

Generating Volatility Forecasts from Value at Risk Estimates

James W. Taylor

Saïd Business School

University of Oxford

Management Science, 2005, Vol. 51, pp. 712-725.

Address for Correspondence:

James W. Taylor
Saïd Business School
University of Oxford
Park End Street
Oxford OX1 1HP, UK

Tel: +44 (0)1865 288927
Fax: +44 (0)1865 288805
Email: james.taylor@sbs.ox.ac.uk

Generating Volatility Forecasts from Value at Risk Estimates

Generating Volatility Forecasts from Value at Risk Estimates

Abstract

Statistical volatility models rely on the assumption that the shape of the conditional distribution is fixed over time and that it is only the volatility that varies. The recently proposed conditional autoregressive value at risk (CAViaR) models require no such assumption, and allow quantiles to be modelled directly in an autoregressive framework. Although useful for risk management, CAViaR models do not provide volatility forecasts, which are needed for several other important applications, such as option pricing and portfolio management. It has been found that, for a variety of probability distributions, there is a surprising constancy of the ratio of the standard deviation to the interval between symmetric quantiles in the tails of the distribution, such as the 0.025 and 0.975 quantiles. This result has been used in decision and risk analysis to provide an approximation of the standard deviation in terms of quantile estimates provided by experts. Drawing on the same result, we construct financial volatility forecasts as simple functions of the interval between CAViaR forecasts of symmetric quantiles. Forecast comparison, using five stock indices and 20 individual stocks, shows that the method is able to outperform GARCH models and moving average methods.

Key words: Volatility Forecasting; Value at Risk; CAViaR Models.

1. Introduction

Volatility forecasting is important for many financial market applications, including option pricing and investment decisions. The empirical finding that series of returns often exhibit volatility clustering has led to the development of a variety of univariate time series methods for volatility forecasting. The popular GARCH class of models, as well as stochastic volatility models, rely on the assumption that the shape of the conditional distribution is fixed over time and that it is only the conditional volatility, and sometimes the conditional mean, that are assumed to vary. For example, the Gaussian or Student- t distribution is usually used within the GARCH maximum likelihood estimation procedure. If there is variation over time in the shape of the distribution, this is a likely source of error for the volatility forecasts produced by these models. By contrast, the recently proposed conditional autoregressive value at risk (CAViaR) models of Engle and Manganelli (2004) require no distributional assumptions. These models allow quantiles to be modelled directly in an autoregressive framework. The θ quantile of a financial return, r_t , is known as the value at risk (VaR), and is defined as the value, $Q_t(\theta)$, for which $P(r_t \leq Q_t(\theta)) = \theta$. As VaR is a risk management tool, the quantiles of interest are in the tails of the distribution. Although useful for risk management, CAViaR models do not provide volatility estimates. However, the appeal of CAViaR models in describing the behaviour of prices motivates consideration of how volatility forecasts might be derived from CAViaR quantile forecasts.

The problem of estimating the variance of a distribution from a small number of quantile estimates exists in several decision and risk analysis applications (Keefer and Bodily 1983). For example, in PERT analysis, estimates of the mean and variance of a distribution must often be derived from judgementally assessed quantile estimates. The results of Pearson and Tukey (1965) are frequently used to address this problem. They show that, for a variety of probability distributions, there is a surprising constancy of the ratio of the standard deviation to the interval between symmetric tail quantiles, $Q(\theta)$ and $Q(1-\theta)$. For example, they conclude that a simple approximation to the standard deviation is provided by the interval between $Q(0.025)$ and $Q(0.975)$ divided by 3.92. Clearly, this value would be appropriate if the distribution was Gaussian, but the interesting point is that it is also approximately correct for a variety of distributions. This suggests that, even though the

conditional volatility and distribution of financial returns may vary over time, the conditional volatility can be approximated by a constant simple function of the interval between symmetric conditional quantiles. This provides us with a basis for constructing volatility forecasts from quantile forecasts produced by CAViaR models or, indeed, other VaR methods. If CAViaR models are used, our proposed method constructs volatility forecasts from separate autoregressive models for the left-tail and right-tail quantiles, $Q(\theta)$ and $Q(1-\theta)$, respectively. By contrast, GARCH models use only an autoregressive model for the variance. If the left and right tails of the conditional distribution are driven by different forces over time, our approach should capture the evolution of the variance better than GARCH models.

In Sections 2 and 3, we briefly review the literatures on volatility forecasting and VaR estimation, respectively. Section 4 describes the new volatility forecasting method. In Section 5, we consider the method's news impact curve and evaluate the method's forecasting performance. Section 6 provides a summary and concluding comments.

2. Volatility Forecasting

Volatility forecasts are produced by either market-based methods or time series methods. Market-based forecasting involves the calculation of implied volatility from current option prices by solving the Black and Scholes option pricing model for the volatility that results in a price equal to the market price. In this paper, our focus is on the development of a new time series method. These methods provide estimates of the conditional variance, $\sigma_t^2 = \text{var}(r_t|I_{t-1})$, of the log return, r_t , at time t conditional upon I_{t-1} , the information set of all observed returns up to time $t-1$. This can be viewed as the variance of an error (or residual) term, ε_t , defined by $\varepsilon_t = r_t - E(r_t|I_{t-1})$, where $E(r_t|I_{t-1})$ is a conditional mean term, which is often assumed to be zero or a constant. ε_t is often referred to as the price "shock" or "news". In the next two sections, we review the two most popular times series approaches: moving averages and GARCH models.

2.1. Moving Averages

The simplest approach to volatility forecasting is to estimate the variance as a simple moving average of past squared shocks. A problem with this method is that the choice of how many past periods to include in the moving average is arbitrary. Including too few observations will lead to large sampling error, while using too many will result in predictions that are slow to react to changes in the true volatility. This issue and the strong appeal in giving more weight to more recent observations motivate the use of an exponentially weighted moving average of past squared shocks. If a long history of observations is used, the one step-ahead variance estimator, $\hat{\sigma}_{t+1}^2$, can be written in the simple exponential smoothing recursive form with smoothing parameter, α :

$$\hat{\sigma}_{t+1}^2 = \alpha \varepsilon_t^2 + (1-\alpha)\hat{\sigma}_t^2$$

Although for daily returns a value of 0.06 has been recommended for α (*RiskMetrics* 1996), a more appealing approach is to optimise the parameter value by minimising the in-sample sum of squared deviations between the variance forecasts, $\hat{\sigma}_{t+1}^2$, and the squared error, ε_{t+1}^2 , which serves as a proxy for actual variance, which is unobservable. For moving average and exponential smoothing methods, the multiperiod variance forecast, $\hat{\sigma}_{t,k}^2$, made from origin t , for the return over a holding period consisting of the next k -periods is calculated by simply multiplying the one step-ahead forecast, $\hat{\sigma}_{t+1}^2$, by k .

2.2. GARCH Models

Generalized autoregressive conditional heteroskedasticity (GARCH) models (see Engle 1982 and Bollerslev 1986) are the most widely used statistical models for volatility. GARCH models express the conditional variance as a linear function of lagged squared error terms and lagged conditional variance terms. For example, the GARCH(1,1) model is shown in the following expression

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$$

where ω , α and β are parameters. The multiperiod variance forecast, $\hat{\sigma}_{t,k}^2$, is calculated as the sum of the variance forecasts for each of the k periods making up the holding period.

$$\hat{\sigma}_{t,k}^2 = \frac{\omega k}{1-\alpha-\beta} + \left(\hat{\sigma}_{t+1}^2 - \frac{\omega}{1-\alpha-\beta} \right) \left(\frac{1-(\alpha+\beta)^k}{1-\alpha-\beta} \right) \quad (1)$$

where $\hat{\sigma}_{t+1}^2$ is the one step-ahead variance forecast. Empirical results for the GARCH(1,1) model have shown that often $\beta \approx (1-\alpha)$. The model in which $\beta = (1-\alpha)$ is termed *integrated* GARCH (IGARCH) (see Nelson 1990). Exponential smoothing has the same formulation as the IGARCH(1,1) model with the additional restriction that $\omega = 0$. The IGARCH(1,1) multiperiod forecast is written as

$$\hat{\sigma}_{t,k}^2 = \frac{1}{2}k(k-1)\omega + k\hat{\sigma}_{t+1}^2 \quad (2)$$

