The Forecasting Performance of Implied Volatilities on Individual Equity Options

John Garvey
The University of Limerick

May 9, 2006

Abstract

For equity fund managers the attribution and control of risk is a central function requiring appropriate volatility forecasts on individual assets over a number of forecast horizons. This paper assesses the predictive ability of artificial implied volatility estimates against forecasts provided by GARCH and E-GARCH for a range of FTSE-100 stocks. From the evidence shown here, implied volatility forecasts are superior to alternative methodologies over both short- and medium-term forecast horizons. This finding supports the use of individual equity options data in formulating an optimal rebalancing strategy for equity portfolios.

1 Introduction

Generating accurate estimates of the future asset volatility has been an area of some success in finance theory. The academic progress in this area has not easily translated into standard industry practice despite the obvious need for good forecasts with model complexity often cited as a central reason. Volatility forecasting literature is dominated by two principal methodologies the most common of which is a time-series approach that utilises historical returns pattern to generate forecasts of return volatility. The rate of innovation that has characterised time-series forecasting has led to the development of often successful but often complex forecast models that are currently applied in only a limited context within the financial markets. These models are essentially backward-looking using previous return patterns and projecting this behaviour forward. An alternative though less widely studied forecasting method extracts expectations of volatility from observed option prices. The forward-looking perspective of option traders implies that observed prices should provide some information on the returns behaviour of the underlying asset, this is limited by the efficiency of the options market and the ability of option pricing models to adequately capture trader's expectations. The relative ease with which implied volatilities can be backed-out of observed option prices is an attractive feature for some market preactitioners, such as portfolio managers who have a requirement for easily available and accurate volatility forecasts for individual assets over short and medium term horizons.

This paper provides a comprehensive examination of individual equity options quoted on the London International Financial Futures Exchange (LIFFE) and compares their predictive ability to the most commonly cited time-series approaches, namely GARCH (1,1) and E-GARCH (1,1). The use of a sample of individual equity options represents a signifcant extension of previous work in the area which has primarily focused on index options. The majority of studies examining the forecast performance of implied volatility have focused on index options where issues around liquidity and data availability are less of a concern. The Chicago Board of Trade Volatility Index (VIX) is a weighted average of 8 OEX (S&P 100 Index) put and call options available since 1993 and its application as a forecasting instrument by market participants and is a good example of how this data has been used to provide an indication of future volatility at the market level. At the level of individual assets models such as GARCH (1,1) and its extensions dominate the forecasting literature and have proved to be relatively successful in a univariate context. The increase over the last ten years in the number of securities on which options trade now means that there is an extensive source of information in the options market for portfolio managers that is an attractive alternative to these using GARCH-based approaches.

The difference between forecasts provided by both methodologies can be understood in the context of information arriving to the market as either expected announcements or unexpected news. A specific volatility pattern is associated with each type of event. For many securities volatility often declines in advance of expected announcements, such as an earnings report or the result of a corporate legal action, as traders and institutional investors adopt a 'wait and see' approach. In the trading period immediately following the announcement volatility will usually spike with a more pronounced rise in volatility in the case of negative news. In both cases volatility will then gradually return to some normal level as market participants reassess the value of the security. Forecasting approaches based on historical volatility fail to adequately capture this dynamic. The decline in volatility prior to an announcement is included in near term forecasts and should a short-term volatility jump occur on the announcement date it will remain dominant in the short-term forecasts leading to an overestimation of actual volatility in subsequent time-periods. Despite the limitations of the underlying assumptions of standard pricing models, the pricing rationale of traders should intuitively recognise the difference between short-term rises in volatility levels due to company specific announcements and more general increases in volatility that reflect fundamental conditions affecting stock volatility. In the case of individual options the challenge is to exploit this potentially valuable information content in a portfolio setting.

The forward-looking expectations reflected in option prices has potentially strong applications in some specific financial decision-making, such as those taken by institutional equity investors. In actively managed equity porfolios a set of assets is selected based on the manager's preference. The timing and extent of the rebalancing decision over the holding period is generally undertaken in an ad hoc manner, either calendar-based (rebalancing may occur monthly or quarterly) or volatility-based. The development of an optimal rebalancing strategy where the cost of transacting is a primary input requires the fund manager to have access to the most appropriate volatility forecasts for individual assets over a number of forecast horizons.

As mentioned earlier, the majority of previous studies on implied volatility have focused on highly liquid datasets of index options. The objective of this analysis is to provide a comprehensive analysis of the pricing of a range of individual equity options with a view to extracting any forward-looking information on future idiosyncratic or stock-specific volatility patterns. The predictive ability of options on FTSE-100 stocks are tested over 1-, 5-, 10- and 22-day forecast horizons and are ranked against popular statistical forecasting methods. The next section traces the development of the literature on volatility forecasting and is followed by a detailed analysis of the forecasts provided respectively by the time-series and the implied volatility approaches. The testing procedure can be separated into two parts, a preliminary analysis carried out over an extended sample period ranks forecasts using mean absolute error (MAE) and root mean square error (RMSE) statistics. In this instance squared returns are used as the proxy for actual volatility proxy, limiting the tests to a simple ranking procedure. The use of intra-day (five-minute) returns data facilitates a more comprehensive analysis of each of the forecasting methods. The information content of volatility forecasts can be tested more extensively through the running of basic and encompassing regressions. In addition, test statistics developed by Diebold and Mariano (1995) [8] and Harvey-Leybourne-Newbold (1998) [**?**, **?**] provide a pairwise comparison of each of the forecasting models.

2 Forecasting Volatility

The evolution of time-series models can be traced to the use of simple historical averages as an estimate of future volatility. This first generation of models was extended to moving average and exponentially weighted moving averages (EWMA) specifications, such as the popular EWMA specification provided by Riskmetrics^{TM} which were designed to capture volatility clustering. More sophisticated attempts to forecast volatility were developed and can be outlined collectively as ARCH-type models. Pioneered by Engle [11] this approach uses a maximum likelihood procedure to estimate the conditional variance of returns. The $ARCH(q)$ model put forward by Engle [11] generated h_t , the one-step ahead variance forecast, based on q past squared returns. The generalised ARCH ($GARCH(p,q)$) model (Bollerslev [5]) additionally captures the volatility persistence and provides a specification much more suited to modelling financial data. The plain vanilla GARCH (1,1) remains a popular forecasting tool both for its efficiency and the relative ease of implementation. The pattern of asymmetric volatility observed in financial data presented an additional requirement on forecasting models and an early solution is provided by [18] in the form of exponential GARCH (E-GARCH) which captured the negative correlation between stock returns and volatility. This specification now represents only one of a wide variety of sophisticated ARCH-type models.

