that a five year look-back at volatility can be short-sighted in terms of anticipation of future volatility shocks. Nevertheless, when concerning ourselves with individual companies, it is relatively rare to have the luxury of decades of stock price trading history. Furthermore, even a long-lived company is likely to vary its financing and operating business model over time, implying its stock price volatility history may entail regime shifts that could be “lost” within an excessively long look-back period for volatility estimation. With such regard we apply two year and five year look-back volatility estimation windows to our cross-sectional samples of individual companies. This is with the understanding that the time series of mean volatility forecast errors will demonstrate significant variability; because rolling forward of the time series shifts the “framing” of disparate historic volatility shocks from the forecast window to the estimation window, leading to a cycle of realised volatility being underestimated by a forecast from look-back estimation, to realised volatility being overestimated.

Figures 2 and 3 display the time series of cross-sectional mean volatility forecast errors for the various volatility models for two plus five year and five plus five year estimation plus forecast rolling windows. Recall that, for the ARCH family and EWMA models, for individual cross-sectional companies, if maximum likelihood estimation

fails, or if the volatility forecast is more than twice the simple volatility forecast (also considered failure), the simple volatility forecast is used as a substitute. The bottom row, panels (Biv) to (Fiv) of both Figures 2 and 3 show the time series substitution rates for the ARCH family and EWMA models, thus indicating the prevalence of dependence on the simple volatility forecast as a substitute. Panel (Aiv) shows the time series cross-sectional sample size of eligible companies; here it can be seen that the database presents a large jump upwards in cross-sectional sample size for the two plus five years and five plus five years ending January 1998 and January 2001 respectively. Looking across each of the first three rows of Figures 2 and 3 allows comparison of the volatility forecast performance of the different volatility models in terms of the three error metrics (MRE, MRAE and MRSE).

Comparing (the first three rows of) the first and second columns of both Figures 2 and 3 shows little variation in the volatility forecasting performance of the simple and GARCH/simple models, although the GARCH/simple model uniformly entails moderately superior median error metrics. Moving then to the third column, the EGARCH/simple model generally entails further moderate improvement of median error metrics (median MRAE being the exception); the time series error metric patterns remain broadly similar to those of the simple and GARCH/simple models, but tend to be “fuzzy” due to the simple-for-EGARCH substitution rate hovering around 5% to 10% entailing considerable switching between simple and EGARCH volatility forecasts for individual cross-sectional companies.

Next the IGARCH/simple model in the fourth column has an extremely high simple-for-IGARCH substitution rate (generally around 70%); despite this high dependence on simple volatility forecasts substituting for IGARCH forecasts, the IGARCH/simple model is comfortably outperformed by the simple model in terms of median error metrics. The FIGARCH/simple model in the fifth column has a low simple-for-FIGARCH substitution rate, but nevertheless generally underperforms the simple model, but outperforms the IGARCH/simple model, again in terms of median error metrics.

Although the IGARCH/simple and FIGARCH/simple models tend to produce median error metrics inferior to the simple model, their time series error metric patterns are still somewhat similar to those of the simple model (and the GARCH/simple and

EGARCH/simple models). In the sixth column, the time series error metric patterns of the EWMA/simple model are much less similar. For the two year estimation windows (Figure 2) the EWMA/simple model produces, by far, the worst median error metrics. For the five year estimation windows (Figure 3) the EWMA/simple model produces the worst median MRAE.

### 5.1. Error metric time series parameters

Tables 1 and 2 summarise the time series parameters (minimum, maximum, mean, median and standard deviation) of the error metrics (MRE, MRAE and MRSE). The rows of Tables 1 and 2 are grouped into complete time series results and partial time series results. The large jump upwards in cross-sectional sample size distinctly visible in panel (Aiv) of both Figures 2 and 3 is used as the splitting point to commence the partial time series; the partial time series results therefore avoid any influence from this large jump in cross-sectional scope, and also happen to avoid any influence from the October 1987 stock market crash. The complete and partial time series results are furthermore given *with and without* substitution of simple volatility forecasts for failed ARCH family and EWMA forecasts. That is, we summarise the volatility forecasting performance of the GARCH/simple, EGARCH/simple, IGARCH/simple, FIGARCH/simple and EWMA/simple models (where the simple volatility forecast is substituted for a failed ARCH family or EWMA forecast) versus simple realised volatility; and we also summarise the “pure” volatility forecasting performance of the GARCH, EGARCH, IGARCH, FIGARCH and EWMA models (discarding failed forecasts) versus simple realised volatility. The (variable) failure rates of the ARCH family and EWMA forecasts mean that the cross-sectional fraction of eligible sample companies for which a valid volatility forecast can be obtained will often be less than unity when substitution of the simple volatility forecast is not used (this is especially an issue for the IGARCH model).

Therefore we have two “forecast error perspectives” for calculating the time series of mean cross-sectional volatility forecasting errors. For each cross-sectional sample company: (i) each volatility model’s forecast, with substitution of the simple forecast in event of estimation failure, is assessed against the realised simple volatility (e.g.

GARCH/simple versus simple); and (ii) each volatility model's forecast, with estimation failures discarded, is assessed against the realised simple volatility (e.g. GARCH versus simple).

In Tables 1 and 2, for each combination of, the complete and partial time series, forecast error perspective (i.e. with and without substitution of simple volatility forecasts for failed ARCH family and EWMA forecasts), five time series parameters (minimum, maximum, mean, median and standard deviation), and three error metrics (MRE, MRAE and MRSE): there is a column-group of six results indicating the relative performance of each volatility model (i.e. there are 60 column-groups of six,  $2*2*5*3=60$ , for Table 1 and for Table 2). Within each column-group of six, the "winner" result is shaded (this is the value closest to zero); if a single model provides the winner for all three error metrics, it is declared the "clear winner" for that particular forecast error perspective and time series parameter.