Stock return volatility is often found to be greater following a negative return than a positive return of equal size. This “leverage effect” has prompted the development of a number of GARCH models that allow for asymmetry. The first asymmetric formulation was the exponential GARCH model of Nelson (1991). In this log formulation for volatility, the impact of lagged squared residuals is exponential, which may exaggerate the impact of large shocks. A simpler asymmetric model is the GJR GARCH model of Glosten et al. (1993). The GJR GARCH(1,1) model is given by

$$\sigma_t^2 = \omega + (1 - I[\varepsilon_{t-1} > 0])\alpha \varepsilon_{t-1}^2 + (I[\varepsilon_{t-1} > 0])\gamma \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2$$

where ω , α , γ and β are parameters; and $I[\cdot]$ is the indicator function. Typically, it is found that $\alpha > \gamma$, which indicates the presence of the leverage effect. The assumption that the median of the distribution of ε_t is zero implies that the expectation of the indicator function is 0.5, which enables the derivation of the following multiperiod forecast expression:

$$\hat{\sigma}_{t,k}^2 = \frac{\omega k}{1 - \frac{1}{2}(\alpha + \gamma) - \beta} + \left(\hat{\sigma}_{t+1}^2 - \frac{\omega}{1 - \frac{1}{2}(\alpha + \gamma) - \beta} \right) \left(\frac{1 - \left(\frac{1}{2}(\alpha + \gamma) + \beta\right)^k}{1 - \frac{1}{2}(\alpha + \gamma) - \beta} \right) \quad (3)$$

GARCH parameters are estimated by maximum likelihood, which requires the assumption that the standardised errors, ε_t/σ_t , are independent and identically distributed (i.i.d.). Although a Gaussian assumption is common, the distribution is often fat tailed, which has prompted the use of the Student- t distribution (Bollerslev 1987) and the generalised error distribution (Nelson 1991).

Stochastic volatility models provide an alternative statistical volatility modelling approach (see Ghysels et al. 1996). However, estimation of these models has proved difficult and consequently, they are not as widely used as GARCH models. Andersen et al. (2003) show how daily exchange rate volatility can be forecasted by fitting long-memory, or fractionally-integrated, autoregressive and vector autoregressive models to the log of realised daily volatility constructed from half-hourly

returns. Although results for this approach are impressive, such high frequency data is not available to many forecasters and so there is still great interest in methods applied to daily data. A useful review of the volatility forecasting literature is provided by Poon and Granger (2003).

3. Value at Risk Methods

Manganelli and Engle (2004) describe how the VaR literature contains three different categories of methods: parametric, nonparametric and semiparametric. *Parametric* approaches involve a parameterisation of the behaviour of prices. Quantiles are estimated using a volatility forecast with an assumption for the type of the distribution, such as Gaussian. Typically, exponential smoothing or a GARCH model is used to forecast the volatility.

The most widely used *nonparametric* method is historical simulation, which requires no distributional assumptions and estimates the VaR as the quantile of the empirical distribution of historical returns from a moving window of the most recent periods. The same issues that motivated the use of exponentially weighted moving averages for volatility forecasting (see Section 2.1) prompted Boudoukh et al. (1998) to propose the analogy of this method for quantiles. We term this the BRW method. It involves allocating to the most recent n returns, exponentially decreasing weights, which sum to one, $\left(\frac{1-\lambda}{1-\lambda^n}\right), \left(\frac{1-\lambda}{1-\lambda^n}\right)\lambda, \left(\frac{1-\lambda}{1-\lambda^n}\right)\lambda^2, \dots, \left(\frac{1-\lambda}{1-\lambda^n}\right)\lambda^{n-1}$. The returns are then ordered in ascending order and, starting at the lowest return, the weights are summed until a value of θ is reached. The θ quantile estimate is set as the return that corresponds to the final weight used in this summation. Linear interpolation is used if the estimate falls between two returns. Boudoukh et al. experimented with arbitrary choices of 0.97 and 0.99 for the parameter λ .

Included in the *semiparametric* VaR category are methods that use extreme value theory and methods that use quantile regression, such as the CAViaR models introduced by Engle and Manganelli (2004). CAViaR models involve direct autoregressive modelling of the conditional quantiles and thus do not involve any distributional assumptions. Engle and Manganelli present the following four CAViaR models:

Indirect GARCH(1,1) CAViaR:
$$Q_t(\theta) = (1 - 2I[\theta < 0.5])\left(\omega + \alpha Q_{t-1}(\theta)^2 + \beta \varepsilon_{t-1}^2\right)^{\frac{1}{2}} \quad (4)$$

Adaptive CAViaR:
$$Q_t(\theta) = Q_{t-1}(\theta) + \alpha(\theta - I[\varepsilon_{t-1} \leq Q_{t-1}(\theta)]) \quad (5)$$

Symmetric Absolute Value CAViaR:
$$Q_t(\theta) = \omega + \alpha Q_{t-1}(\theta) + \beta |\varepsilon_{t-1}| \quad (6)$$

Asymmetric Slope CAViaR:
$$Q_t(\theta) = \omega + \alpha Q_{t-1}(\theta) + \beta_1(\varepsilon_{t-1})^+ + \beta_2(\varepsilon_{t-1})^- \quad (7)$$

where $Q_t(\theta)$ is the conditional θ quantile; ω , α , β and β_i are parameters; and $(x)^+ = \max(x,0)$ and $(x)^- = \min(x,0)$.

If the error term standardised by the GARCH(1,1) volatility estimate, ε_t/σ_t , is i.i.d., the Indirect GARCH(1,1) CAViaR model is the same as the GARCH(1,1) model. The indicator function in the Adaptive CAViaR model has the effect of reducing the next quantile estimate if, in the current period, the quantile estimate is greater than the error. If the error exceeds the Adaptive model quantile estimate, the next estimate is increased. The Symmetric Absolute Value and the Asymmetric Slope CAViaR models rely on the magnitude of the error, rather than the squared error as in GARCH models. The Asymmetric Slope model was designed specifically to model the asymmetric leverage effect.

CAViaR model parameters are estimated using the quantile regression minimisation in expression (8), which was introduced by Koenker and Bassett (1978).

$$\min \left(\sum_{t | y_t \geq Q_t(\theta)} \theta |y_t - Q_t(\theta)| + \sum_{t | y_t < Q_t(\theta)} (1-\theta) |y_t - Q_t(\theta)| \right) \quad (8)$$

where $Q_t(\theta)$ is the model for the θ quantile of the dependent variable y_t . White (1994) presents theoretical support for the use of quantile regression to estimate non-linear quantile models. Engle and Manganelli (2004) provide results for the asymptotic distribution of the CAViaR parameter estimates so that tests of significance can be performed. In this paper, we consider how CAViaR quantile forecasts can be used to construct volatility forecasts.

4. Using Value at Risk Estimates to Generate Volatility Forecasts

4.1. Approximating Standard Deviations Using the Interval Between Symmetric Quantiles

In decision and risk analysis, it is often useful to approximate a probability distribution by a small number of quantile estimates. For example, a distribution may be needed for a Monte Carlo simulation, but the only information available may be judgementally assessed quantiles (see Keefer 1994, Clemen 1996). In other applications, such as PERT analysis, expert assessments of quantiles are sometimes used to estimate the mean and variance (see Keefer and Verdini 1993). The work of Pearson and Tukey (1965) has been influential in addressing these issues.

Pearson and Tukey found that the ratio of the standard deviation to the interval between symmetric quantiles, $Q(\theta)$ and $Q(1-\theta)$, in the tails of the distribution is remarkably constant for a variety of distributions. Their analysis considered 98%, 95% and 90% intervals. They propose the following simple approximations for the standard deviation in terms of estimated tail quantiles:

$$\hat{\sigma} = \frac{\hat{Q}(0.99) - \hat{Q}(0.01)}{4.65} \quad \hat{\sigma} = \frac{\hat{Q}(0.975) - \hat{Q}(0.025)}{3.92} \quad \hat{\sigma} = \frac{\hat{Q}(0.95) - \hat{Q}(0.05)}{3.25} \quad (9)$$

Note that, for a Gaussian distribution, the correct denominators in this expression would be: $2 \times 2.326 = 4.653$, $2 \times 1.960 = 3.920$ and $2 \times 1.645 = 3.290$, respectively. Pearson and Tukey show that the accuracy of these approximations depends on the values of the skewness and kurtosis of the distribution. They found that the approximation based on the 90% interval was the most robust to different skewness and kurtosis values.

