



QUANTITATIVE FINANCE
RESEARCH CENTRE



UNIVERSITY OF
TECHNOLOGY SYDNEY



QUANTITATIVE FINANCE
RESEARCH CENTRE

UTS

THINK.CHANGE.DO

QUANTITATIVE FINANCE RESEARCH CENTRE

Research Paper 127

June 2004

Forecasting Stock Index Volatility: The Incremental Information in the Intraday High-Low Price Range

Charles Corrado and Cameron Truong

ISSN 1441-8010

www.qfrc.uts.edu.au

**Forecasting Stock Index Volatility:
The Incremental Information in the
Intraday High-Low Price Range**

**Charles J. Corrado
University of Technology - Sydney
Sydney, Australia**

and

**Cameron Truong
University of Auckland
Auckland, New Zealand**

June 2004

Abstract

We compare the incremental information content of implied volatility and intraday high-low range volatility in the context of conditional volatility forecasts for three major market indexes: the S&P 100, the S&P 500, and the Nasdaq 100. Evidence obtained from out-of-sample volatility forecasts indicates that neither implied volatility nor intraday high-low range volatility subsumes entirely the incremental information contained in the other. Our findings suggest that intraday high-low range volatility can usefully augment conditional volatility forecasts for these market indexes.

Please direct inquiries to: charles_corrado@hotmail.com

Forecasting Stock Index Volatility: The Incremental Information in the Intraday Price Range

1. Introduction

Since the development of autoregressive conditional heteroscedasticity (ARCH) models by Engle (1982) and their generalization (GARCH) by Bollerslev (1986), ARCH modeling has become the bedrock for dynamic volatility models. While originally formulated to forecast conditional variances as a function of past variances, the inherent flexibility of ARCH modeling allows ready inclusion of other volatility measures as well. Consequently, extensive research has focused on evaluating other volatility measures that might improve conditional volatility forecasts. One popular volatility measure used to augment ARCH forecasts is implied volatility from option prices. Lamoureux and Lastrapes (1993) find that an ARCH model provides superior volatility forecasts than implied volatility alone in a sample of 10 stock return series. However, Day and Lewis (1992) report that a mixture of implied volatility and ARCH forecasts of future return volatility for the S&P 100 stock index outperforms separate forecasts from implied volatility or ARCH alone. More recently, Mayhew and Stivers (2003) find that implied volatility improves GARCH volatility forecasts for individual stocks with high options trading volume. They report that for stocks with the most actively traded options, implied volatility reliably outperforms GARCH and subsumes all information in return shocks beyond the first lag.

Another volatility measure that has become popular with the increasing availability of intraday security price data is an intraday variance computed by summing the squares of intraday returns sampled at short intervals. Essentially, if the security price path is continuous then increasing the sampling frequency yields an arbitrarily precise estimate of return volatility (Merton, 1980). The usefulness of intraday return variances has been demonstrated with foreign exchange data by Andersen et al. (2001b), Andersen, Bollerslev, and Lange (1999), Andersen and Bollerslev (1998), and Martens (2001) and with stock market data by Andersen et al. (2001a), Areal and Taylor (2002), Fleming,

Kirby, and Ostdiek (2003), and Martens (2002). Indeed as a competitor to implied volatility, Taylor and Xu (1997), Pong, Shackleton, Taylor, and Xu (2003), and Neely (2002) report that intraday return variances in the foreign exchange market provide incremental information content beyond that provided by implied volatility forecasts. By contrast, Blair, Poon, and Taylor (2001) find that the incremental information content of intraday return variances for the S&P 100 stock index is scant and that an implied volatility index published by the Chicago Board Options Exchange (CBOE) provides the most accurate forecasts at all forecast horizons.

This paper extends the volatility forecasting literature cited above with the specific objective of demonstrating the usefulness of the intraday high-low price range for improving volatility forecasts for three major stock market indexes: the S&P 100, the S&P 500, and the Nasdaq 100. This study represents the first attempt to compare the effectiveness of the intraday high-low price range and implied volatility as forecasts of future realized volatility for these market indexes.

We find that the intraday high-low range volatility estimator provides incremental information content beyond that already contained in implied volatility indexes published by the Chicago Board Options Exchange (CBOE). This is demonstrated by comparing augmented volatility forecasts based around the asymmetric GARCH model developed by Glosem et al. (1993) and Zakoian (1990), hereafter referred to as GJR-GARCH. Our findings suggest that intraday high-low range volatility can usefully augment conditional volatility forecasts for the three broad market indexes examined.