Tables 1 and 2 show the EGARCH/simple and EGARCH models to generally dominate as winners and clear winners. Across both Tables 1 and 2: out of 120 column-groups of six, EGARCH/simple and EGARCH together have 52 wins, followed by GARCH/simple and GARCH together with 30 wins; and out of 40 contests for clear winner, EGARCH/simple and EGARCH together are clear winners seven times, followed by IGARCH as clear winner twice. If we consider only the mean and median as the most important time series parameters: out of 48 column-groups of six, EGARCH/simple and EGARCH together have 30 wins, followed by GARCH/simple and GARCH together with 17 wins; and out of 16 contests for clear winner, EGARCH/simple and EGARCH together are clear winners five times, followed by GARCH as clear winner once.

However, if we focus only on the partial time series results, the EGARCH/simple and EGARCH performance falls behind GARCH/simple and GARCH. Considering only the partial time series results of Tables 1 and 2, out of 60 column-groups of six, GARCH/simple and GARCH together have 25 wins, followed by EGARCH/simple and EGARCH together with 17 wins. If we consider only the mean and median of the partial time series error metrics: out of 24 column-groups of six, GARCH/simple and GARCH together have 13 wins, followed by EGARCH/simple and EGARCH to-

gether with 11 wins; and out of eight contests for clear winner, GARCH and EGARCH achieve clear winner once each.

Seemingly in contrast to the error metric patterns shown in Figures 2 and 3, the EWMA/simple and EWMA models often have the lowest error metric standard deviations, suggesting strong volatility forecasting precision, even if mean and median forecasting performances are poor. Referring to Figures 2 and 3, it can be seen that the EWMA/simple error metric patterns entail a lot of high frequency volatility, but are less subject to big cyclical swings than is the case for the other volatility models.

## **5.2. What happened to FIGARCH?**

Our results differ from the small number of other long-run volatility forecasting studies, which find that fractionally integrated models provide superior volatility forecasts. We actually find the FIGARCH model to be a poor performer. Previous studies have had the luxury of estimating their models from very long time series of returns for very liquid assets. Our purpose is to forecast volatility for individual companies that will frequently have had relatively short and illiquid stock trading histories. To assuage market illiquidity concerns we use weekly stock return series. In conjunction with our two or five year estimation windows, our stock return series for volatility estimation have a mere 104 or 261 observations. This sparsity of data particularly hampers the FIGARCH estimation because 20 observations are used as return innovation lags for application of the fractional difference series. That is, the data-hungry FIGARCH model does not happily accommodate reasonable practical limitations of stock return series.

## **5.3. Two year or five year estimation window?**

Focussing only on the means and medians of the error metrics, and only on EGARCH/simple, EGARCH, GARCH/simple and GARCH models, the relative optimality of two year and five year estimation windows is not clear-cut. For the complete time series results, the 1987 stock market crash confounds a large subset of vola-

tility forecasts, and thus the two year rolling estimation window is superior; because a shorter rolling estimation window does not overlap the 1987 event as often. For the partial time series results, the five year rolling estimation window is generally superior; because a longer estimation window gives better model-parameter estimation (in the absence of extreme events).