4.2. Forecasting Volatility Using the Interval Between Symmetric VaR Quantile Forecasts

The Pearson and Tukey (1965) approximations in expression (9) provide a convenient basis from which to generate volatility forecasts for financial returns from quantile estimates produced by VaR methods, most notably CAViaR models. By contrast with the use of the approximations in decision and risk analysis, in the context of financial returns, we have historical times series of quantile estimates and of realisations from the returns distribution. This enables estimation of parameters for the approximations in expression (9), rather than reliance on the Pearson and Tukey values in each denominator. We propose a least squares (LS) regression of the squared errors, ε_i^2 ,

which act as a proxy for actual variance, on the square of the interval between symmetric quantile estimates. One step-ahead quantile forecasts, $\hat{Q}_{t+1}(1-\theta)$ and $\hat{Q}_{t+1}(\theta)$, can then be substituted into the resultant model to deliver one step-ahead variance predictions as in the following expression:

$$\hat{\sigma}_{t+1}^2 = \alpha_1 + \beta_1 \left(\hat{Q}_{t+1}(1-\theta) - \hat{Q}_{t+1}(\theta) \right)^2 \quad (10)$$

where α_1 and β_1 are the parameters estimated by the LS regression.

The simplest approach to generating multiperiod variance forecasts is to multiply the one step-ahead forecasts by the duration, k , of the holding period, which is the approach used for moving average methods discussed in Section 2.1. However, this approach requires the assumption that the variance is constant for each day in the holding period. This assumption is reasonable if the quantile forecasts are produced by the historical simulation method, the BRW method, or the Adaptive CAViaR model. However, it is inappropriate if the quantile forecasts come from one of the other CAViaR models. Unfortunately, analytical formulae for multiperiod quantile forecasts from these models do not exist. A rather complex simulation could be used to produce these multiperiod forecasts. But as a simpler alternative, we propose that, for each holding period, a separate LS regression is run of the realised multiperiod variance, $\sigma_{Rt,k}^2$, on the square of the interval between symmetric quantile estimates, $\left(\hat{Q}_{t+1}(1-\theta) - \hat{Q}_{t+1}(\theta) \right)^2$. For daily returns, if we make the reasonable assumptions that the conditional mean is constant over the k days and that there is no autocorrelation between successive daily shocks, then, for the holding period of duration k days starting in period $t+1$, the realised multiperiod variance can be calculated as

$$\sigma_{Rt,k}^2 = \sum_{i=1}^k \varepsilon_{t+i}^2 \quad (11)$$

The multiperiod variance forecasts are then produced by substituting one step-ahead quantile forecasts into the resultant model as in the following expression:

$$\hat{\sigma}_{t,k}^2 = \alpha_k + \beta_k \left(\hat{Q}_{t+1}(1-\theta) - \hat{Q}_{t+1}(\theta) \right)^2 \quad (12)$$

where α_k and β_k are the parameters estimated by the LS regression.

The volatility forecasting approach that we have described is a two-stage method. First, the parameters of the θ and $(1-\theta)$ quantile models are estimated, and then, in a second stage, the

parameters in expression (10) or (12) are estimated. This procedure could possibly be made more efficient if all the parameters are estimated simultaneously. This would require the extension of the CAViaR models to the bivariate case where θ and $(1-\theta)$ quantile processes are modelled together in order to capture their potential interactions. We reserve consideration of this for future work, and in all empirical work in this paper we use the simpler two-stage estimation approach.

It is worth noting that the regressions in expressions (10) and (12) can be viewed as restricted versions of more general regressions. For example, expression (10) can be written as the following:

$$\hat{\sigma}_{t+1}^2 = \alpha_1 + \beta_1 \hat{Q}_{t+1}^2(1-\theta) + \beta_2 \hat{Q}_{t+1}^2(\theta) + \beta_3 \hat{Q}_{t+1}(1-\theta) \hat{Q}_{t+1}(\theta) \quad (13)$$

where the restrictions are $\beta_2 = \beta_1$ and $\beta_3 = -2\beta_1$. Firstly, these restrictions could be tested, and, secondly, there may be benefit in using the unrestricted regressions to produce forecasts.

Instead of basing the approach on the Pearson and Tukey approximations, one might also consider one of the proposed modifications (e.g., Moder and Rogers 1968, Keefer and Bodily 1983, Johnson 2002). However, several of these modifications use the median, and this is unlikely to be beneficial as the median is very often close to zero for daily financial returns.

Our proposal to estimate volatility by using the interval between quantiles in the tails of the distribution has similarities with range-based volatility estimation (e.g., Parkinson 1980, Garman and Klass 1980, Alizadeh et al. 2002, Brandt and Diebold 2005). This class of methods bases estimation on the difference between the highest and lowest log price, which are essentially the quantiles corresponding to $\theta=1$ and $\theta=0$, respectively. High and low price quotes are widely available in the financial pages of newspapers. Interestingly, Brandt and Diebold (2005) show how covariance forecasts can be constructed from range-based volatility forecasts. This same approach could be used to construct covariance predictions from our VaR-based volatility forecasts.

5. Empirical Comparison of Volatility Forecasting Methods

We compared the accuracy of the volatility forecasts from our new VaR-based method with those from moving average methods and GARCH models. We used the following stock indices: the French CAC40, the German DAX30, British FTSE100, Japanese Nikkei225 and the US S&P500. The sample period used in our study consisted of 10 years of daily data, from 29 April 1993 to 28

April 2003. This period delivered 2,608 log returns. We used the first 2,089 returns to estimate method parameters and the remaining data to evaluate 500 post-sample forecasts for the volatility over the following holding periods: 1 day, 10 days and 20 days. Following common practice, we did not estimate models for the conditional mean of each series (see Poon and Granger 2003). For all five series, we subtracted from each return, r_t , the mean, μ , of the 2,089 in-sample returns. The volatility forecasting methods were applied to the resultant errors, $\varepsilon_t = r_t - \mu$. In the next two sections, we present the forecasting methods considered in our study.

5.1. Moving Average Methods and GARCH Volatility Models

We produced volatility forecasts from a 30-day simple moving average. We also implemented exponential smoothing. We found that optimising the exponential smoothing α parameter gave slightly better results than the fixed value of 0.06, suggested by *RiskMetrics* (1996). For simplicity, in Section 5.3, we report in detail only the results for the optimised method.

We included in the study the following three GARCH models, which were described in Section 2.1: GARCH(1,1), IGARCH(1,1) and GJR(1,1). Our choice of the (1,1) specification for all three models was based on our analysis of the in-sample period of 2,608 returns and on the general popularity of this order for GARCH models. We derived the model parameters using maximum likelihood based on a Student- t distribution with optimised degrees of freedom. We produced multi-period variance forecasts from these three models using the formulae in expressions (1)-(3).

The initial work of Koenker and Bassett (1978) on quantile regression emphasised its robustness to non-Gaussian, especially long-tailed, situations. As CAViaR models are essentially quantile regression models, if they are used within our new VaR-based approach, then the approach would seem to have an appeal of robustness. In view of this, we also included in our comparative study two robust benchmark approaches. Both of these estimated the GARCH models described above using winsorized datasets (see Hoaglin et al. 1983). The first approach was simplistic and set the largest 1% of the 2,089 in-sample observations to the value of the unconditional 0.99 quantile of these returns, and set the lowest 1% of these 2,089 observations to the value of the unconditional 0.01 quantile. The second approach based the winsorization on CAViaR quantile models for the 0.99 and 0.01 quantiles. All in-sample observations

larger than their corresponding in-sample fitted conditional 0.99 quantile were set equal to this value, and all in-sample observations lower than their corresponding in-sample fitted conditional 0.01 quantile were set equal to this value. For brevity, in Section 5.3, we discuss only the results of these approaches using the GJR-GARCH(1,1) model and the Asymmetric Slope CAViaR model, as this led to the best results. We describe CAViaR model parameter estimation in the next section.

5.2. VaR-Based Volatility Forecasting

Using our proposed new method, variance forecasts were produced as in expressions (10) and (12) based on 98%, 95% and 90% intervals constructed from symmetric quantile forecasts produced from the following VaR methods, which were discussed in Section 3: historical simulation, BRW and the four CAViaR models in expressions (4) to (7). Note that, although the new method was motivated by the appeal of using CAViaR models, it can be used with quantile forecasts from any method.