The primary alternative to this methodology is to back-out the forward-looking expectations of volatility implied in option prices. The development of the Black and Scholes (1973) [16] model marked a significant breakthrough in option pricing. Although the model makes strong assumptions that diverged from reality, such as constant volatility, σ, no transaction costs, divisible securities and no arbitrage, it as a widely popular tool for option pricing. Alternative option pricing models accomodated the possibility of early exercise and were more suited to the pricing of American-style options such as the binomial approach adopted by Cox Ross and Rubinstein (1976) [16], the CRR Model.

2.1 Implied Volatility Forecasting

An analysis undertaken by Lamoureux and Lastrapes (1993) looked at the forecast performance of implied volatilities drawn from ten European-style individual equity options traded on the CBOE for the two-year period from April 1982 to March 1984. Their conclusions indicate that implied volatilities contain information above that available in historical prices for forecasts over a 180- to 90–calendar-day horizon. The findings presented in this paper hint at the potential for implied volatility as the optimal forecasting approach despite the absence of the "correct" equilibrium option pricing model, however the results are not entirely consistent with other contemporary studies.

An analysis of the predictive ability of S&P Index Options by Canina and Figlewski [6] concluded that implied volatilities are of no use when providing forecasts of index volatility explaining the lack of information content as a result of option market inefficiency and the assumptions underlying the Black and Scholes (1973) option-pricing model. Numerous subsequent studies are more in line with the findings of Lamoreux and Lastrapes which depicts a useful but biased forecast of future volatility. The bias discovered in implied volatility forecasts is attributed to the limiting assumption of constant volatility required by standard option pricing models. The constant estimate of volatility required by the Black and Scholes (1973) formula is not a close approximation of reality and therefore results in a biased estimate of asset volatility over the remaining life of the option. In reality the stochastic nature of stock return volatilities, misspecification of the terminal stock price distribution and the presence of early exercise possibilities may all have an impact on the forecast bias of implied volatilities.

Figlewski [9] provides a number of reasons for the weak results from tests on implied volatiltiy carried out in those early studies. One of the key reasons is that the equilibrium option price is difficult to observe due to bid-ask spreads. In addition there are issues resulting from non-continuous trading.

He states that these market frictions can be exaggerated when using deep-ITM and deep-OTM options where even a narrow bid-ask spread has a significant impact on the implied volatility estimate. Another issue surrounding the use of option prices to forecast volatility is the clientele effect where groups demand options for specific reasons. One of the commonly cited examples of this type of 'clinteleism' is that of fund managers who purchase put options to protect against downside risk in their portfolio. The demand for specific types of options can impact on the market price and result in a bias in forecasts derived from observed market prices.

Feinstein (1989) and later Corrado and Miller (1994) both provide positive results on implied volatilities. They find that the Black-Scholes option-pricing model can recover virtually unbiased stock return volatility estimates when volatility behaves stochastically identifying the greater sensitivity of at-the-money options to volatility of underlying stocks than options that are away-from-the-money. In addition they show that the apparent bias in implied volatilities that result from misspecification of stock price dynamics can also be minimized if implied volatilities are obtained from at-the-money options. Their work also suggests that implied volatility estimates are nearly indistinguishable across option pricing models for at-the-money options. In a similar finding Bodie and Merton (1995) find that the bias associated with implied volatility forecasts can be mitigated by using short-dated options that are close-to or at-the-money. Using S&P 500 options for the period 1985-1987 Bates [4] examined the relative prices of out-of-the-money (OTM) put options and out-of-the-money (OTM) call options in the expectation that "unusually" expensive OTM puts are indicative of a market assessment of an imminent downturn in the market. In the case of the October 1987 stock market crash he found expectations of a market downturn were reflected in options prices in the year before the crash, however in the period immediately preceding the crash downside risk was not very pronounced.

Christensen and Prabhala [7] - henceforth CP - looked at S&P index option prices in the broader context of their predictive ability for the volatility of the underlying S & P Stock Index. Their results contradicted earlier findings put forward by Canina and Figlewski [6] and they discount their conculsion of option market inefficiency. They also note that a regime shift occurred around the time of the October 1987 crash, explaining the bias that was inherent in option prices in the time immediately preceding the crash. The results achieved by CP affirm the superiority of implied volatilities over past volatility in forecasting the volatility of the S&P 100. A significant conclusion of their study was that previous studies were weakened by their use of overlapping sampling method. Their analysis represented a step forward in the analysis of the relationship between implied and realised volatility. They use a longer sample period of $11\frac{1}{2}$ years, significantly longer than previous studies. In addition the volatility series are constructed with non-overlapping data, providing more reliable regression estimates that would otherwise be distorted using an overlapping dataset. Empirical studies that assess the informational content of implied volatility by regressing ex post realized volatility on ex ante implied volatility have generally found a bias in forecasts provided.

This study differs from those of Lamoureux and Lastrapes [17] and Christensen and Prabhala [7] where implied volatility is backed out of option data based on a single strike price. This is traditionally attributed to wrong choice of option pricing model, data may not be measured correctly and statistical problems arise because the sample periods is too short A recent addition to the literature provided by Penttinen [19] looks at the reasons why implied volatility slightly exceeds realised volatility most of the time and is lower than realised volatility during periods of relatively high market volatility.

Pentinnen [19] argues that the bias between ex-ante implied and ex-post realised volatility results from unrealised expectations of infrequently occurring jumps in volatility. His analysis supports the idea that the option pricing process is an example of rational behaviour on the part of options traders whose expectations remained unrealised for a period of time. In this study the sample period provides an opportunity to investigate the comparative performance of implied volatility over an extended time period that can be divided into approximately equal time periods chatacterised by low volatility and prolonged heightened volatility respectively. A review of nineteen papers on this topic by Poon and Granger [20] finds that despite model weaknesses and issues around option market efficiency, implied volatility performed well and in some cases performed better than models based on historical data.

3 Methodology

Options data on individual FTSE-100 equities was purchased from the London International Financial Futures Exchange (LIFFE). Sixteen individual equity options were selected wit a continuous sample period available from September 1997 to December 2003. The preliminary analyses uses end-of-day price data on the underlying assets provided by Datastream to generate squared returns. Intra-day data purcased from data provider Olsen was used to calculate realised volatility. This section outlines the estimation process for daily implied volatility (IV) used as a forecast and sets out the testing procedure used to IV forecasts against forecasts provided by GARCH (1,1) and E-GARCH (1,1).