## **6. Conclusion**

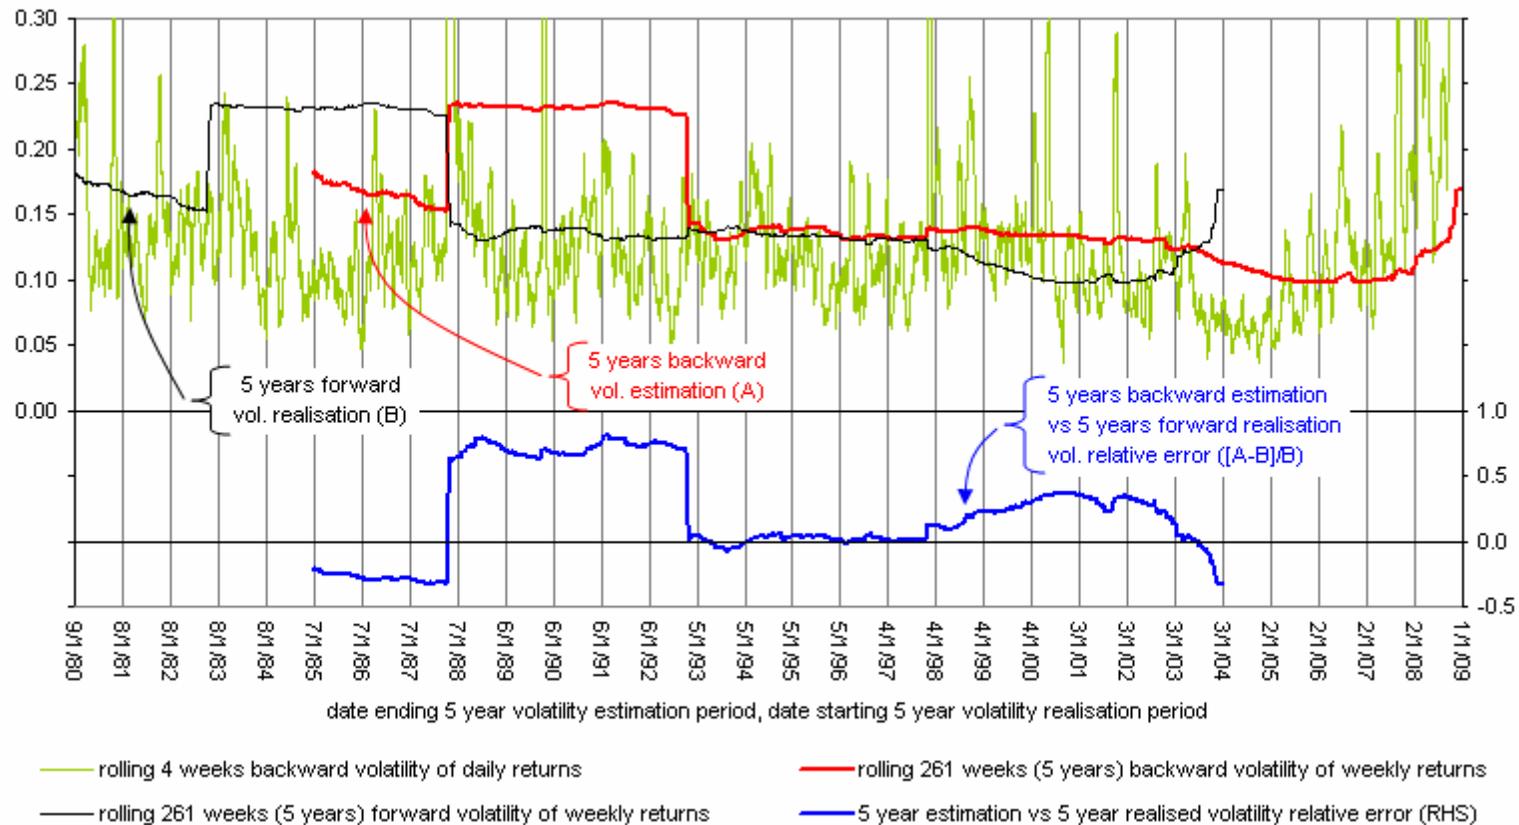
Employee stock options (ESOs) are prominently used as part of the remuneration offered by many companies. For the sake of financial transparency, it is important that these ESOs be fairly valued. A very important, but certainly a vague or “rubbery” input for ESO valuation is the expected future volatility of the underlying stock price. This vagueness stems in large part from the fact that ESOs are generally long-lived. Although financial market academics and participants have a long history of valuing and trading stock options, the vast majority of this experience has been concerned with short-lived options. Thus this paper is motivated to test the predictive ability of several volatility estimation models to see which best forecasts future volatility over a long term of five years, using as data a diverse time series set of Australian weekly stock returns.

From a range of six volatility models, mostly from the ARCH family, our results indicate that the EGARCH model generally produces the best five year volatility forecasts, in terms of cross-sectional mean relative (forecast versus realised) volatility errors.

Our results also demonstrate the confounding effect of extreme events on volatility forecasting. The temporal spacing between extreme events (e.g. the 1987 stock market crash and the current global financial crisis) means that, depending on the chosen estimation and forecast time windows, some extreme event will often be present in the estimation window but not in the realisation (i.e. forecast) window, or vice versa; this leads to time series cyclicalities of volatility forecast errors.