We used one year of data in the moving windows for both the historical simulation and the BRW methods. We experimented with the fixed values of λ proposed by Boudoukh et al. (1998) for the BRW method, but found that greater quantile forecast accuracy resulted when we optimised λ , on the in-sample data, using the quantile regression summation (QR Sum) presented in expression (8). In the next section, we report only the results for the BRW method with optimised parameter.

We estimated the parameters for the CAViaR models using a procedure similar to that described by Engle and Manganelli (2004). For each model, we first generated 10^5 vectors of parameters from a uniform random number generator between 0 and 1. We then evaluated the QR Sum for each of the vectors. The 10 vectors that produced the lowest values for the function were used as initial values in a quasi-Newton algorithm. The QR Sum was then calculated for each of the 10 resulting vectors, and the one producing the lowest value of the function was chosen as the final parameter vector. The software Gauss was used for all computational work in this study.

In Table 1, for each of the five stock indices, we present the LS regression parameters, α_1 and β_1 , in expression (10), estimated for one step-ahead variance prediction from the VaR-based method using symmetric quantile forecasts produced by the Asymmetric Slope CAViaR model. Interestingly, many of the estimated parameters are close to the Pearson and Tukey (1965) values, which were given in

expression (9) and are shown in the final column of the table. In 13 of the 15 LS regressions, the constant is not significantly different from zero (at the 5% level). For all five indices, the parameter estimates corresponding to the 95% interval are not significantly different from the Pearson and Tukey values. This is also true for all three LS regressions involving the DAX30 index. In spite of this, we found that our LS regression parameters led to noticeably better volatility forecast accuracy than the Pearson and Tukey values, and so, in the remainder of the paper, we refer only to the results for our parameterisation.

----- Table 1 -----

For each of the five stock indices, we performed the Wald test of the restrictions $\beta_2=\beta_1$ and $\beta_3=2\beta_1$ in the regression of expression (13), which we discussed in Section 4.2. We considered, in turn, $\alpha=0.01, 0.025$ and 0.05 . The results are reported in detail in an appendix on the Online Supplement website. Interestingly, in the majority of cases, the restrictions were rejected (at the 5% level), suggesting that the unrestricted regression should be used to produce forecasts. However, in this introductory paper, we limit our focus to the simpler regressions in expressions (10) and (12), which link more intuitively to the approximations of Pearson and Tukey in expression (9).

The “news impact curve” (NIC) of Engle and Ng (1993) has been widely used to compare different GARCH models. The curve shows the impact of shocks, or news, ε_{t-1} , on the next period’s volatility prediction, $\hat{\sigma}_t$. Figure 1 compares the NICs for the new VaR-based method applied to the S&P500 index. The method uses 90% intervals based on symmetric quantiles estimated by the four CAViaR models in expressions (4) to (7). The NICs in Figure 1 are conditional on the quantile estimates in the previous period. We set these to be equal to the average of the estimation sample quantile estimates from the corresponding model. The x -axis in Figure 1 extends in both directions by three times the unconditional standard deviation of the estimation sample of returns.

----- Figure 1 -----

The symmetric convex shape of the NIC for the method based on the Indirect GARCH(1,1) model is similar to that of a GARCH(1,1) model. This is perhaps not surprising, given the similarities of these two models. The NIC for the method based on the Adaptive CAViaR model shows discrete jumps between three volatility levels. This is the result of the indicator function in the Adaptive model formulation in expression (5). An interesting feature of the NIC for the Adaptive model is that,

for relatively large shocks, it displays the asymmetric leverage effect, with greater volatility following a negative shock than a positive shock of equal size. As we might have anticipated, this effect is very apparent in the NIC for the new method based on the Asymmetric Slope CAViaR model.

The NICs corresponding to the method based on the Asymmetric Slope and Symmetric Absolute Value models each consist of two lines meeting at $\varepsilon_{t-1} = 0$. These lines are almost straight, which is due to the relatively small estimated values for α_1 in expression (10). For the Asymmetric Slope model, the values of α_1 are given in Table 1. If α_1 is zero in expression (10), the volatility estimate is proportional to the magnitude of the interval, which will be piecewise linear if the quantile estimates are from either the Asymmetric Slope or Symmetric Absolute Value CAViaR models.

In Figure 2, we plot the S&P500 returns and day-ahead forecasts from two implementations of the VaR-based method using 90% intervals between symmetric quantiles. The first implementation used the Adaptive CAViaR model to estimate the quantiles of the 90% intervals, and the second used the Asymmetric Slope CAViaR model. The figure shows both series of volatility predictions reacting to changes in the magnitude of the returns, with the forecasts based on the Asymmetric Slope CAViaR model being substantially more responsive. Of the two sets of forecasts, those produced using the Asymmetric Slope CAViaR model are far more similar to those from the GARCH models.

----- Figure 2 -----

5.3. Post-Sample Volatility Forecast Evaluation for Stock Indices

In Tables 2 to 4, we summarize the post-sample forecasting performance from the various methods for the three different holding periods. The summary measure reported in the tables is the coefficient of determination, R^2 , from the LS regression of realised multiperiod variance on the post-sample variance forecasts. The realised multiperiod variance is calculated as in expression (11). The R^2 is a measure of informational content, with larger values being better. The values in bold in each column of the tables indicate the best performing method for each index. The average of the five R^2 values for each method is presented in the final column of each table. Note that for the R^2 measure, the values of the VaR-based method parameters, α_i and β_i , in expressions (10) and (12), are irrelevant because a new constant and coefficient are produced by the R^2 LS regression. For the VaR-based methods, the R^2

values, therefore, reflect the informational content in the squared interval between symmetric quantile forecasts. This is of interest because better approaches may exist than ours for incorporating, in a variance forecast, the information provided by the squared interval.

----- Tables 2 to 4 -----

The results show that GJRGARCH outperformed GARCH and IGARCH, as well as the moving average methods. The CAViaR-based winsorized GJRGARCH approach performed only slightly better than the standard GJRGARCH model. The superiority of GJRGARCH over GARCH confirms the existence of the leverage effect in the stock indices. This point is supported by the results for the VaR-based methods, which show the method based on the Asymmetric Slope CAViaR model outperforming the others. The superiority of the method increases, as the holding period gets longer. Indeed, the differences between the performances of all the methods is more pronounced in Tables 3 and 4 for the 10-day and 20-day holding periods, respectively. Interestingly, the results for the new method based on 95% or 90% intervals from the Asymmetric Slope CAViaR model were better than for GJRGARCH for all five indices. Although the methods based on the Indirect GARCH and the Symmetric Absolute Value CAViaR models did not match the performance of GJRGARCH, it is encouraging to see that these symmetric VaR-based methods did, overall, outperform the symmetric GARCH models and the moving average methods. There seems little potential for the VaR-based method when quantiles are estimated by historical simulation, the BRW method or the Adaptive CAViaR model. The poor performance of the method based on the Adaptive CAViaR model is, perhaps, not surprising, given the NIC for the method in Figure 1, which in our view is an unappealing representation of the behaviour of stock return volatility.

Because the R^2 measure is unaffected by the approach used to estimate the parameters, α_i and β_i , in expressions (10) and (12), the results in Tables 2 to 4 imply that there is more informational content for volatility forecasting in the magnitude of the 90% and 95% intervals, constructed from the Asymmetric Slope CAViaR model, than in any other method considered in our study. Evaluating the methods using root mean squared error, we found that the relative performances of the methods were similar to those for the R^2 measure. These additional results are available in an appendix on the Online Supplement website.

To gain insight into the differences between the methods, we used the volatility forecasts with the Black-Scholes model to price an at-the-money call option, 20 days from expiration. The resulting

option prices differed noticeably depending upon the volatility forecasting method used. For example, for the final 500 periods of the S&P500 series, the option prices calculated using the 20-day volatility forecasts from the new method, based on 90% intervals from the Asymmetric Slope CAViaR model, differed from the option prices calculated using the corresponding volatility forecasts from the GJRARCH model by, on average, 8.5% of the average of the two prices.