3.1 Measuring Volatility

The correct ranking of volatility forecasts requires a consistent proxy for the 'true' volatility process. The difficulty associated with ascertaining the correct proxy is addressed by Andersen and Bollerslev [2] - hereafter AB - who studied the conditional variance provided by a forecasting model against r_t^2. Squared returns are a noisy estimate of volatility and in all cases a regression of the volatility forecast against squared returns would indicate low explanatory power, in fact the R^2 resulting from such an analysis would never exceed $\frac{1}{3}$. In an analysis where the objective is to assess the fraction of volatility explained by a forecasting model this would prove to be a severe limitation, however squared returns are sufficient when ranking the predictive accuracy of various models as in the preminary test here. The availability of intra-day price data allows us to extract more detailed information on volatility forecasts. If $\ln S_t, t = 1, ..., T$ is a series of daily stock prices, and let $\ln S_{t+k\xi}, k = 1, ..., m$ and $\xi = 1/m$ denote a series of five minute observations then a daily estimate of realised volatility can be constructed as $RV_{t,m} = \sum_{k=0}^{m-1} (\ln S_{t+(k+1)\xi} - \ln S_{t+(k)\xi})^2$. If $\ln S_t$ is a continuous semimartingale process, and $\sigma^{\dagger 2}(t)$ is the instantaneous volatility of that process then as $m \longrightarrow \infty (\xi \longrightarrow 0), RV_{t,m} \xrightarrow{pr} \int_{t-1}^{t} \sigma^{\dagger 2}(s) ds$, for $t = 1, 2, ...$where the r.h.s. is the (daily) integrated volatility. As the time interval goes towards zero, realised volatility provides a consistent estimator of the integrated volatility process and extends range of comparisons between forecasting models.

However as stated above, Awartani and Corradi [3] show that even if the true unobservable volatility, $\sigma_t^{\dagger 2}$, is replaced with squared returns, r_t^2, then the correct ranking of models based on any quadratic loss function will be maintained and so realised volatility does not offer significant benefit over squared returns in this context.

3.2 Deriving Implied Volatilities from Individual Equity Options

LIFFE provide data on American-style equity options priced using the Cox-Ross-Rubinstein - henceforth CRR - binomial model with each option having rights on 1000 shares. The CRR framework splits the time to expiry into sub periods in which the share value can rise or fall by a known proportion resulting in a range of final possible share price outcomes. Associated with each of the final price outcomes is an associate option price and the option value is obtained by working backwards over the number of periods from expiry, thus allowing for early exercise and the payment of dividends over the life of the option. The opportunity of early exercise means that American option prices will trade at a mark-up over European option prices and this premium is dependant on the cash flows of the underlying asset. When used to price an option this the necessary inputs are the current stock price, the exercise price, the risk-free interest rate and an estimate of the stocks volatility. All of the factors are known except for the volatility estimate, therefore the forward-looking expectation of volatility is implied in the observed trading price of options.

For each trading day, equity options have maturities corresponding to the two near-term months plus two additional months from the January, February or March quarterly cycles and from which a synthetic implied volatility for each trading day can be constructed using a simple weighted average approach based on the option's 'moneyness'[1]. Ederington and Guan [12] show that the approach of commercial providers, who use just a few at-the-money or close-to-the-money options provides a marginally better forecast than estimates of implied volatilities that include away from the money options. In this case the options dataset is subject to a daily ranking procedure that yields the two nearest-to-the-money calls and the two nearest-to-the-money puts for each stock option included in the analysis, thus excluding options trading far from the money. A weighted average of these four options is then calculated, thus giving greater emphasis implied volatility drawn from at- or close-to-the-money options which have been shown to have greater sensitivity to movements in the underlying asset. The weighting scheme can be written as

$$\begin{aligned} IV_{t,T} &= 0.25\left(\frac{X_{c2}-F}{X_{c2}-X_{c1}}\right).IV_{c1} + 0.25\left(1-\frac{X_{c2}-F}{X_{c2}-X_{c1}}\right).IV_{c2} \\ &\quad +0.25\left(\frac{X_{p2}-F}{X_{p2}-X_{p1}}\right).IV_{p1} + 0.25\left(1-\frac{X_{p2}-F}{X_{p2}-X_{p1}}\right).IV_{p2} \qquad (1) \end{aligned}$$

where F is the underlying stock price (face value), X_{c1} and X_{c2}($X_{p1}$and X_{p2}) are the strike prices of the closest-to-the-money call (put) options and IV_{c1} and IV_{c2} ($IV_{p1}$and IV_{p2}) corresponding implied volatilities. The implied volatilities for the options at various strike prices are made available by LIFFE and the Cox-Ross-Rubinstein model is used to estimate the volatility implied by a quoted option prices. This is done by inverting the model and using the option mid prices as an input.

There are infrequent cases of gaps in the dataset. In these circumstances implied volatility for that day is taken to be the average implied volatility over the entire sample. At every point in time across the sample period the model provides a one-step ahead forecast of volatility. Volatility forecasts for longer horizons are estimated by applying the "square root of time rule" is applied. In the case of the implied volatility forecast, the n day volatility forecast is given by:

$$IV_{n,t} = \sqrt{\frac{n}{360}} IV_t.$$

where $IV_{n,t}$ is the expected volatility over the $[t+1, t+n]$ period.

[1]The 'moneyness' of an option refers to the distance between the strike price, the price at which the option can be exercised, and the current price on the underlying asset.

3.3 Statistical Approaches to Volatility Forecasting

As mentioned in the previous section the development of methodologies that can capture the volatility pattern of an individual asset over time and use it to forecast future volatility have been the relatively successful in volatility forecasting. In each of the following cases $r_n = \ln(P_n) - \ln(P_{n-1})$ is the daily return on a FTSE 100 stock.

One of the most straightforward methods used to model volatility in financial data is the exponentially weighted moving average (EWMA) approach used by Riskmetrics. The estimate of volatility is constructed as the exponentially weighted moving average of squared returns. It can be written as

$$\sigma_n^2 = \lambda\sigma_{n-1}^2 + (1-\lambda)u_{n-1}^2 \tag{2}$$

The volatility estimate, σ_n for day n is calculated from σ_{n-1} available at the of the previous day, while $\sigma_{n-1}$was the volatility estimate for $n-1$ made at the end of day $n-2$, $u_{n-1}^2$is the the most recently observed market return for the asset. This methodology has the advantage of being easily implemented and provides a simple measure for tracking volatility, however it has been superceded my models that capture additional volatility patterns. The ARCH(p) process developed by Engle [11] was the earliest model to successfully model the conditional heteroskedasticity of financial returns by assuming the conditional variance is a weighted average of the squared average of up to p previous squared unexpected returns. The GARCH model developed by Bollerslev [5] is a generalization of this model and has proved relatively successful in capturing additional volatility dynamics. The GARCH(p,q) model includes q autoregressive terms in addition to the ARCH(p). The simple GARCH (1,1) model has become the most widely used version due to its simplicity and efficacy in providing short-term forecasts. It can be written as follows:

$$\sigma_n^2 = \gamma V_L + \alpha u_{n-1}^2 + \beta\sigma_{n-1}^2 \tag{3}$$

where γ is the weight assigned to V_L , α is the weight assigned to u_{n-1}^2,and β is the weight assigned to σ_{n-1}^2.The weights must sum to one, so that, $\gamma+\alpha+\beta = 1$. The model specification in this case indicates that σ_n^2 is based on the most recent observation of u^2 and the most recente estimate of the variance. The model used to estimate the parameters can be written as

$$\sigma_n^2 = \varpi + \alpha u_{n-1}^2 + \beta\sigma_{n-1}^2 \tag{4}$$

where $\varpi = pV_L$.Once ϖ, α and β are known then the long term variance V_L can be calculated as ϖ/γ. As with the EWMA model the "decay rate", β in this case has a significant impact on how the future volatiltiy estimate utilises past observations. Maximum likelihood is used to estimate the appropriate parameters for the model, where the underlying distribution is a student-t form that takes account of the fat-tailed or kurtotic distribution pattern common to equity returns. The parameters are re-estimated on a rolling basis. The GARCH (1,1) model does not capture the asymmetric volatility observed in financial data, but despite this it remains one of the most popular GARCH specifications. GARCH specifications. Nelson [18] developed the first asymetric GARCH model, known as exponential or E-GARCH. The conditional variance equation in the E-GARCH model introduces leverage terms as follows:

$$\ln \sigma_n^2 = \gamma + g(z_{n-1}) + \beta \ln \sigma_{n-1}^2 \tag{5}$$

where $g(\cdot)$ is an asymmetric response function defined by

$$g(z_n) = \lambda z_n + \varphi(|z_n| - \sqrt{2/\pi}). \tag{6}$$

The standard normal variable z_n is the standardised unexpected return ε_n/σ_n. When $\varphi > 0$, and $\lambda < 0$ negative shocks to returns ($z_{n-1} < 0$) induce larger conditional variance responses than positive shocks. One of the advantages of GARCH models is that a single model can provide volatility forecasts for a number of maturities. Forecasts generated from GARCH models will mean-revert to the a long-term level volatility level that is determined by the GARCH parameters. Significant differences can occur in the short-term forecasts provided by symmetric and asymmetric GARCH. The GARCH forecast over h-period horizon is the sum of the instantaneous GARCH forecast variances, plus double the sum of the forecast autocovariances between returns. However, the second part of the equation will be very small compared to the first part. The conditional mean equation is generally a constant so that the double sum is zero. Eliminating this part of the equation means that h-day forecasts can be estimated by adding the j-step-ahead GARCH variance forecasts. These are then square rooted and annualised to give GARCH h-day volatility forecasts. It can be written as

$$\sigma_{n,t}^2 = \sqrt{\sum_{k=1}^{h} g_{t+k}} \tag{7}$$

where $\sigma_{n,t}^2$ is the expected volatility over the period [t+1, t+n] according to the GARCH model.

4 Empirical Results

The preliminary test in this analysis ranks forecasts provided by alternative models based on the out-of-sample mean absolute error (MAE) and the root mean square error (RMSE). The implied volatility forecast is ranked against forecasts provided by the GARCH (1,1) and E-GARCH (1,1) model for one-day and 5-, 10- and 22-day forecasting horizon. As mentioned above squared returns provide a sufficient proxy of 'true' for the simple ranking procedure. Given daily returns $r_n = \ln(P_n) - \ln(P_{n-1})$ the forward-looking realised volatility over a time horizon of n days can be estimated for different forecast horizons by calculating the square root of the sum of the squared returns over the n-day period. Squared returns are calculated ex-post for each of the 16 stocks under analysis so that at time t, the squared return, $SR_{n,t}$ for the period [t+1, t+n] is

$$SR_{n,t} = \sqrt{\sum_{j=1}^{n} r_{t+j}^2} \tag{8}$$

As pointed out in Christensen and Prabhala [7] the use of overlapping data results in strongly correlated volatility and leads to problems when used to test volatility forecasting models. To avoid this a non-overlapping dataset is used, in this instance a subset of k times is used, where the sampling periods k are: {1, k, 2k, ...}. The models are tested on a sample of sixteen FTSE-100 stocks and the volatility forecasts provided are tested for the sample period from 1^{st} March 1997 to 31^{st} December 2003 providing 1,630, one-day forecasts, 326, five-day forecasts, 163, ten-day forecasts and 75, 22-day horizon forecasts. Occasional gaps appear in the options data and in such cases the previous day's observation is used. The sample provided belong to a range of industries and would be expected to show different volatility patterns across the sample period.

The first part of the testing procedure ranks each of the forecasting techniques for each asset and across each of the forecast horizons using the mean absolute error (MAE) and the root mean square error (RMSE). The mean absolute error (MAE) is given as

$$MAE = \frac{1}{T}\sum_{t=1}^{T} |SR_{n,t} - Y_{n,t}|$$

where $SR_{n,t}$ is the squared returns for each asset and $Y_{n,t}$ is the forecast of volatility on the individual asset provided by either implied volatility, the GARCH (1,1) model or the E-GARCH (1,1) model. The root mean square error (RMSE) of the volatilty forecast and is given by :

$$RMSE = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(SR_{n,t} - Y_{n,t})^2}$$

and the results provided by this measure are broadly similar to those stated above for the MAE statistic. The average implied volatility one-day ahead MAE and RMSE forecast is 0.0278 and 0.0350 respectively, giving a marginally lower error statistic than that that provided by either the GARCH (1,1)(MAE is 0.0281, RMSE is 0.0351) or E-GARCH (MAE is 0.0281, RMSE is 0.0352) specification. Over longer forecast horizons a broadly similar result can be found.

For the five day forecast horizon the MAE for the implied volatility forecast is 0.535 while the forecast error for the GARCH and E-GARCH models are the same at 0.0537 and this ranking of the forecast approaches is similarly reflected in the results for ten -day and 22-day forecasts. Ranked according the MAE and RMSE using non-overlapping data implied volatilities backed out of individual equity options provide better forecasts than the GARCH and E-GARCH models. Tables 1 to 4 detail the performance of each of the forecast models for each of the sixteen assets over the respective forecasting horizons.