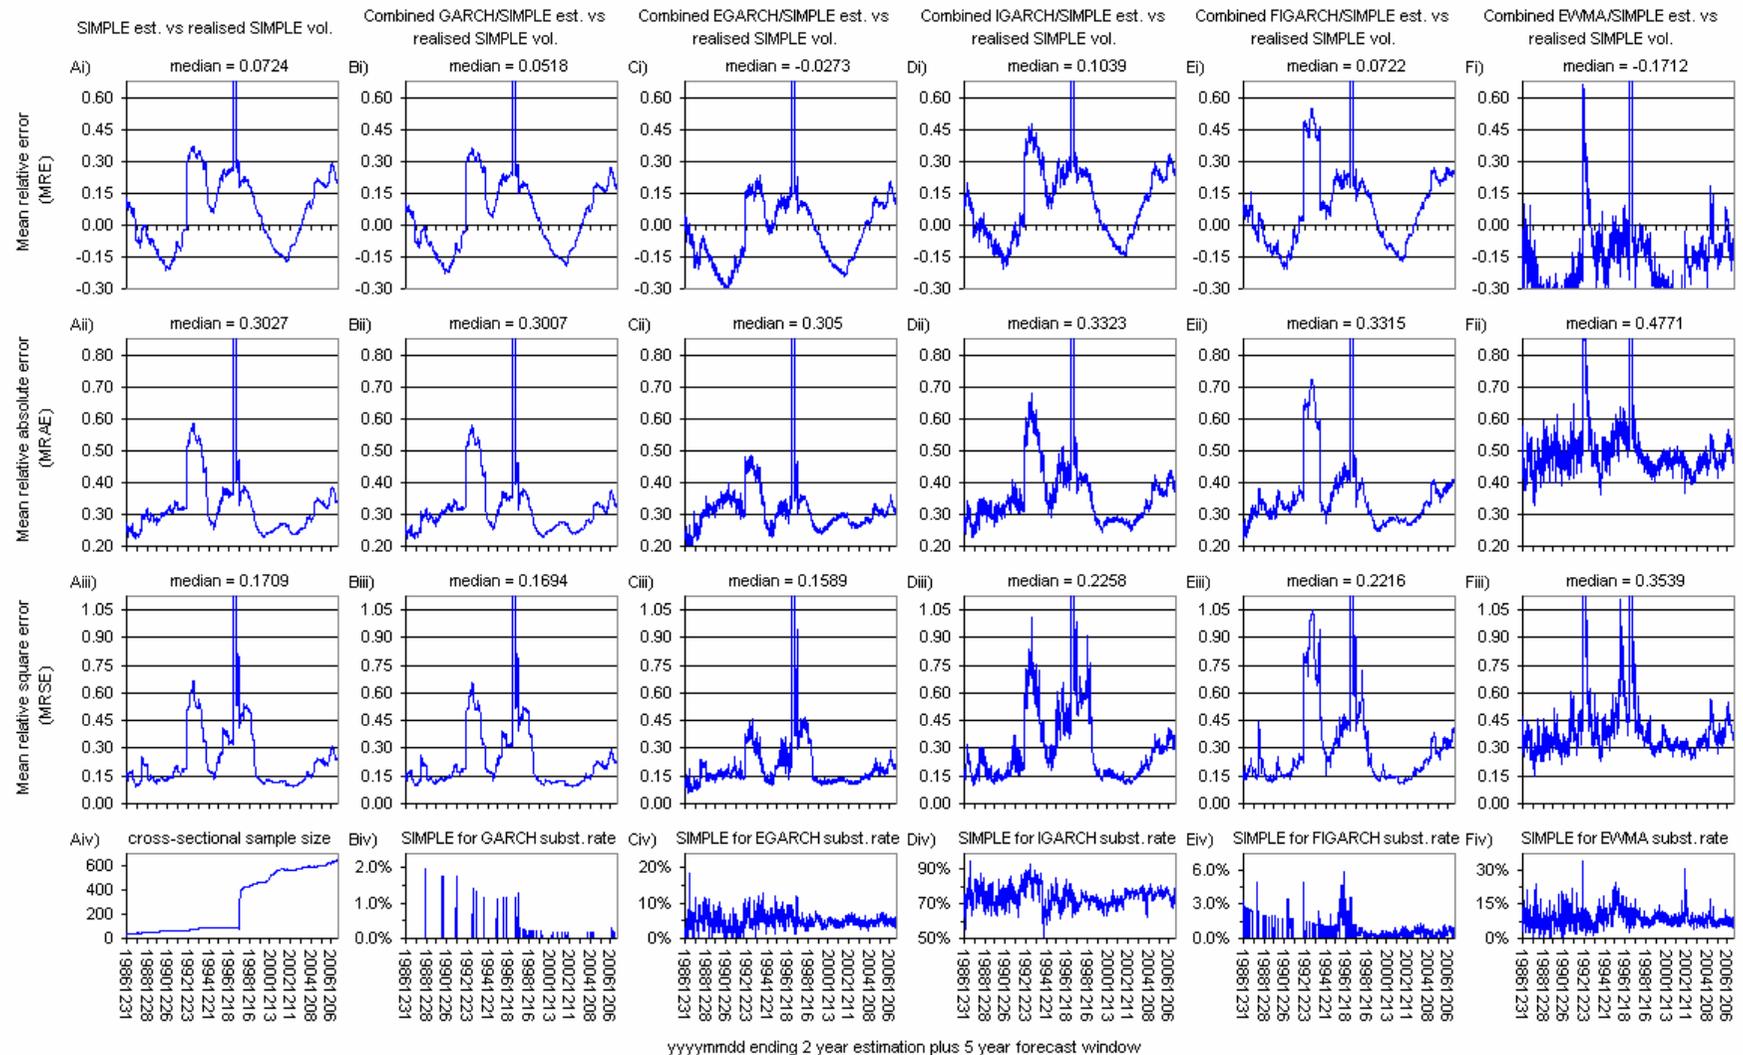
**Figure 1: AOAI backward and forward volatility**

Rolling five years backward and five years forward annualised sample weekly return volatility for Australia's All Ordinaries Accumulation Index; and the relative error associated with using the backward volatility to forecast the forward volatility. That is, for a chosen date, the chart shows historic volatility, future volatility, and the relative error that results from using historic volatility to predict future volatility.