Poon and Granger (2003) argue for the use of statistical tests when evaluating relative volatility forecasting performance. We performed encompassing tests to investigate whether the post-sample performance of the method based on the Asymmetric Slope CAViaR 90% intervals (in the bottom row of Tables 2 to 4) was significantly better than that of (non-winsorized) GJRARCH. With these tests, a combined forecast is formed as a weighted average of the two forecasts (see Granger and Newbold 1973 and 1986, Chong and Hendry 1986). If the weight on one method is zero, that method is said to be encompassed by the other. The model used for the test is of the following form

$$\sigma_{Rt,k}^2 = w\hat{\sigma}_{Ct,k}^2 + (1-w)\hat{\sigma}_{Gt,k}^2 + e_t$$

where $\sigma_{Rt,k}^2$ is realised multiperiod variance; $\hat{\sigma}_{Ct,k}^2$ is the variance forecast from the CAViaR-based method; $\hat{\sigma}_{Gt,k}^2$ is the GJRARCH variance forecast; e_t is a residual term; and w is the combining weight estimated by the LS regression of $(\sigma_{Rt,k}^2 - \hat{\sigma}_{Gt,k}^2)$ on $(\hat{\sigma}_{Ct,k}^2 - \hat{\sigma}_{Gt,k}^2)$. We used non-overlapping realised and forecasted volatility data in the regression. If overlapping data is used, the models will suffer from considerable autocorrelation, rendering the test invalid (Christensen and Prabhala 1998).

Table 5 presents the results of the encompassing tests. For each series and each holding period, the table shows the estimated weight as well as p-values corresponding to tests of $w=1$ and $w=0$. Inability to reject $w=1$, implies that we cannot reject that the CAViaR-based method encompasses GJRARCH. In only one out of the 15 cases is the hypothesis $w=1$ rejected (at the 5% significance level). This one case is the 20-day holding period for the DAX30 index. Inability to reject $w=0$, implies that we cannot reject that GJRARCH encompasses the CAViaR-based method. The hypothesis is rejected in 10 out of the 15 cases (at the 5% level). It cannot be rejected for the DAX30 index at any of the three holding periods.

----- Table 5 -----

5.4. Post-Sample Quantile Forecast Evaluation for Stock Indices

In the previous section, we found that the new CAViaR-based volatility forecasting method is more successful when based on the 90% or 95% intervals than when based on the 98% interval. We also found that, when based on the 90% or 95% intervals, the method is able to outperform GARCH models. In this section, we report the results of a study that investigated whether these results are due to quantile forecasting performance. More specifically, we compared the day-ahead post-sample quantile forecasts produced by the Asymmetric Slope CAViaR model with those based on the GJRGARCH(1,1) model with Student-t distribution. We used the same 10 years of stock index data that we used in our volatility forecasting study of Sections 5.1 to 5.3, with the first 2,089 returns used for parameter estimation and the next 500 returns used for evaluation.

To evaluate the quantile forecasts, we use the three measures employed by Engle and Manganelli (2004): hit percentage, dynamic quantile test statistic and QR Sum. The hit percentage assesses the unconditional coverage of a θ quantile estimator. It is the percentage of observations falling below the estimator. Ideally, the percentage should be θ . With a sufficiently large sample, significance tests can be performed on the percentage using a Gaussian distribution and the standard error formula for a proportion. The Engle and Manganelli (2004) dynamic quantile test for conditional coverage evaluates the dynamic properties of a quantile estimator. It involves the joint test of whether the hit variable, defined as $Hit_t \equiv I[\varepsilon_t \leq \hat{Q}_t(\theta)] - \theta$, is distributed i.i.d. Bernoulli with probability θ , and is independent of the quantile estimator, $\hat{Q}_t(\theta)$. Ideally, Hit_t will have zero unconditional and conditional expectations. We included five lags of Hit_t in the test's regression framework to deliver a dynamic quantile test statistic (DQ), which, under the null hypothesis of perfect conditional coverage, is distributed $\chi^2(7)$. The third measure, the QR Sum, was defined in Section 5.2. It can be viewed as the equivalent of the root mean squared error for evaluating quantile forecast accuracy. Lower values of DQ and QR Sum are better.

In Table 6, we report the results for prediction of the 0.01 and 0.99 quantiles. We consider these two quantiles in the same table because they are used together within our new volatility forecasting approach. The values in bold indicate the best performing method for each quantile according to the

evaluation measure under consideration. The GJRGARCH and Asymmetric Slope CAViaR models perform similarly in terms of hit percentage, and, while GJRGARCH has better DQ, the CAViaR method seems to perform better in terms of QR Sum. These results are consistent with those in Table 2, where the volatility forecast evaluation measures were very similar for the standard GJRGARCH model and the new approach based on the Asymmetric Slope CAViaR 98% intervals.

----- Tables 6 and 7 -----

Table 7 reports results for the 0.05 and 0.95 quantiles. Overall, in this table, the CAViaR model outperforms the GJRGARCH approach. This is consistent with the volatility forecasting results in Table 2, where the CAViaR-based volatility forecasting method, using the 90% interval, tended to outperform GJRGARCH. Interestingly, the results in Table 7 show that for estimation of the DAX30 quantiles, the GJRGARCH method is not outperformed. This has strong similarities with the results of the encompassing tests, reported in Table 5, which were generally favourable for the CAViaR-based method using the 90% interval, except for the DAX30 series. The relative performances of the two methods for estimation of the 0.025 and 0.975 quantiles were similar to those reported in Table 7 for the 0.05 and 0.95 quantiles. These results are available in an appendix on the Online Supplement website.

The results in this section indicate that the success of the new CAViaR-based volatility forecasting method is related to the quality of the CAViaR quantile forecasts. The results also suggest that the new approach, based on 90% and 95% intervals, outperforms GJRGARCH because the corresponding CAViaR models are better able to model the tail dynamics of the conditional distribution.

5.5. Post-Sample Volatility Forecast Evaluation for Individual Stocks

To explore the robustness of our findings for the five stock indices, we repeated our comparison of volatility forecasting methods for the 20 individual S&P500 stocks that had highest market capitalization at the end of the year 2003. We used the same in-sample and post-sample daily periods employed in our study of the stock indices. The 20 individual stocks are listed in Table 8, in descending order of market capitalization, along with values of skewness and excess kurtosis calculated for each series using all 2,608 returns. The Procter and Gamble returns series contains a very large outlier, and this

is reflected in the very large values for the skewness and excess kurtosis. For completeness, we include at the bottom of Table 8 the skewness and excess kurtosis values for the five stock indices.

----- Tables 8 and 9 -----

In Table 9, for brevity, we report only the mean R^2 measure for each of the three holding periods. For the stock indices, these summary measures were presented in the final columns of Tables 2 to 4. Table 9 shows that the relative performance of the methods was similar for the individual stocks to that described earlier for the stock indices. The best performing approach is again the new volatility forecasting approach based on the Asymmetric Slope CAViaR intervals. A summary of the RMSE results for the 20 individual stocks is available in an appendix on the Online Supplement website.

As the accuracy of the Pearson and Tukey (1965) approximations depends on the degree of skewness and kurtosis, one might have surmised that the relative performance of our new approach would be related to the extent of skewness and kurtosis in each series of returns. However, we were unable to find any such relationship in our results for the 20 individual stocks.

6. Summary and Concluding Comments

Volatility forecasts are often used as a basis for estimating VaR. In this paper, we have shown how VaR estimates can be used as a basis for producing volatility forecasts. The motivation for this is the recently proposed CAViaR models, which provide an appealing way to model financial returns. The autoregressive nature of CAViaR quantile models has similarities with the widely used GARCH models, but unlike these models, CAViaR models require no distributional assumptions.

Drawing on the work of Pearson and Tukey (1965), our proposed new method involves generating variance forecasts as linear functions of the square of the interval between symmetric quantiles, which have been estimated by a VaR method. As in the decision and risk analysis literature, we found that basing the method on 95% and 90% intervals tended to be more successful than the use of 98% intervals. Our best results were achieved using the Asymmetric Slope CAViaR model, which accommodates the leverage effect in stock returns. Using error summary measures and encompassing tests to evaluate post-sample forecasting accuracy for five stock indices and 20 individual stocks, we found that, overall, this method outperformed GARCH models and moving average methods.

Although our proposed method involves an approximation, the same is also true of other univariate time series methods. Poon and Granger (2003) comment that GARCH models can be thought of as approximating a deeper time-varying construction, possibly involving several economic variables. The same authors also note that GARCH models fail to account for all the tail thickness in returns, even when a Student- t distribution is used within the maximum likelihood procedure. By directly modelling the tail quantiles, our CAViaR-based volatility forecasting method should be better able to account for the characteristics of the tails of the distribution. In particular, if the left and right tails of the conditional distribution are driven by different forces over time, our approach should capture the evolution better than GARCH models, which rely on a single autoregressive model for the variance.