The analysis is extended by substituting daily squared returns as a proxy for actual volatility with realised volatility estimated from intra-day price data [2]. Consistent with the findings of Christensen and Prabhala (1998) [7] non-overlapping data is again used to avoid the problems associated with serial correlation. Using ordinary regression analysis the forecasts provided by the three forecast models can be regressed against the ex-post observed realized volatility over the forward-looking h-day horizon as follows,

$$RV_{h,t} = \beta_0 + \beta_1 f_{\mathrm{mod}\,el,h,t} + e_t, \tag{9}$$

where $RV_{h,t}$ is the realized volatility estimated from five-minute price data, and $f_{\mathrm{mod},h,t}$ is the forecast is provided one of the selected models. In addition to the basic regression analysis specified above an encompassing regression will provide information on the efficiency and level of bias of the forecasts.

$$RV_{h,t} = \beta_0 + \beta_1 f_{\mathrm{mod}\,el,h,t} + \beta_2 RV_{h,t-1} + e_t, \tag{10}$$

An inital test for efficiency can be carried out by estimating a t-statistic on the β_2 coefficient. A H0: $\beta_2 = 0$ that cannot be rejected indicates that the forecast captures all information contained in observed realized volatility, $RV_{h,t}$. A test of the level of bias in the forecast can be carried out using a Fisher test

[2] The second part of the analysis uses the same set of FTSE-100 stocks as used in the preliminary analysis. The exception in this case is Dixons Group, five-minute price data was not available on this stock at the time of writing.

where H0: $\beta_0 = 0$ and $\beta_1 = 1$ cannot be rejected. A third test examines whether the forecast is unbiased and efficient if H0: $\beta_0 = 0$ and $\beta_1 = 1$ and $\beta_2 = 0$ cannot be rejected. This analysis was carried out for the period 1st October 2002 to 31st December 2003 (14 months) for fifteen stocks from the FTSE 100 and a summary of the aggregate results are provided in tables 5 and 6. In almost all cases the R^2 of the basic and encompassing regressions increases as the forecast horizon lengthens. The results also indicate that the implied volatility forecasts for the most part outperform the two GARCH specifications over every forecasting horizon (the only exception to this is a marginal outperformance by E-GARCH (1,1) over the 22-day forecast horizon, shown in the results for the encompassing regression). The slightly negative intercept term, β_0, exhibited by

the implied volatility forecasts indicates an overstimation of volatility consistent with previous findings, however this bias is not excessive. The regression results indicate efficiency in all cases, that is, all information contained in past realised volatiltity is contained in the forecasts as β_2 is for the most part not signficantly different from 0.

To facilitate a pairwise comparison of each of the models we estimated a t-type test associated with Diebold and Mariano (1995) [8] where the differential loss in period i from using model 1 versus model 2 is written as $d_{t+1} = \widehat{u}^2_{1,t+1} - \widehat{u}^2_{2,t+1}$, and $\overline{d} = P^{-1}\sum_t d_{t+1} = MSE_1 - MSE_2$. The Diebold and Mariano (1995) [8] test for equal MSE is formed as

$$\begin{aligned} DM &= \frac{MSE_1 - MSE_2}{\sqrt{\widehat{v}ar(MSE_1 - MSE_2)}} = \frac{\overline{d}}{\sqrt{\widehat{v}ar(\overline{d})}} \\ &= \frac{\overline{d}}{\sqrt{P^{-2}\sum_t (d_{t+1} - \overline{d})^2}} \end{aligned} \quad (11)$$

The DM statistics for both the implied volatility and E-GARCH (1,1) forecasts are compared to a t-statistic with equivalent degrees of freedom for each forecast horizon. An encompassing test using a methodology similar to that of Diebold and Mariano (1995) [8] is suggested by Harvey, et. al (1998) [**?**, **?**]. This latter test is more appropriate as Diebold and Mariano (1995) [8] have shown that their test is oversized over longer forecast horizons (h) and the problem presents itself in come cases when h = 2. The HLN-test is an attempt to mitigate the problem by modifying the methodology developed by Diebold and Mariano (1995).

The test focuses on the covariance between $\widehat{u}_{1,t+1}$ and $\widehat{u}^2_{1,t+1} - \widehat{u}_{1,t+1}\widehat{u}_{2,t+1}$ where $c_{t+1} = \widehat{u}_{1,t+1}(\widehat{u}_{1,t+1} - \widehat{u}_{2,t+1}) = \widehat{u}^2_{1,t+1} - \widehat{u}_{1,t+1}\widehat{u}_{2,t+1}$ and $\overline{c} = P^{-1}\sum_t c_t$. The encompassing test proposed by Harvey et al. (1998) [**?**, **?**] is given as

$$
\begin{aligned}
HLN &= \frac{\overline{c}}{\sqrt{\widehat{v}ar(\overline{c})}} = \frac{\overline{c}}{\sqrt{P^{-1}\sum_t (c_t - \overline{c})^2}} \\
&= P^{\frac{1}{2}} . \frac{P^{-1}\sum_t \widehat{u}_{1,t+1} - P^{-1}\sum_t \widehat{u}_{1,t+1}\widehat{u}_{2,t+1}}{\sqrt{P^{-1}\sum_t \{(\widehat{u}_{1,t+1}^2 - \widehat{u}_{1,t+1}\widehat{u}_{2,t+1}) - \overline{c}\}^2}}
\end{aligned} \tag{12}
$$

Under the null that model 1 (GARCH (1,1)) forecast encompasses model 2 (Implied Volatility or E-GARCH (1,1)), the covariance between $u_{1,t}$ and $u_{1,t} - u_{2,t}$,will be less than or equal to 0, while under the alternative that model 2, contains addtional information, the covariance should be positive. Tables 7 and 8 summarizes the results for the DM-test for equal accuracy and the HLN encompassing test for added information for both the implied volatility forecasts and the E-GARCH (1,1) forecasts where the GARCH (1,1) specification is the benchmark (Model 1) in this instance.

The DM-test is inconclusive providing no evidence of a dominant model, however the HLN test is more informative. The HLN-statistic comparing the implied volatility forecast with GARCH (1,1) is in all cases positive and for the majority of cases signficant additional information is contained in the implied volatility forecast not captured by GARCH (1,1). The average HLN-statistic exceeds the critcal values from the Student's t-distribution with (P-1) degrees of freedom for the 1-, 5- and 10– day forecast horizons and is strongly positive at the 22-day horizon. The positive HLN-statistic for the GARCH versus E-GARCH forecasts shows that the E-GARCH forecast also contains additional information over all forecast horizons, however it is only significant at the one-day horizon. The encompassing test indicates that additional forward-looking information not captured in either GARCH or E-GARCH models, is incorporated into the pricing of individual equity options.

5 Conclusions

This paper is a comprehensive analysis of the most effective volatility forecasting procedures for individual equities trading on the FTSE-100. Using a number of testing procedures we show how synthetic implied volatilities derived from individual equity options provide better forecasts than the popular and often more complex, statistical approaches over different forecast horizons. The forward-looking information contained in traded option prices is compared against GARCH(1,1) and E-GARCH (1,1) over 1-, 5-, 10- and 22-day forecast horizons.