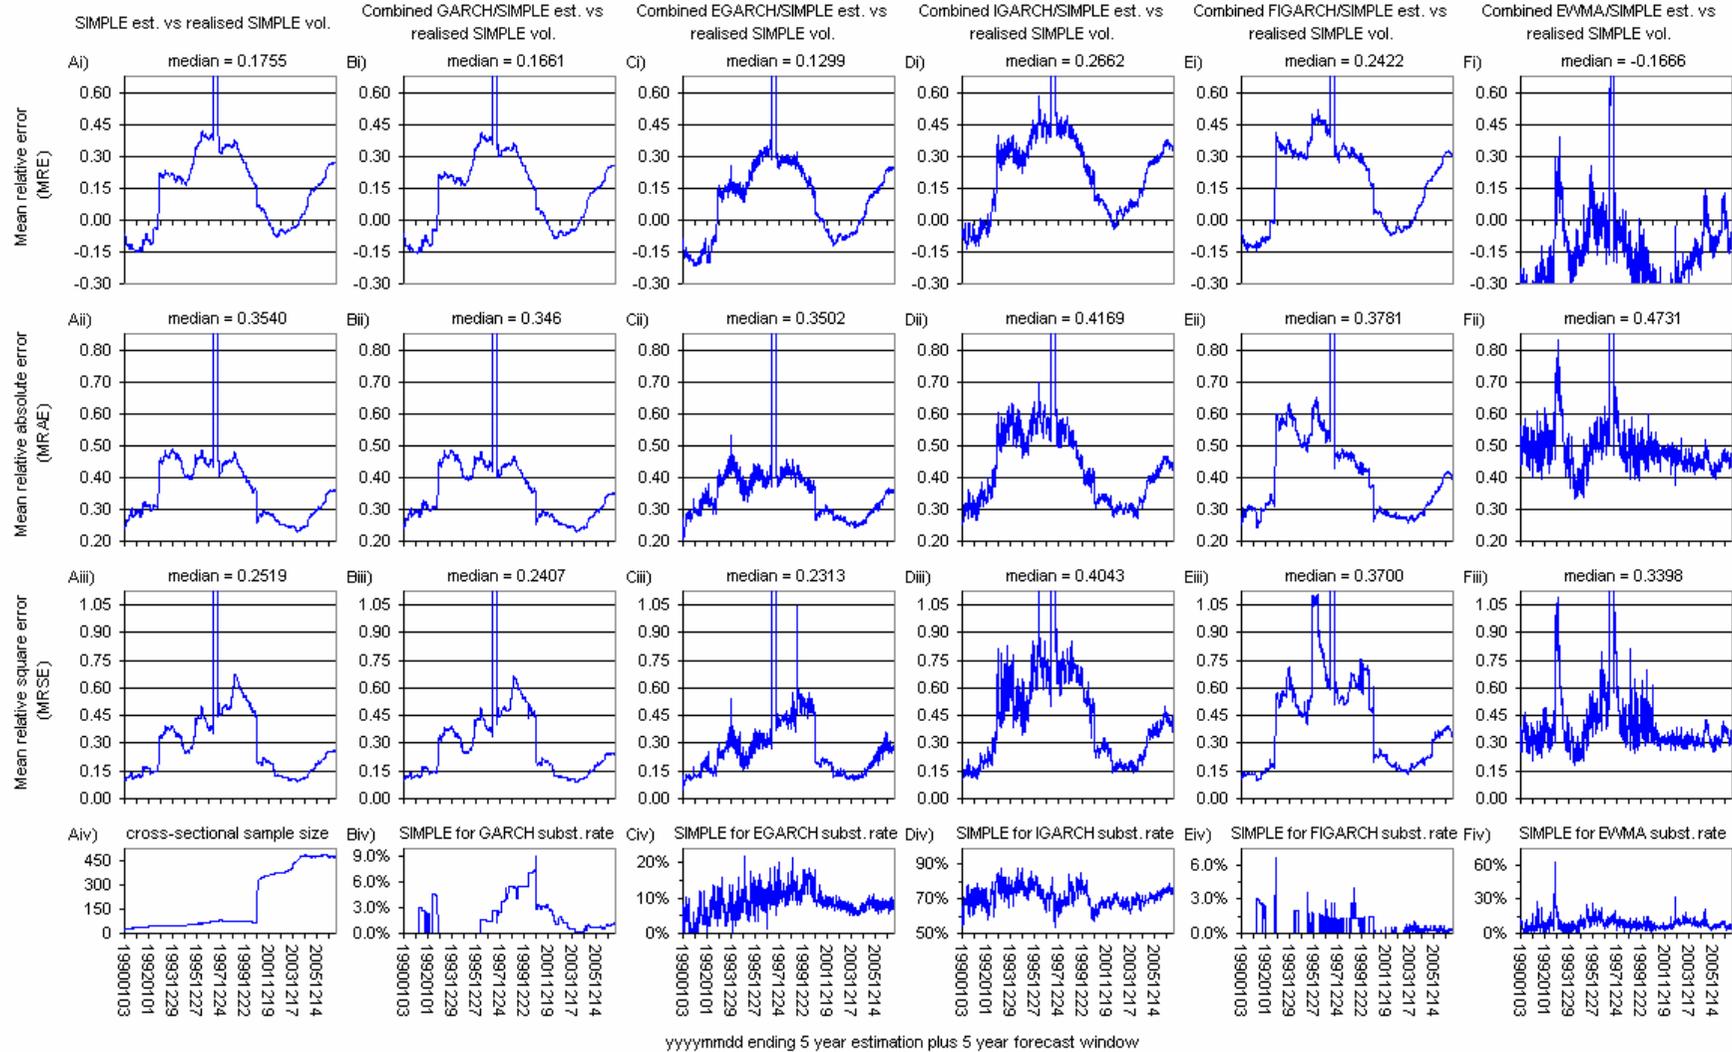
**Figure 2: Complete time series of volatility forecasting errors for a 2+5 year estimation plus forecast rolling window**

Complete time series of cross-sectional volatility forecasting errors for the simple, ARCH family and EWMA estimators, for a two plus five year estimation plus forecast rolling window; for the two plus five years ending 31 December 1986 to 1 August 2007 (1075 windows, weekly steps).



**Figure 3: Complete time series of volatility forecasting errors for a 5+5 year estimation plus forecast rolling window**

Complete time series of cross-sectional volatility forecasting errors for the simple, ARCH family and EWMA estimators, for a five plus five year estimation plus forecast rolling window; for the five plus five years ending 3 January 1990 to 1 August 2007 (918 windows, weekly steps).