There are several reasons to consider the intraday high-low price range for volatility measurement and forecasting. Firstly, high-low price range data has long been available in the financial press and is often available when high-frequency intraday returns data are not. Secondly, Andersen and Bollerslev (1998) point out that market microstructure issues such as nonsynchronous trading effects, discrete price observations, and bid-ask spreads, etc. may limit the effectiveness of intraday return variances as volatility forecasts. For example, Andersen et al. (1999) report that sampling intraday returns at one-hour intervals provided better results than sampling at 5-minute intervals in their study of foreign exchange market volatility. The intraday high-low price range may offer a useful alternative to an intraday return variance when market microstructure

effects are severe. Indeed, Alizadeh et al. (2002) suggest that, “Despite the fact that the range is a less efficient volatility proxy than realized volatility under ideal conditions, it may nevertheless prove superior in real-world situations in which market microstructure biases contaminate high-frequency prices and returns.”

Thirdly, in addition to potential market microstructure biases Bai, Russell, and Tiao (2001) point out that the estimation efficiency of an intraday return variance estimator can be sensitive to non-normality in intraday returns data. As a basic demonstration of potential sensitivity to non-normality, let r_d and r_h denote a one-day return and an intraday return, respectively, such that the one-day return is the sum of n intraday returns, i.e., $r_d = \sum_{h=1}^n r_h$. Assuming that the n intraday returns are identically, independently distributed (*iid*), with an expected value of zero, i.e., $E(r_h) = 0$, then the sum of the squared intraday returns is an unbiased estimator of the daily return variance.

$$E\left(\sum_{h=1}^n r_h^2\right) = \text{Var}\left(\sum_{h=1}^n r_h\right) = \text{Var}(r_d) \quad (1)$$

Theoretically, the efficiency of the squared intraday returns volatility estimator specified in equation (1) above increases monotonically by dividing the trading day into finer increments. A general statement of this proposition is provided by the following theorem:

Theorem

The variance of the squared intraday returns volatility estimator, i.e., $\text{Var}\left(\sum_{h=1}^n r_h^2\right)$, assuming *iid* squared intraday returns with zero expected value is given by the expression immediately below in which $Kurt(r_d)$ and $Kurt(r_h)$ denote the kurtosis of daily returns and intraday returns, respectively.

$$\begin{aligned}
\text{Var}\left(\sum_{h=1}^n r_h^2\right) &= \sum_{h=1}^n \text{Var}(r_h^2) \\
&= \sum_{h=1}^n \left(E(r_h^4) - (\text{Var}(r_h))^2\right) \\
&= n \times (\text{Var}(r_h))^2 (Kurt(r_h) - 1) \\
&= (\text{Var}(r_d))^2 \times \frac{(Kurt(r_h) - 1)}{n} \\
&= (\text{Var}(r_d))^2 \times \left(Kurt(r_d) - 3 + \frac{2}{n}\right)
\end{aligned} \tag{2}$$

The last equality on the right-hand side of equation (2) above is an immediate consequence of the assumption of *iid* intraday returns, for which the following relationship holds as an adjunct to the Central Limit Theorem:¹

$$Kurt(r_h) - 3 = n \times (Kurt(r_d) - 3) \tag{3}$$

Thus with given values for the variance and kurtosis of daily returns, i.e., $\text{Var}(r_d)$ and $Kurt(r_d)$, the variance of the squared intraday returns volatility estimator declines monotonically as n increases.

However as shown in the last line of equation (2), the variance of the squared intraday returns volatility estimator is bounded away from zero for non-normally distributed returns with $Kurt(r_d) > 3$. The theoretical relative efficiency of the squared intraday returns volatility estimator to the squared daily return volatility estimator as a function of return kurtosis is stated in equation (4) immediately below.

$$\frac{\text{Var}(r_d^2)}{\text{Var}\left(\sum_{h=1}^n r_h^2\right)} = \frac{Kurt(r_d) - 1}{Kurt(r_d) - 3 + \frac{2}{n}} \tag{4}$$

With exactly normally distributed returns, i.e., $Kurt(r_d) = 3$, this relative efficiency is bounded only by the number of intraday return intervals n . However for plausible kurtosis values, the relative efficiency in equation (4) can be severely bounded. For

¹ An appendix provides a derivation.

example, a daily return kurtosis of $Kurt(r_d) = 4$ with $n = 79$ intraday return intervals yields a theoretical relative efficiency of just 2.93.²