A preliminary ranking procedure using the mean absolute error (MAE) and root mean squared error (RMSE) is carried out using ex-post squared returns as a proxy for actual volatility. In almost all cases, the forecasts provided by implied volatilities perform strongly over all forecast horizons. Intra-day (five-minute) stock prices can be used to estimate a closer approximation of 'true' realised volatility, facilitating a more detailed forecast analysis. A basic regression and an encompassing regression both provide strong overall evidence in favour of using implied volatility despite evidence of a slight forecast bias. Implied volatiltiy forecasts in almost all cases provide the highest R^2, while E-GARCH which allows for asymmetric volatility ranks second-best in terms of predictive accuracy.

The final set of tests carried out are the Diebold-Mariano (DM) test for equal accuracy and the Harvey et al. (HLN) encompassing test. While the results using the DM-test do not suggest a dominant forecasting approach, the positive HLN-statistic across all stocks and forecast horizons indicates that implied volatility and E-GARCH consistently provide additional forward-looking information than that contained in the GARCH (1,1) forecast. As the forecast-horizon increases the HLN-statistic in both cases also increases, however the HLN-statistic for the implied volatility forecasts is the most strongly positive in almost all cases, providing evidence that this is the most appropriate forecast methodology over short and medium term horizons.

Accurate forecasts of stock price volatility are a central decision-making input for a number of market practitioners. For equity fund managers volatility forecasts on individual equities over different forecast horizons can provide a valuable role in the operation of an optimal rebalancing strategy. Although the predictive ability of option implied volatility has been extensively analysed in the context of equity indices less attention has been directed at the role of individual equity options. This paper outlines a simple method for deriving a hypothetical implied volatility forecast from available options and provides evidence that it outperforms sophisticated forecasting procedures over a number of forecast horizons. The speed and simplicity with which implied forecasts can be obtained and the accuracy of their forecasts provides strong motivation for their increased use in formulating a rebalancing strategy for partially diversified equity funds.

References

[1] Alexander, C., 2001, Market models. (Wiley).

[2] Andersen, T. & Bollerslev T., 1998, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, International Economic Review 39, 885-905.

[3] Awartani, B. M. A., & Corradi, V., Predicting the volatility of the S&P-500 stock index via GARCH models: the role of asymmetries, Journal of Forecasting 21, 167-183.

[4] Bates, D.S., 1991, The crash of '87:Was it expected? The evidence from options markets, Journal of Finance 46 (3), 1009-1044.

[5] Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity, Journal of Econometrics 31, 307-327.

[6] Canina, L., & Figlewski, S., 1993, The informational content of implied volatility, Review of Financial Studies 6(3), 659-681.

[7] Christensen, B.J., & Prabhala, N.R., 1998, The relation between implied and realized volatility, Journal of Financial Economics 50, 125-150.

[8] Diebold, F. X. & Mariano, R.S., 1995, Comparing predictive accuracy, Journal of Business & Economic Statistics 20 134-144.

[9] Figlewski, S., 1997, Forecasting Volatility, Financial Markets, Institutions and Instruments Vol. 6 No. 2.

[10] Fleming, J., 1998, The quality of market volatility forecasts implied by S&P 100 index option prices, Journal of Empirical Finance 5, 317-345.

[11] Engle, R. F., 1982, Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation, Econometrica 50,987-1008.

[12] Ederington, L. & Guan, W., 2001, Measuring implied volatility: Is an average better?, Working Paper University of Oklahoma.

[13] Giot, P., 2002, The information content of implied volatility indexes for forecasting volatility and market risk, Working Paper.

[14] Harvey, D.I., Leybourne, S.J. and Newbold, P., 1997, Testing the equality of prediction mean squared errors, International Journal of Forecasting, 13,281-291.

[15] Harvey, D. I., Leybourne, S. J. and Newbold, P., 1998, Tests for Forecast Encompassing, Journal of Business and Economic Statistics, 16(2), 254-59.

[16] Hull, J.C., 2003, Options, Futures, and Other Derivatives, Prentice Hall, 5th Ed.

[17] Lamoureux, C.G., & Lastrapes, WD., 1993, Forecasting stock-return variance:Toward an understanding of stochastic implied volatilities, The Review of Financial Studies 6(2), 293-326.

[18] Nelson, D.B., 1991, Conditional heteroskedasticity in asset returns: A new approach, Econometrica 59 (2), 347-370.

[19] Penttinen, A., 2001, The sensitivity of implied volatility to expectations of jumps in volatility: An explanation for the illusory bias in implied volatility as a forecast, Working Paper.

[20] Poon, SH., & Granger, C.W.J., 2003, Forecasting volatility in financial markets: A review, Journal of Economic Literature XLI, 478-539.

Table 1: Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) for the implied volatility/GARCH (1,1)/E-GARCH(1,1) volatility forecasts against non-overlapping squared returns at the one-day horizon. Results are for 16 FTSE-100 stocks for the 1st March 1997 - 31st December 2003 time period.

	MAE			RMSE		
Stock	Implied Volatility	GARCH(1,1)	E-GARCH(1,1)	Implied Volatility	GARCH (1,1)	E-GARCH(1,1)
Aviva	0.0293	0.0296	0.0296	0.0365	0.0366	0.0367
British Aerospace	0.0329	0.0330	0.0331	0.0412	0.0411	0.0412
British Airways	0.0290	0.0292	0.0291	0.0388	0.0388	0.0388
Cadburys	0.0206	0.0211	0.0211	0.0256	0.0260	0.0260
Diageo	0.0227	0.0231	0.0231	0.0286	0.0289	0.0289
Dixons	0.0314	0.0315	0.0315	0.0405	0.0406	0.0406
GlaxoSmithKline	0.0235	0.0239	0.0239	0.0290	0.0293	0.0293
Hilton Group	0.0284	0.0286	0.0287	0.0361	0.0362	0.0363
Hanson	0.0252	0.0256	0.0256	0.0311	0.0313	0.0313
HSBC	0.0239	0.0242	0.0243	0.0298	0.0300	0.0300
Kingfisher	0.0269	0.0271	0.0271	0.0335	0.0336	0.0336
Lonmin	0.0254	0.0257	0.0258	0.0319	0.0322	0.0322
Marks & Spencers	0.0249	0.0253	0.0253	0.0313	0.0316	0.0316
Prudential	0.0293	0.0296	0.0296	0.0366	0.0368	0.0368
Royal & Sun Alliance	0.0344	0.0346	0.0345	0.0431	0.0431	0.0431
Reuters	0.0367	0.0367	0.0367	0.0462	0.0461	0.0460

Table 2: Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) for the implied volatility/GARCH (1,1)/E-GARCH(1,1) volatility forecasts against non-overlapping squared returns at the five-day horizon. Results are for 16 FTSE-100 stocks for the 1st March 1997 - 31st December 2003 time period.