**Table 1: Time series parameters of volatility forecasting errors for a 2+5 year estimation plus forecast rolling window**

Volatility forecast versus volatility realisation, estimation methods		Column-groups categorised by time series parameter, and cross-sectional volatility forecast versus realisation error metric *															Cross-sectional fraction of eligible companies with valid forecasted volatility				
		minimum			maximum			mean			median			standard deviation			min	max	mean	st. dev.	
		MRE	MRAE	MRSE	MRE	MRAE	MRSE	MRE	MRAE	MRSE	MRE	MRAE	MRSE	MRE	MRAE	MRSE					
Complete time series of 2+5 year estimation plus forecast windows, 2+5 years ending 19861231 to 20070801 (1075 windows with weekly steps)	simple vs. simple	-0.215	0.222	0.094	4.103	4.230	1244	0.088	0.343	3.656	0.072	0.303	0.171	0.276	0.224	44.49	1.000	1.000	1.000	1.000	
	GARCH/simple vs. simple	-0.226	0.223	0.094	4.088	4.230	1248	0.069	0.339	3.661	0.052	0.301	0.169	0.275	0.222	44.65	1.000	1.000	1.000	1.000	
	EGARCH/simple vs. simple	-0.316	0.176	0.056	3.734	3.956	1107	-0.014	0.331	3.283	-0.027	0.305	0.159	0.254	0.202	40.20	1.000	1.000	1.000	1.000	
	IGARCH/simple vs. simple	-0.205	0.229	0.098	4.126	4.279	1244	0.124	0.380	3.719	0.104	0.332	0.226	0.276	0.229	44.50	1.000	1.000	1.000	1.000	
	FIGARCH/simple vs. simple	-0.205	0.225	0.110	4.250	4.494	1405	0.103	0.379	4.150	0.072	0.332	0.222	0.284	0.241	50.21	1.000	1.000	1.000	1.000	
	EWMA/simple vs. simple	-0.528	0.329	0.150	5.171	5.760	2294	-0.145	0.523	7.718	-0.171	0.477	0.354	0.340	0.311	87.92	1.000	1.000	1.000	1.000	
	clear winner?				EGARCH/simple			EGARCH/simple			EGARCH/simple										
	simple vs. simple	-0.215	0.222	0.094	4.103	4.230	1244	0.088	0.343	3.656	0.072	0.303	0.171	0.276	0.224	44.49	1.000	1.000	1.000	1.000	
	GARCH vs. simple	-0.226	0.223	0.094	4.088	4.230	1248	0.069	0.339	3.666	0.052	0.301	0.169	0.275	0.223	44.71	0.980	1.000	1.000	0.002	
	EGARCH vs. simple	-0.343	0.175	0.054	3.893	4.127	1161	-0.032	0.317	2.347	-0.034	0.299	0.151	0.228	0.174	38.88	0.816	1.000	0.951	0.021	
	IGARCH vs. simple	-0.247	0.166	0.041	1.060	1.204	2.662	0.181	0.441	0.427	0.164	0.419	0.339	0.201	0.144	0.339	0.051	0.551	0.259	0.058	
	FIGARCH vs. simple	-0.205	0.224	0.109	4.298	4.544	1422	0.103	0.378	4.180	0.072	0.332	0.221	0.285	0.243	50.70	0.942	1.000	0.995	0.008	
EWMA vs. simple	-0.571	0.354	0.173	5.866	6.500	2600	-0.160	0.541	8.226	-0.186	0.493	0.371	0.368	0.339	98.85	0.661	1.000	0.912	0.038		
clear winner?				IGARCH			IGARCH			EGARCH			IGARCH								
Partial time series of 2+5 year estimation plus forecast windows, 2+5 years ending 19980107 to 20070801 (500 windows with weekly steps)	simple vs. simple	-0.174	0.229	0.094	0.288	0.388	0.538	0.057	0.292	0.202	0.072	0.270	0.142	0.139	0.049	0.124	1.000	1.000	1.000	1.000	
	GARCH/simple vs. simple	-0.190	0.229	0.094	0.268	0.380	0.525	0.037	0.288	0.194	0.052	0.275	0.136	0.137	0.043	0.118	1.000	1.000	1.000	1.000	
	EGARCH/simple vs. simple	-0.244	0.238	0.099	0.203	0.361	0.471	-0.029	0.292	0.180	-0.017	0.288	0.139	0.127	0.029	0.091	1.000	1.000	1.000	1.000	
	IGARCH/simple vs. simple	-0.145	0.245	0.107	0.333	0.436	0.912	0.095	0.328	0.262	0.107	0.302	0.197	0.142	0.057	0.151	1.000	1.000	1.000	1.000	
	FIGARCH/simple vs. simple	-0.173	0.248	0.110	0.272	0.414	0.725	0.066	0.319	0.238	0.088	0.297	0.190	0.139	0.049	0.117	1.000	1.000	1.000	1.000	
	EWMA/simple vs. simple	-0.426	0.365	0.240	0.187	0.567	0.584	-0.179	0.467	0.350	-0.162	0.465	0.334	0.112	0.031	0.061	1.000	1.000	1.000	1.000	
	clear winner?																				
	simple vs. simple	-0.174	0.229	0.094	0.288	0.388	0.538	0.057	0.292	0.202	0.072	0.270	0.142	0.139	0.049	0.124	1.000	1.000	1.000	1.000	
	GARCH vs. simple	-0.190	0.229	0.094	0.268	0.379	0.525	0.037	0.288	0.193	0.052	0.274	0.136	0.137	0.043	0.118	0.997	1.000	1.000	0.001	
	EGARCH vs. simple	-0.246	0.239	0.098	0.192	0.360	0.476	-0.037	0.288	0.170	-0.026	0.287	0.133	0.124	0.028	0.084	0.921	0.980	0.954	0.010	
	IGARCH vs. simple	-0.103	0.271	0.129	0.530	0.654	2.121	0.187	0.414	0.388	0.184	0.405	0.326	0.156	0.089	0.245	0.184	0.375	0.266	0.036	
	FIGARCH vs. simple	-0.172	0.248	0.109	0.276	0.415	0.684	0.067	0.318	0.235	0.087	0.296	0.190	0.139	0.048	0.113	0.986	1.000	0.996	0.003	
EWMA vs. simple	-0.439	0.407	0.253	0.179	0.591	0.609	-0.193	0.484	0.367	-0.174	0.482	0.352	0.114	0.032	0.062	0.693	0.958	0.919	0.022		
clear winner?							EGARCH														

\* Notes.