Pearson and Tukey (1965) comment that there is potential for estimating the skewness and kurtosis from a knowledge of the quantiles. This suggests that quantile forecasts from CAViaR models could be used to construct forecasts for time-varying higher moments in financial returns, which are needed in a variety of finance applications (see Jondeau and Rockinger 2003).

Acknowledgements

We would like to acknowledge the insightful comments of two anonymous referees. We are also grateful to Mike Staunton for his help with the data.

References

- Alizadeh, S., M.W. Brandt, F.X. Diebold. 2002. Range-based estimation of stochastic volatility models. *Journal of Finance* **57** 1047-1091.
- Andersen, T.G., T. Bollerslev, F.X. Diebold, P. Labys. 2003. Modeling and forecasting realized volatility. *Econometrica* **71** 529-626.
- Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* **31** 307-327.
- Bollerslev, T. 1987. A conditionally heteroskedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics* **69** 542-547.
- Boudoukh, J., M. Richardson, R.F. Whitelaw. 1998. The best of both worlds. *Risk* **11** May 64-67.

- Brandt, M.W., Diebold, F.X. 2005. A no-arbitrage approach to range-based estimation of return covariances and correlations. *Journal of Business* forthcoming.
- Chong, Y.Y., D.F. Hendry. 1986. Econometric evaluation of linear macroeconomic models. *Review of Economic Studies* **53** 671-690.
- Christensen, B.J., N.R. Prabhala. 1998. The relation between implied and realized volatility. *Journal of Financial Economics* **50** 125-150.
- Clemen, R.T. 1996. *Making Hard Decisions: An Introduction to Decision Analysis* (2nd ed.). Duxbury Press, Belmont, CA.
- Engle, R.F. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* **50** 987-1008.
- Engle, R.F., S. Manganelli. 2004. CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business and Economic Statistics* **22** 367-381.
- Engle R.F., V.K. Ng. 1993. Measuring and testing the impact of news on volatility. *Journal of Finance* **48** 1749-1778.
- Garman, M.B., M.J. Klass. 1980. On the estimation of security price volatilities from historical data. *Journal of Business* **53** 67-78.
- Ghysels, E., A.C. Harvey, E. Renault. 1996. Stochastic volatility. G.S. Maddala, C.R. Rao, H.D. Vinod, eds. *Handbook of Statistics: Statistical Methods in Finance, Vol. 14*. Elsevier Science:, Amsterdam.
- Glosten, L.R., R. Jagannathan, D.E. Runkle. 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* **48** 1779-1801.
- Granger, C.W.J., P. Newbold. 1973. Some comments on the evaluation of economic forecasts. *Applied Economics* **5** 35-47.
- Granger, C.W.J., P. Newbold. 1986. *Forecasting Economic Time Series*, 2nd ed. Academic Press, Orlando.
- Hoaglin, D.C., F. Mosteller, J.W. Tukey. 1983. *Understanding Robust and Exploratory Data Analysis*, John Wiley, New York.

- Johnson, D. 2002. Triangular approximations for continuous random variables in risk analysis. *Journal of the Operational Research Society* **53** 457-467.
- Jondeau, E., M. Rockinger. 2003. Conditional volatility, skewness, and kurtosis: Existence, persistence, and comovements. *Journal of Economic Dynamics and Control* **27** 1699-1737.
- Keefer, D.L. 1994. Certainty equivalents for three-point discrete-distribution approximations. *Management Science* **40** 760-773.
- Keefer, D.L., S.E. Bodily. 1983. Three-point approximations for continuous random variables. *Management Science* **29** 595-609.
- Keefer, D.L., W.A. Verdini. 1993. Better estimation of PERT activity time parameters. *Management Science* **39** 1086-1091.
- Koenker, R.W., G.W. Bassett. 1978. Regression quantiles. *Econometrica* **46** 33-50.
- Manganelli S., R.F. Engle. 2004. A comparison of value-at-risk models in finance. Szegö G., ed., *Risk Measures for the 21st Century*, Wiley: Chichester.
- Moder, J.J, E.G. Rogers. 1968. Judgement estimates of the moments of PERT type distributions. *Management Science* **15** B76-B83.
- Nelson, D.B. 1990. Stationarity and persistence in the GARCH(1,1) model. *Econometric Theory* **6** 318-334.
- Nelson, D.B. 1991. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* **59** 347-370.
- Parkinson, M. 1980. The extreme value method for estimating the variance of the rate of return. *Journal of Business* **53** 61-65.
- Pearson, E.S., J.W. Tukey. 1965. Approximate means and standard deviations based on distances between percentage points of frequency curves. *Biometrika* **52** 533-546.
- Poon, S., C.J.W. Granger. 2003. Forecasting volatility in financial markets: A review. *Journal of Economic Literature* **41** 478-639.
- RiskMetrics*. 1996. *Technical Document*. Morgan Guaranty Trust Company of New York.
- White, H. 1994. *Estimation, Inference and Specification Analysis*, Cambridge University Press, Cambridge.

Table 1 Parameters, α_1 and β_1 , in Expression (10), from LS Regression of ε_i^2 on the Interval Between Symmetric Quantiles Estimated by Asymmetric Slope CAViaR Model for In-Sample Stock Index Data. Standard Errors in Parentheses.

Interval	Parameter	CAC40	DAX30	FTSE100	NIKKEI225	S&P500	Pearson and Tukey Values
0.98	$\alpha_1 \times 10^6$	-53.1 (15.0)	-1.77 (11.7)	-13.1 (6.58)	24.0 (16.5)	3.28 (8.03)	0
	β_1	0.0575 (0.0038)	0.0429 (0.0023)	0.0551 (0.0029)	0.0327 (0.0027)	0.0364 (0.0021)	$4.65^{-2} = 0.0462$
0.95	$\alpha_1 \times 10^6$	-11.5 (12.1)	5.78 (11.4)	-5.99 (6.06)	19.6 (16.7)	-4.78 (8.56)	0
	β_1	0.0681 (0.0043)	0.0642 (0.0034)	0.0693 (0.0035)	0.0562 (0.0046)	0.0653 (0.0040)	$3.92^{-2} = 0.0651$
0.90	$\alpha_1 \times 10^6$	-4.89 (11.7)	5.36 (11.3)	7.78 (5.55)	14.1 (17.0)	6.67 (7.96)	0
	β_1	0.0924 (0.0059)	0.0944 (0.0050)	0.0800 (0.0041)	0.0884 (0.0071)	0.0824 (0.0049)	$3.25^{-2} = 0.0947$

Table 2 R^2 Measure of Informational Content for 500 Post-Sample 1 Day-Ahead Variance Forecasts for Stock Indices. R^2 Values are Percentages.

	CAC40	DAX30	FTSE100	NIKKEI225	S&P500	Mean
Moving Average and GARCH Methods						
Simple Moving Average	10.9	8.9	9.0	1.2	5.2	7.0
Exponential Smoothing	12.7	12.9	13.3	1.6	8.9	9.9
GARCH	11.9	12.9	12.6	1.9	9.2	9.7
IGARCH	11.2	12.8	12.3	1.8	9.2	9.5
GJRGARCH	15.0	14.9	16.0	2.8	16.6	13.0
Simplistic Winsorized GJRGARCH	14.0	13.9	15.7	2.7	16.0	12.5
CAViaR Winsorized GJRGARCH	15.0	14.9	16.7	2.8	16.7	13.2
VaR-Based Methods						
Historical Simulation 98% Interval	4.5	5.8	3.1	1.0	0.3	2.9
Historical Simulation 95% Interval	5.6	5.8	2.6	1.0	0.3	3.1
Historical Simulation 90% Interval	3.1	4.3	2.4	0.2	0.5	2.1
BRW 98% Interval	5.0	5.7	4.7	0.6	2.6	3.7
BRW 95% Interval	9.0	7.7	6.8	1.0	4.2	5.8
BRW 90% Interval	8.0	8.8	6.3	0.5	3.7	5.5
Indirect GARCH CAViaR 98% Interval	13.0	12.8	14.4	1.9	10.4	10.5
Indirect GARCH CAViaR 95% Interval	12.3	13.0	14.2	2.2	10.0	10.3
Indirect GARCH CAViaR 90% Interval	13.0	13.0	13.9	2.2	10.1	10.5
Adaptive CAViaR 98% Interval	5.3	6.4	5.5	0.2	3.7	4.2
Adaptive CAViaR 95% Interval	6.8	8.2	4.7	0.1	3.2	4.6
Adaptive CAViaR 90% Interval	8.9	7.8	3.2	0.9	6.3	5.4
Sym Abs Value CAViaR 98% Interval	11.8	12.5	12.3	1.8	9.5	9.6
Sym Abs Value CAViaR 95% Interval	11.7	12.3	13.0	2.0	9.0	9.6
Sym Abs Value CAViaR 90% Interval	12.0	12.3	10.3	1.9	8.5	9.0
Asym Slope CAViaR 98% Interval	14.5	14.0	15.1	2.9	17.4	12.8
Asym Slope CAViaR 95% Interval	15.4	15.5	16.4	3.0	18.7	13.8
Asym Slope CAViaR 90% Interval	15.4	15.4	16.3	3.1	19.5	13.9

Table 3 R² Measure of Informational Content for 500 Post-Sample Variance Forecasts for 10-Day Holding Period for Stock Indices. R² Values are Percentages.