	MAE			RMSE		
Stock	Implied Volatility	GARCH(1,1)	E-GARCH(1,1)	Implied Volatility	GARCH (1,1)	E-GARCH(1,1)
Aviva	0.0561	0.0564	0.0564	0.0634	0.0636	0.0636
British Aerospace	0.0616	0.0617	0.0008	0.0703	0.0702	0.0703
British Airways	0.0566	0.0569	0.0008	0.0691	0.0692	0.0691
Cadburys	0.0402	0.0406	0.0003	0.0447	0.0451	0.0451
Diageo	0.0431	0.0434	0.0005	0.0491	0.0494	0.0494
Dixons	0.0626	0.0628	0.0008	0.0726	0.0727	0.0727
GlaxoSmithKline	0.0449	0.0453	0.0005	0.0500	0.0503	0.0503
Hilton Group	0.0551	0.0553	0.0007	0.0634	0.0635	0.0635
Hanson	0.0477	0.0480	0.0005	0.0531	0.0534	0.0534
HSBC	0.0460	0.0463	0.0005	0.0516	0.0519	0.0520
Kingfisher	0.0520	0.0522	0.0006	0.0588	0.0589	0.0589
Lonmin	0.0487	0.0491	0.0007	0.0544	0.0547	0.0547
Marks & Spencers	0.0484	0.0487	0.0005	0.0551	0.0554	0.0554
Prudential	0.0560	0.0564	0.0007	0.0635	0.0637	0.0638
Royal & Sun Alliance	0.0655	0.0657	0.0009	0.0750	0.0751	0.0750
Reuters	0.0708	0.0709	0.0010	0.0815	0.0814	0.0814

Table 3: Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) for the implied volatility/GARCH (1,1)/E-GARCH(1,1) volatility forecasts against non-overlapping squared returns at the ten-day horizon. Results are for 16 FTSE-100 stocks for the 1st March 1997 - 31st December 2003 time period.

	MAE			RMSE		
Stock	Implied Volatility	GARCH(1,1)	E-GARCH(1,1)	Implied Volatility	GARCH (1,1)	E-GARCH(1,1)
Aviva	0.0787	0.0790	0.0790	0.0866	0.0868	0.0868
British Aerospace	0.0865	0.0866	0.0867	0.0965	0.0964	0.0965
British Airways	0.0785	0.0786	0.0786	0.0909	0.0910	0.0910
Cadburys	0.0566	0.0570	0.0570	0.0608	0.0612	0.0613
Diageo	0.0609	0.0613	0.0613	0.0675	0.0678	0.0678
Dixons	0.0857	0.0858	0.0858	0.0947	0.0948	0.0948
GlaxoSmithKline	0.0633	0.0636	0.0636	0.0687	0.0690	0.0690
Hilton Group	0.0775	0.0777	0.0777	0.0864	0.0866	0.0866
Hanson	0.0676	0.0679	0.0679	0.0732	0.0735	0.0735
HSBC	0.0641	0.0644	0.0645	0.0698	0.0701	0.0702
Kingfisher	0.0727	0.0729	0.0729	0.0797	0.0798	0.0798
Lonmin	0.0677	0.0680	0.0681	0.0730	0.0733	0.0734
Marks & Spencers	0.0678	0.0681	0.0681	0.0748	0.0751	0.0751
Prudential	0.0778	0.0781	0.0782	0.0863	0.0865	0.0866
Royal & Sun Alliance	0.0901	0.0903	0.0903	0.0997	0.0998	0.0997
Reuters	0.0982	0.0982	0.0982	0.1089	0.1089	0.1088

Table 4: Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) for the implied volatility/GARCH (1,1)/E-GARCH(1,1) volatility forecasts against non-overlapping squared returns at the 22-day horizon. Results are for 16 FTSE-100 stocks for the 1st March 1997 - 31st December 2003 time period.

	MAE			RMSE		
Stock	Implied Volatility	GARCH(1,1)	E-GARCH(1,1)	Implied Volatility	GARCH (1,1)	E-GARCH(1,1)
Aviva	0.1150	0.1153	0.1153	0.1242	0.1241	0.1245
British Aerospace	0.1286	0.1286	0.1287	0.1402	0.1401	0.1402
British Airways	0.1181	0.1183	0.1182	0.1315	0.1316	0.1316
Cadburys	0.0831	0.0835	0.0835	0.0876	0.0881	0.0881
Diageo	0.0910	0.0913	0.0913	0.0981	0.0984	0.0984
Dixons	0.1258	0.1259	0.1259	0.1360	0.1360	0.1360
GlaxoSmithKline	0.0944	0.0947	0.0947	0.0995	0.0998	0.0999
Hilton Group	0.1124	0.1126	0.1126	0.1225	0.1227	0.1227
Hanson	0.1004	0.1008	0.1008	0.1061	0.1064	0.1064
HSBC	0.0948	0.0951	0.0952	0.1014	0.1017	0.1017
Kingfisher	0.1062	0.1064	0.1064	0.1137	0.1138	0.1138
Lonmin	0.1037	0.1040	0.1040	0.1088	0.1091	0.1091
Marks & Spencers	0.1003	0.1006	0.1006	0.1070	0.1073	0.1073
Prudential	0.1155	0.1158	0.1159	0.1255	0.1257	0.1258
Royal & Sun Alliance	0.1337	0.1339	0.1339	0.1471	0.1472	0.1471
Reuters	0.1445	0.1445	0.1445	0.1559	0.1558	0.1558

Table 5: Results for Basic and Encompassing Regression: 1-Day and 5-Day Forecast Horizon. Sample period from 01 October 2002-31 December 2003 using five-minute stock price data to estimate daily realised volatility.

Average Regression Statistics								
Forecast Horizon	1-day				5-day			
Basic Regression	R^2	β_0	β_1		R^2	β_0	β_1	
Implied Volatility	23.64%	-0.0019	3.3291		31.53%	-0.0023	3.6742	
GARCH (1,1)	18.96%	0.0002	1.6222		25.52%	-0.0002	1.6207	
E-GARCH (1,1)	23.07%	0.0001	1.7850		31.48%	-0.0000	1.9064	
Encompassing Regression	R^2	β_0	β_1	β_2	R^2	β_0	β_1	β_2
Implied Volatility	28.15%	-0.0015	2.7083	0.1952	34.55%	-0.0021	2.9038	0.0066
GARCH (1,1)	25.81%	0.0001	1.3220	0.2354	29.68%	-0.0003	1.0520	0.0124
E-GARCH (1,1)	27.80%	0	1.4803	0.1978	33.63%	-0.0002	1.5422	0.0064

Table 6: Results for Basic and Encompassing Regression: 10-Day and 22-Day Forecast Horizon.