Parkinson (1980) shows that the intraday high-low price range volatility estimator has a theoretical relative efficiency of 4.762 compared to a squared daily return. However, this value assumes normally distributed returns. To assess relative efficiency with non-normally distributed returns, we use Monte Carlo simulation experiments with various return kurtosis values. We then simulate intraday returns over $n = 79$ intraday intervals for each of 100,000 trading days. Kurtotic intraday returns are generated by random sampling from a mixture of normals, where with probability p a random normal variate is drawn with variance σ_p^2 and with probability $1-p$ is drawn with variance σ_{1-p}^2 . The probability p and the ratio of variances determine the kurtosis of the normals' mixture:

$$Kurtosis = \frac{3(p\sigma_p^4 / \sigma_{1-p}^4 + 1 - p)}{(p\sigma_p^2 / \sigma_{1-p}^2 + 1 - p)^2}$$

In each simulated trading day, we compute the sum of squared intraday returns, the squared daily return, and the squared high-low range. Relative efficiencies computed from these daily statistics averaged over 100,000 days are reported in the panel immediately below.

Relative efficiencies of intraday variance estimators to squared daily return estimator with varied kurtosis.

Daily return kurtosis	Squared intraday returns estimator	Squared intraday high-low range estimator	Ratio
3.5	4.786	2.960	1.617
4.0	2.944	2.632	1.118
4.5	2.289	2.462	0.930
5.0	1.997	2.385	0.837

Comparing relative efficiencies for the squared intraday returns estimator and the squared intraday high-low range estimator as shown in the panel above, we see that for plausible

² Bai, Russell, and Tiao (2001) provide an extensive analysis of efficiency losses due to kurtosis and other effects with non-*iid* intraday returns.

kurtosis values the squared intraday returns volatility estimator may not be greatly more efficient than the squared high-low range estimator. Indeed, for daily kurtosis values higher than about 4.3 the squared high-low range estimator is more efficient than the squared intraday returns estimator. Further, Alizadeh et al. (2002) suggest that the intraday high-low range is robust to microstructure noise, while the squared intraday returns estimator can be quite sensitive to such noise.

This study is organized as follows. Section 2 discusses the data used for this research. Section 3 presents the augmented ARCH volatility forecasting methods. Empirical assessments of the extended ARCH models are provided in Section 4. Section 5 summarizes and concludes this paper.

2. Data sources

This study is based on returns for the S&P 100, S&P 500, and Nasdaq 100 stock market indexes, along with daily implied volatilities for these indexes published by the Chicago Board Options Exchange (CBOE). Ticker symbols for the implied volatility indexes are VIX for the S&P 500, VXO for the S&P 100, and VXN for the Nasdaq 100.³ Our data set spans the period January 1990 through December 2003 for the S&P 100 and S&P 500 stock indexes, and from January 1995 through December 2003 for the Nasdaq 100 stock index.

2.1. Daily index returns

Daily index returns are calculated as the natural logarithm of the ratio of consecutive daily closing index levels.

$$r_t = \ln(c_t/c_{t-1}) \quad (5)$$

In equation (5), r_t denotes the index return for day t based on index levels at the close of trading on days t and day $t-1$, i.e., c_t and c_{t-1} , respectively.

³ The CBOE previously used the ticker VIX for S&P 100 implied volatility, but began using VXO for S&P 100 implied volatility and VIX for S&P 500 implied volatility with the introduction of the latter series.

2.2. Daily high-low price range

“..intuition tells us that high and low prices contain more information regarding to volatility than do the opening and closing prices.” (Garman and Klass, 1980) For example, by only looking at opening and closing prices we may wrongly conclude that volatility on a given day is small if the closing price is near the opening price despite large intraday price fluctuations. Intraday high and low values may bring more integrity into an estimate of actual volatility.

In this study, we use the intraday high-low volatility measure specified in equation (6), in which hi_t and lo_t denote the highest and lowest index levels observed during trading on day t .

$$RNG_t^2 = \frac{(\ln hi_t - \ln lo_t)^2}{4 \ln 2} \quad (6)$$

This intraday high-low price range was originally suggested by Parkinson (1980) as a measure of security return volatility.⁴

2.3 CBOE implied volatility indexes

Implied volatilities have long been used by academics and practitioners alike to provide forecasts of future return volatility. In addition to studies cited earlier, Christensen and Prabhala (1998) overcome the methodological difficulties in Canina and Figlewski (1993) and show that by using non-overlapping data and an instrumental variables econometric methodology that implied volatility outperforms historical volatility as a forecast of future return volatility for the S&P 100 index. Corrado and Miller (2004) update the Christensen and Prabhala study and suggest that implied volatility continued to provide a superior forecast of future return volatility during the period 1995 through 2002.