	CAC40	DAX30	FTSE100	NIKKEI225	S&P500	Mean
Moving Average and GARCH Methods						
Simple Moving Average	34.7	25.9	19.1	3.2	12.8	19.1
Exponential Smoothing	43.8	39.5	29.0	4.9	23.1	28.1
GARCH	39.3	39.3	27.3	5.9	23.8	27.1
IGARCH	36.5	38.9	26.4	5.5	23.8	26.2
GJRGARCH	52.2	46.1	34.6	9.1	44.8	37.4
Simplistic Winsorized GJRGARCH	47.3	41.3	33.9	9.1	43.3	35.0
CAViaR Winsorized GJRGARCH	52.4	46.6	36.5	9.3	45.4	38.0
VaR-Based Methods						
Historical Simulation 98% Interval	14.7	19.2	6.9	5.1	0.2	9.2
Historical Simulation 95% Interval	18.6	18.8	5.7	4.9	0.3	9.6
Historical Simulation 90% Interval	11.3	14.8	4.8	1.3	0.7	6.6
BRW 98% Interval	13.4	17.4	9.4	0.7	6.6	9.5
BRW 95% Interval	29.1	22.2	12.7	0.6	11.2	15.2
BRW 90% Interval	24.5	24.3	12.4	2.3	8.5	14.4
Indirect GARCH CAViaR 98% Interval	48.2	41.8	32.2	6.0	27.2	31.1
Indirect GARCH CAViaR 95% Interval	41.4	40.0	31.1	7.5	26.4	29.3
Indirect GARCH CAViaR 90% Interval	46.3	40.2	30.5	7.2	26.5	30.2
Adaptive CAViaR 98% Interval	17.2	18.8	11.7	0.3	10.3	11.7
Adaptive CAViaR 95% Interval	21.3	24.7	9.6	0.5	8.6	13.0
Adaptive CAViaR 90% Interval	28.2	23.6	6.2	3.3	16.4	15.5
Sym Abs Value CAViaR 98% Interval	40.9	42.4	28.5	7.4	25.2	28.9
Sym Abs Value CAViaR 95% Interval	40.6	39.5	30.1	9.0	24.0	28.6
Sym Abs Value CAViaR 90% Interval	42.3	39.5	22.9	8.6	22.4	27.2
Asym Slope CAViaR 98% Interval	52.5	48.6	35.0	13.4	48.9	39.7
Asym Slope CAViaR 95% Interval	57.1	50.9	37.6	14.1	50.8	42.1
Asym Slope CAViaR 90% Interval	56.1	51.4	37.8	14.8	52.9	42.6

Table 4 R^2 Measure of Informational Content for 500 Post-Sample Variance Forecasts for 20-Day Holding Period for Stock Indices. R^2 Values are Percentages.

	CAC40	DAX30	FTSE100	NIKKEI225	S&P500	Mean
Moving Average and GARCH Methods						
Simple Moving Average	28.7	23.2	14.2	2.2	11.3	15.9
Exponential Smoothing	36.0	32.6	21.4	4.5	18.9	22.7
GARCH	32.1	32.5	20.1	5.7	19.3	21.9
IGARCH	29.6	32.3	19.5	5.2	19.3	21.2
GJRGARCH	44.9	37.9	26.1	10.7	37.5	31.4
Simplistic Winsorized GJRGARCH	40.1	35.7	25.6	10.7	36.2	29.7
CAViaR Winsorized GJRGARCH	45.0	38.2	27.6	11.0	37.9	31.9
VaR-Based Methods						
Historical Simulation 98% Interval	12.6	19.9	5.6	7.7	0.3	9.2
Historical Simulation 95% Interval	16.0	17.8	4.6	7.0	0.0	9.1
Historical Simulation 90% Interval	10.0	14.1	3.1	1.5	0.1	5.8
BRW 98% Interval	9.0	15.6	6.7	0.6	4.1	7.2
BRW 95% Interval	21.3	20.4	8.6	0.9	8.6	11.9
BRW 90% Interval	18.8	22.9	8.5	2.6	7.0	12.0
Indirect GARCH CAViaR 98% Interval	40.0	33.4	23.6	5.9	20.9	24.8
Indirect GARCH CAViaR 95% Interval	33.9	32.7	23.0	7.6	20.6	23.6
Indirect GARCH CAViaR 90% Interval	38.2	32.9	22.5	7.5	20.7	24.4
Adaptive CAViaR 98% Interval	14.1	16.8	8.8	0.5	7.8	9.6
Adaptive CAViaR 95% Interval	17.9	23.7	6.8	1.6	6.1	11.2
Adaptive CAViaR 90% Interval	24.2	21.5	3.8	4.0	13.6	13.4
Sym Abs Value CAViaR 98% Interval	34.6	35.7	21.4	7.7	19.3	23.7
Sym Abs Value CAViaR 95% Interval	34.4	34.0	22.7	10.1	18.5	23.9
Sym Abs Value CAViaR 90% Interval	36.0	34.2	17.1	9.4	17.7	22.9
Asym Slope CAViaR 98% Interval	45.9	41.3	27.0	16.0	38.2	33.7
Asym Slope CAViaR 95% Interval	50.4	44.0	29.4	17.0	40.9	36.3
Asym Slope CAViaR 90% Interval	49.4	44.1	29.2	18.1	41.7	36.5

Table 5 Results of the Encompassing Test, $\sigma_{Rt,k}^2 = w\hat{\sigma}_{Ct,k}^2 + (1-w)\hat{\sigma}_{Gt,k}^2 + e_t$, for Stock Index Data. $\sigma_{Rt,k}^2$ is Realised Multiperiod Variance; $\hat{\sigma}_{Ct,k}^2$ is the VaR-Based Variance Forecast Using 90% Interval Estimated by the Asymmetric Slope CAViaR Model; $\hat{\sigma}_{Gt,k}^2$ is the GJRGARCH Variance Forecast; and e_t is a Residual Term.

	CAC40	DAX30	FTSE100	NIKKEI225	S&P500
1 step-ahead					
\hat{w}	1.36	0.51	0.70	1.13	1.20
P-value for $H_0: w = 1 \quad H_1: w < 1$	0.80	0.10	0.34	0.60	0.77
P-value for $H_0: w = 0 \quad H_1: w > 0$	0.00	0.09	0.17	0.01	0.00
10-day holding period					
\hat{w}	1.45	0.55	2.61	2.42	1.04
P-value for $H_0: w = 1 \quad H_1: w < 1$	0.80	0.16	0.87	0.99	0.55
P-value for $H_0: w = 0 \quad H_1: w > 0$	0.00	0.11	0.03	0.00	0.00
20-day holding period					
\hat{w}	2.55	0.22	1.23	3.06	0.61
P-value for $H_0: w = 1 \quad H_1: w < 1$	0.93	0.04	0.59	1.00	0.12
P-value for $H_0: w = 0 \quad H_1: w > 0$	0.01	0.31	0.12	0.00	0.03

Note. Test uses multi-period forecasts for non-overlapping holding periods in the post-sample period.

Table 6 Hit Percentage, Dynamic Quantile Test Statistic and QR Sum for 1 Day-Ahead Forecasting of the 0.01 and 0.99 Quantiles for the Stock Index Data Using GJRGARCH and Asymmetric Slope CAViaR. QR Sum Values Have Been Multiplied by 10^5 .