Average Regression Statistics								
Forecast Horizon	10-day				22-day			
Basic Regression	R^2	β_0	β_1		R^2	β_0	β_1	
Implied Volatility	36.87%	-0.0018	3.1314		47.66%	-0.0029	4.3108	
GARCH (1,1)	26.00%	-0.0003	1.4028		32.75%	0.0004	1.4234	
E-GARCH (1,1)	35.14%	-0.0001	1.7544		42.06%	-0.0005	2.9199	
Encompassing Regression	R^2	β_0	β_1	β_2	R^2	β_0	β_1	β_2
Implied Volatility	42.69%	-0.0019	2.8092	0.0039	59.54%	-0.0027	2.6079	0.0082
GARCH (1,1)	35.40%	-0.0002	0.9259	0.0072	59.80%	-0.0013	-0.0925	0.0155
E-GARCH (1,1)	41.61%	-0.0001	1.4348	0.0031	60.84%	-0.0013	0.6160	0.0125

Table 7: Summary of empirical results for Diebold and Mariano DM Test for equal predictability and Harvey, et al HLN Encompassing Test. Actual realised volatility are calculated from 5-minute price data from 01 October 2002 to 31 December 2003. A HLN-stat = 0 indicates the model 1 forecast encompasses model 2, a positive HLN-stat indicates that model 2 contains added information.

	Diebold-Mariano Test: GARCH (1,1) vs Implied Volatility							
	1-Day (300 Obs) t-stat = 0.0956		5-Day (60 Obs) t-stat = 0.2143		10-Day (29 Obs) t-stat = 0.3118		22-Day (13 Obs) t-stat = 0.4763	
	DM-stat	HLN-stat	DM-stat	HLN-stat	DM-stat	HLN-stat	DM-stat	HLN-stat
Aviva	0.1043	0.2658	0.02193	0.3486	0.3974	0.5012	0.3812	0.5378
British Aerospace	0.1366	0.2740	0.1202	0.2569	0.2030	0.4524	-0.2389	0.0483
British Airways	-0.0321	0.0036	0.0793	0.2014	0.2054	0.3300	0.2356	0.3922
Cadburys	0.0533	0.1045	0.1466	0.3714	0.3365	0.4723	0.1270	0.2614
Diageo	-0.0044	0.1546	-0.1490	-0.0487	-0.0735	0.0166	-0.1709	0.0416
GlaxoSmithKline	0.1196	0.3092	0.1770	0.3607	0.2174	0.3692	0.2022	0.3750
Hilton Group	0.0658	0.1767	0.0419	0.1562	0.0106	0.3235	0.4629	0.5154
Hanson	-0.0898	0.0407	-0.0832	0.0458	0.0238	0.1558	0.1735	0.2634
HSBC	-0.0193	0.1658	-0.0066	0.2201	0.1671	0.3508	0.2452	0.6169
Kingfisher	0.1487	0.2724	0.1629	0.3400	0.0491	0.2926	0.3131	0.4056
Lonmin	0.0652	0.1688	0.1549	0.3133	0.1817	0.3874	0.1673	0.3270
Marks & Spencers	-0.0142	0.1047	0.0262	0.2870	0.1026	0.3080	0.1934	0.2823
Prudential	0.0831	0.2236	0.0305	0.1992	0.1413	0.3250	0.2173	0.3151
Royal & Sun Alliance	0.0837	0.2122	0.0866	0.2378	0.1762	0.2645	0.2634	0.3809
Reuters	0.588	0.0952	0.0952	0.1444	-0.0984	0.0664	-0.2482	0.5681
Average	**0.0508**	**0.1715**	**0.0735**	**0.2290**	**0.1360**	**0.3077**	**0.1549**	**0.3554**

Table 8: Summary of empirical results for Diebold and Mariano DM Test for equal predictability and Harvey, et al HLN Encompassing Test. Actual realised volatility are calculated from 5-minute price data from 01 October 2002 to 31 December 2003. A HLN-stat = 0 indicates the model 1 forecast encompasses model 2, a positive HLN-stat indicates that model 2 contains added information.

	Diebold-Mariano Test: GARCH (1,1) vs E-GARCH (1,1)							
	1-Day (1708 Obs) t-stat = 0.0956		5-Day (346 Obs) t-stat = 0.2148		10-Day (172 Obs) t-stat = 0.3118		22-Day (80 Obs) t-stat = 0.4763	
	DM-stat	HLN-stat	DM-stat	HLN-stat	DM-stat	HLN-stat	DM-stat	HLN-stat
Aviva	0.2405	0.2998	0.2775	0.3393	0.4018	0.4639	0.3706	0.4369
British Aerospace	0.1170	0.1379	0.1342	0.1696	0.1967	0.2337	-0.2876	-0.2296
British Airways	-0.0247	0.0011	0.1487	0.2447	0.1897	0.3019	0.2218	0.3644
Cadburys	0.0381	0.0749	-0.0239	0.0503	0.2539	0.3274	-0.0073	0.0569
Diageo	0.0867	0.1160	0.0546	0.1011	-0.1461	-0.1160	0.5864	0.05928
GlaxoSmithKline	0.1268	0.1982	0.1743	0.2235	0.2306	0.2884	-0.0772	-0.0064
Hilton Group	0.1199	0.1570	0.2434	0.3105	0.2365	0.3022	0.4246	0.5317
Hanson	0.0527	0.1109	-0.1036	-0.0095	-0.0755	-0.0217	0.1983	0.2521
HSBC	0.0823	0.1239	0.0552	0.1476	0.2108	0.2926	0.0844	0.1707
Kingfisher	0.1600	0.2059	0.1989	0.2385	0.2365	0.2867	0.4202	0.4553
Lonmin	0.0542	0.0733	0.1029	0.1364	0.0776	0.1152	0.1706	0.3142
Marks & Spencers	0.0215	0.0572	0.1429	0.1848	0.1676	0.2090	0.1401	0.1862
Prudential	0.1137	0.1618	0.1461	0.2316	0.1956	0.2597	0.3283	0.3722
Royal & Sun Alliance	0.1370	0.2820	0.0866	0.2763	0.2122	0.3059	0.2537	0.3345
Reuters	0.1490	0.1939	0.0989	0.1458	-0.0156	0.0658	0.2038	0.2838
Average	**0.0983**	**0.1462**	**0.1209**	**0.1860**	**0.1581**	**0.2208**	**0.2020**	**0.2744**