- Each column-group of six categorises six volatility forecast versus realisation methods according to: a time series parameter (minimum, maximum, mean, median or standard deviation); and a cross-sectional volatility forecast versus realisation error metric (MRE, MRAE or MRSE).
- A shaded number indicates a "winner" within a column-group of six; the winner is the number closest to zero.
- If a single volatility forecast versus realisation method produces the winner for a given time series parameter for all three error metrics, the volatility forecast model is indicated as a clear winner for that time series parameter.

**Table 2: Time series parameters of volatility forecasting errors for a 5+5 year estimation plus forecast rolling window**

Volatility forecast versus volatility realisation, estimation methods		Column-groups categorised by time series parameter, and cross-sectional volatility forecast versus realisation error metric *															Cross-sectional fraction of eligible companies with valid forecasted volatility				
		minimum			maximum			mean			median			standard deviation			min	max	mean	st. dev.	
		MRE	MRAE	MRSE	MRE	MRAE	MRSE	MRE	MRAE	MRSE	MRE	MRAE	MRSE	MRE	MRAE	MRSE					
Complete time series of 5+5 year estimation plus forecast windows, 5+5 years ending 1990Q103 to 2007Q801 (918 windows with weekly steps)	simple vs. simple	-0.151	0.230	0.094	3.152	3.195	585.7	0.157	0.374	2.204	0.175	0.354	0.252	0.254	0.188	22.79	1.000	1.000	1.000	1.000	
	GARCH/simple vs. simple	-0.154	0.231	0.092	3.142	3.186	585.7	0.146	0.370	2.198	0.166	0.346	0.241	0.253	0.187	22.79	1.000	1.000	1.000	1.000	
	EGARCH/simple vs. simple	-0.217	0.168	0.046	3.071	3.142	585.6	0.109	0.362	2.179	0.130	0.350	0.231	0.249	0.178	22.79	1.000	1.000	1.000	1.000	
	IGARCH/simple vs. simple	-0.138	0.253	0.108	3.256	3.304	585.9	0.237	0.448	2.340	0.266	0.417	0.404	0.257	0.203	22.80	1.000	1.000	1.000	1.000	
	FIGARCH/simple vs. simple	-0.141	0.241	0.096	3.142	3.277	625.5	0.192	0.421	2.445	0.242	0.378	0.370	0.260	0.205	24.33	1.000	1.000	1.000	1.000	
	EWMA/simple vs. simple	-0.497	0.335	0.182	2.675	3.320	586.2	-0.144	0.504	2.305	-0.167	0.473	0.340	0.223	0.176	22.81	1.000	1.000	1.000	1.000	
	clear winner?																				
	EGARCH/simple																				
	simple vs. simple	-0.151	0.230	0.094	3.152	3.195	585.7	0.157	0.374	2.204	0.175	0.354	0.252	0.254	0.188	22.79	1.000	1.000	1.000	1.000	
	GARCH vs. simple	-0.154	0.219	0.091	1.030	1.078	31.31	0.108	0.335	0.270	0.149	0.313	0.186	0.165	0.082	1.033	0.910	1.000	0.982	0.020	
	EGARCH vs. simple	-0.235	0.166	0.045	1.069	1.137	35.15	0.063	0.324	0.243	0.105	0.321	0.194	0.162	0.065	1.157	0.780	1.000	0.914	0.037	
	IGARCH vs. simple	-0.248	0.173	0.048	1.306	1.424	3.537	0.410	0.605	0.732	0.410	0.566	0.618	0.277	0.236	0.519	0.120	0.468	0.288	0.056	
	FIGARCH vs. simple	-0.141	0.246	0.099	3.142	3.277	625.5	0.192	0.420	2.454	0.243	0.377	0.370	0.261	0.206	24.37	0.933	1.000	0.996	0.008	
	EWMA vs. simple	-0.510	0.334	0.184	0.427	0.884	1.254	-0.178	0.501	0.401	-0.182	0.489	0.355	0.157	0.073	0.157	0.378	1.000	0.916	0.052	
clear winner?																					
EWMA																					
EGARCH																					
Partial time series of 5+5 year estimation plus forecast windows, 5+5 years ending 2001Q110 to 2007Q801 (343 windows with weekly steps)	simple vs. simple	-0.079	0.230	0.094	0.271	0.360	0.259	0.061	0.281	0.165	0.031	0.278	0.160	0.111	0.037	0.053	1.000	1.000	1.000	1.000	
	GARCH/simple vs. simple	-0.089	0.231	0.092	0.257	0.350	0.246	0.048	0.277	0.160	0.018	0.275	0.156	0.110	0.034	0.050	1.000	1.000	1.000	1.000	
	EGARCH/simple vs. simple	-0.120	0.243	0.098	0.252	0.365	0.365	0.031	0.290	0.180	0.000	0.284	0.175	0.114	0.032	0.060	1.000	1.000	1.000	1.000	
	IGARCH/simple vs. simple	-0.045	0.273	0.136	0.375	0.468	0.498	0.150	0.351	0.273	0.097	0.329	0.255	0.123	0.057	0.096	1.000	1.000	1.000	1.000	
	FIGARCH/simple vs. simple	-0.074	0.256	0.131	0.323	0.419	0.391	0.084	0.311	0.236	0.038	0.288	0.219	0.131	0.048	0.076	1.000	1.000	1.000	1.000	
	EWMA/simple vs. simple	-0.418	0.378	0.253	0.144	0.541	0.431	-0.181	0.456	0.323	-0.177	0.456	0.318	0.130	0.025	0.034	1.000	1.000	1.000	1.000	
	clear winner?																				
	simple vs. simple	-0.079	0.230	0.094	0.271	0.360	0.259	0.061	0.281	0.165	0.031	0.278	0.160	0.111	0.037	0.053	1.000	1.000	1.000	1.000	
	GARCH vs. simple	-0.097	0.219	0.091	0.246	0.341	0.230	0.035	0.267	0.135	-0.002	0.255	0.117	0.113	0.034	0.044	0.966	0.998	0.987	0.010	
	EGARCH vs. simple	-0.134	0.232	0.093	0.236	0.358	0.374	0.013	0.277	0.154	-0.028	0.265	0.122	0.118	0.032	0.060	0.863	0.950	0.919	0.014	
	IGARCH vs. simple	-0.005	0.301	0.144	0.747	0.815	1.207	0.343	0.509	0.506	0.267	0.436	0.380	0.196	0.145	0.260	0.209	0.424	0.306	0.037	
	FIGARCH vs. simple	-0.074	0.255	0.130	0.322	0.419	0.391	0.084	0.310	0.236	0.038	0.288	0.218	0.131	0.047	0.076	0.989	1.000	0.998	0.002	
	EWMA vs. simple	-0.437	0.401	0.256	0.159	0.545	0.458	-0.193	0.472	0.339	-0.187	0.475	0.333	0.134	0.026	0.036	0.677	0.973	0.923	0.030	
	clear winner?																				
GARCH																					