	CAC40		DAX30		FTSE100		NIKKEI225		S&P500	
	0.01	0.99	0.01	0.99	0.01	0.99	0.01	0.99	0.01	0.99
Hit %										
GJRGARCH	0.6	99.4	0.2	99.8	0.8	99.8	0.2	99.8	0.4	100.0
Asym Slope CAViaR	0.6	97.8	0.8	98.2	1.8	98.6	0.2	99.6	1.8	99.4
DQ										
GJRGARCH	0.9	38.5**	3.2	3.3	0.7	3.2	3.7	3.3	2.2	N/A
Asym Slope CAViaR	33.8**	29.2**	0.8	5.2	46.6**	13.9	4.1	2.1	19.9**	1.3
QR Sum										
GJRGARCH	54	53	61	62	46	44	50	55	45	45
Asym Slope CAViaR	54	57	57	61	46	39	42	51	47	37

Note. Significance at 5% and 1% levels is indicated by * and **, respectively. Tests were performed on DQ but not Hit % because sample size is not sufficiently large. N/A indicates results not available because DQ test regression could not be performed due to dependent variable, Hit_t , being identical for all 500 post-sample periods.

Table 7 Hit Percentage, Dynamic Quantile Test Statistic and QR Sum for 1 Day-Ahead Forecasting of the 0.05 and 0.95 Quantiles for the Stock Index Data Using GJRGARCH and Asymmetric Slope CAViaR. QR Sum Values Have Been Multiplied by 10^5 .

	CAC40		DAX30		FTSE100		NIKKEI225		S&P500	
	0.05	0.95	0.05	0.95	0.05	0.95	0.05	0.95	0.05	0.95
Hit %										
GJRGARCH	5.6	96.6	5.4	96.8	6.0	97.6**	2.4**	97.8**	3.0*	97.6**
Asym Slope CAViaR	6.2	96.0	8.2**	93.2	6.2	97.2*	5.2	94.4	4.6	95.6
DQ										
GJRGARCH	3.5	6.9	7.3	5.6	11.9	8.5	9.1	10.8	7.1	11.2
Asym Slope CAViaR	4.8	5.7	15.9*	13.0	11.2	7.5	4.7	4.4	6.9	6.2
QR Sum										
GJRGARCH	148	153	165	152	210	230	188	202	166	178
Asym Slope CAViaR	142	144	162	147	213	226	184	202	161	170

Note. Significance at 5% and 1% levels is indicated by * and **, respectively. Tests performed on DQ and Hit %.

Table 8 Skewness and Excess Kurtosis for the 20 Individual S&P500 Stocks and the Five Stock Indices.

	Skewness	Excess Kurtosis
Individual Stocks		
General Electric	0.03	3.55**
Microsoft	-0.09	4.13**
Exxon Mobil	0.11*	3.05**
Pfizer	-0.13**	1.85**
Citigroup	0.04	4.00**
Wal Mart Stores	0.12*	1.96**
Intel	-0.40**	5.22**
American International Group	0.17**	2.83**
Cisco Systems	-0.05	3.89**
IBM	0.13**	5.56**
Johnson and Johnson	-0.45**	6.71**
Procter and Gamble	-3.47**	69.80**
Coca Cola	-0.10*	3.45**
Bank of America	-0.11*	2.26**
Altria Group	-0.34**	6.75**
Merck and Co.	-0.02	2.48**
Wells Fargo and Co.	0.11*	1.99**
Verizon Comms.	0.14**	3.64**
Chevron Texaco Co.	0.07	1.61**
Dell	-0.26**	2.90**
Stock Indices		
CAC40	-0.06	2.36**
DAX30	-0.23**	3.14**
FTSE100	-0.17**	2.76**
NIKKEI225	0.10*	2.45**
S&P500	-0.11*	3.69**

Note. Significance at 5% and 1% levels is indicated by * and **, respectively.

Table 9 Mean of the R^2 Measure of Informational Content for 500 Post-Sample Variance Forecasts for 1-, 10- and 20-Day Holding Periods for the 20 Individual Stocks.

	R^2 Mean		
	1-Day	10-Day	20-Day
Moving Average and GARCH Methods			
Simple Moving Average	2.3	6.1	5.5
Exponential Smoothing	2.3	5.4	4.7
GARCH	3.4	7.4	6.4
IGARCH	2.9	7.3	6.5
GJRGARCH	5.0	12.8	11.9
Simplistic Winsorized GJRGARCH	4.9	12.6	11.6
CAViaR Winsorized GJRGARCH	5.6	14.2	13.1
VaR-Based Methods			
Historical Simulation 98% Interval	0.6	3.3	5.9
Historical Simulation 95% Interval	0.5	2.8	5.0
Historical Simulation 90% Interval	0.5	3.1	5.7
BRW 98% Interval	0.7	2.6	3.7
BRW 95% Interval	1.1	3.5	4.4
BRW 90% Interval	0.8	2.6	3.3
Indirect GARCH CAViaR 98% Interval	4.2	9.5	7.7
Indirect GARCH CAViaR 95% Interval	4.1	9.7	8.0
Indirect GARCH CAViaR 90% Interval	4.1	9.8	8.2
Adaptive CAViaR 98% Interval	0.6	2.7	4.3
Adaptive CAViaR 95% Interval	0.7	3.1	4.6
Adaptive CAViaR 90% Interval	0.9	3.7	4.8
Sym Abs Value CAViaR 98% Interval	4.2	10.6	9.2
Sym Abs Value CAViaR 95% Interval	3.9	10.2	9.0
Sym Abs Value CAViaR 90% Interval	3.8	9.9	8.8
Asym Slope CAViaR 98% Interval	6.5	17.8	16.3
Asym Slope CAViaR 95% Interval	6.2	18.3	17.7
Asym Slope CAViaR 90% Interval	6.3	18.1	17.3

Figure 1 News Impact Curves for the VaR-Based Volatility Forecasting Method Using 90% Intervals Estimated by CAViaR Models. Parameters Estimated Using In-Sample Period of the S&P500 Stock Index Returns.

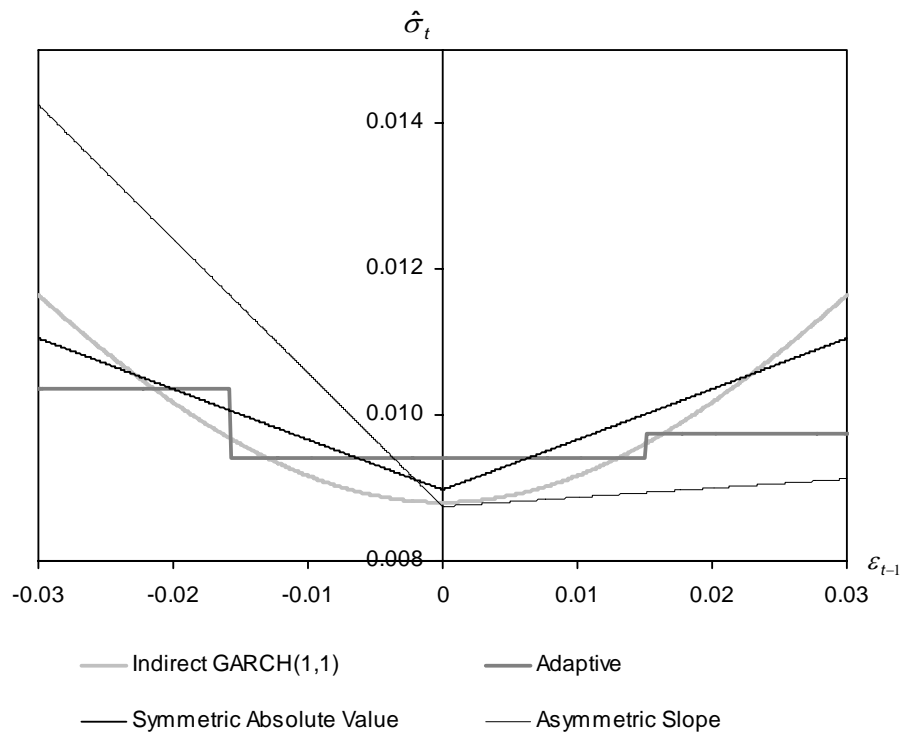


Figure 2 S&P500 Daily Stock Index Returns Plotted on Primary Y-Axis for 29 April 1993 to 28 April 2003. Plotted on Secondary Y-Axis are One Day-Ahead Volatility Forecasts from Two Implementations of the New VaR-Based Volatility Forecasting Method. The First Used the 90% Interval Between Symmetric Quantiles Estimated by the Adaptive CAViaR Model and the Second Used the Asymmetric Slope CAViaR Model. Forecasts for Final 500 Periods are Post-Sample Predictions.