\* Notes.

- Each column-group of six categorises six volatility forecast versus realisation methods according to: a time series parameter (minimum, maximum, mean, median or standard deviation); and a cross-sectional volatility forecast versus realisation error metric (MRE, MRAE or MRSE).
- A shaded number indicates a "winner" within a column-group of six; the winner is the number closest to zero.
- If a single volatility forecast versus realisation method produces the winner for a given time series parameter for all three error metrics, the volatility forecast model is indicated as a clear winner for that time series parameter.

## 7. References

- 1995, RiskMetrics Monitor, (JP Morgan).
- 2004, A Guide to IFRS 2, (Deloitte).
- Andersen, T, T Bollerslev, and S Lange, 1999, Forecasting Financial Market Volatility: Sample Frequency vis-a-vis Forecast Horizon, *Journal of Empirical Finance* 6, 457-477.
- Baillie, R, T Bollerslev, and H Mikkelsen, 1996, Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics* 74, 3-30.
- Black, F, 1976, The Pricing of Commodity Options, *Journal of Financial Economics* 3, 167-179.
- Bollerslev, T, 1986, Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics* 31, 307-327.
- Bollerslev, T, and H Mikkelsen, 1996, Modeling and Long-Memory in Stock Market Volatility, *Journal of Econometrics* 73, 151-184.
- Bollerslev, T, and H Mikkelsen, 1999, Long-term Equity Anticipation Securities and Stock Market Volatility Dynamics, *Journal of Econometrics* 92, 75-99.
- Bollerslev, T, and J Wooldridge, 1992, Quasi-maximum Likelihood Estimation and Inference in Dynamic Models with Time-varying Covariances, *Econometric Reviews* 11, 143-172.
- Brailsford, T, and R Faff, 1996, An Evaluation of Volatility Forecasting Techniques, *Journal of Banking and Finance* 20, 419-438.
- Chou, R, 1988, Volatility Persistence and Stock Valuations: Some Empirical Evidence using GARCH, *Journal of Applied Econometrics* 3, 279-294.
- Chung, C, 2001, Estimating the Fractionally Integrated GARCH Model, *National Taiwan University Discussion Paper*.
- Ding, Z, C Granger, and R Engle, 1993, A Long Memory Property of of Stock Market and a New Model, *Journal of Empirical Finance* 1, 83-106.

- Engle, R, 1982, Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica* 50, 987-1000.
- Engle, R, and T Bollerslev, 1986, Modelling the Persistence of Conditional Variances, *Econometric Reviews* 5, 1-50.
- He, C, T Terasvirta, and H Malmsten, 2002, Moment Structure of a Family of First-order Exponential GARCH Models, *Econometric Theory* 18, 868-885.
- Jorion, P, 1997. *Value at Risk: The New Benchmark for Controlling Market Risk* (Irwin, Burr Ridge, IL).
- Merton, R, 1973, Theory of Rational Option Pricing, *The Bell Journal of Economics and Management Science* 4, 141-183.
- Mikosch, T, and C Starica, 2004, Nonstationarities in Financial Time Series, the Long-range Dependence, and the IGARCH Effects, *Review of Economics and Statistics* 86, 378-390.
- Nelson, D, 1990, Stationarity and Persistence in the GARCH(1, 1) Model, *Econometric Theory* 6, 318-334.
- Nelson, D, 1991, Conditional Heteroskedasticity in Asset Returns: A New Approach, *Econometrica* 59, 347-370.
- Poon, S, and C Granger, 2003, Forecasting Volatility in Financial Markets: A Review, *Journal of Economic Literature* 41, 478-539.
- Weiss, A, 1986, Asymptotic Theory for ARCH Models: Estimation and Testing, *Econometric Theory* 2, 107-131.