In this study, we use data for three implied volatility indexes published by the Chicago Board Options Exchange (CBOE). These implied volatility indexes are

⁴ Interesting extensions to Parkinson (1980) have been developed by Garman and Klass (1980), Ball and Torous (1984), Rogers and Satchell (1991), Kunitomo (1992), and Yang and Zhang (2000).

computed from option prices for options traded on the S&P 100, the S&P, and the Nasdaq 100 stock indexes.

The implied volatility indexes with ticker symbols VIX and VXN are based on European-style options on the S&P 500 and Nasdaq 100 indexes, respectively. These indexes are calculated using the formula stated immediately below, in which $C(K,T)$ and $P(K,T)$ denote prices for call and put options with strike price K and time to maturity T stated in trading days. This formula assumes the option chain has strike prices ordered such that $K_{j+1} > K_j$. The two nearest maturities are chosen with the restriction that $T_2 \geq 22 \geq T_1 \geq 8$.

$$VIX = \sum_{h=1}^2 (-1)^h \frac{(T_h - 22)}{(T_2 - T_1)} \sum_{j=1}^N \frac{K_{j+1} - K_{j-1}}{K_j^2} \min(C(K_j, T_h), P(K_j, T_h)) \quad (7)$$

Theoretical justification for this calculation method is provided by Britten-Jones and Neuberger (2000).

The implied volatility index with ticker symbol VXO is based on American-style options on the S&P 100 index.⁵ This index is calculated using the formula stated immediately below in which $IV_C(K,T)$ and $IV_P(K,T)$ are implied volatilities for call and put options, respectively, with strike K and maturity T . The at-the-money strike K_m denotes the largest exercise price less or equal to the current cash index S_0 . Hence, the volatility index VXO is calculated using only option contracts with strike prices that bracket the current cash index level.

$$VXO = \frac{\sum_{j=0}^1 \sum_{h=1}^2 (-1)^{j+h} (T_h - 22) (S_0 - K_{m+1-j}) (IV_C(K_{m+j}, T_h) + IV_P(K_{m+j}, T_h))}{(T_2 - T_1)(K_{m+1} - K_m)} \quad (8)$$

To be scaled consistently with the other daily volatility measures, the implied volatility indexes VXO, VIX, and VXN are all squared and divided by 252, the assumed number of trading days in a calendar year.

[INSERT TABLE 1]

⁵ Authoritative descriptions of this implied volatility index are Whaley (1993) and Fleming, Ostdiek, and Whaley (1995).

2.4 Descriptive statistics

Table 1 provides a statistical summary of the volatility data used in this study. Panel A reports the mean, maximum, minimum, standard deviation, and skewness and kurtosis coefficients for squared daily returns, squared implied volatilities, and squared high-low price ranges for the S&P 100 index. Panels B and C report corresponding descriptive statistics for the S&P 500 and Nasdaq 100 indexes, respectively.

The period January 1990 through December 2003 yields 3,544 daily observations for the S&P 100 and S&P 500 indexes and the period January 1995 through December 2003 yields 2,266 daily observations for the Nasdaq 100 index. Table 1 reveals noticeable statistical differences among the three volatility measures. For example, in all panels of Table 1 the average squared high-low range volatility is smaller than the average squared daily return, which in turn is smaller than the average squared implied volatility. Comparing volatility measures across S&P 100, S&P 500, and Nasdaq 100 indexes it is evident that volatility for the Nasdaq 100 is highest among the three indexes. Indeed, the average squared daily return for the Nasdaq 100 index is on average four to five times larger in magnitude than average squared daily returns for the S&P 100 and S&P 500 indexes.

3. Forecast methodology

To model market volatility dynamics we draw on the GJR-GARCH model specification developed by Glosten et al. (1993) and Zakoian (1990). This model attempts to capture the asymmetric effects of good news and bad news on conditional volatility. We augment the basic GJR-GARCH model with implied volatility and intraday high-low price range volatility.

3.1 Augmented GJR-GARCH model

Augmented by implied volatility and the intraday high-low range, the GJR-GARCH model for conditional variance is specified in equation (9) immediately below, in which the dummy variable $s_{t-1} = 1$ if $\varepsilon_{t-1} < 0$ and is zero otherwise.

$$\begin{aligned} r_t &= \mu + \varepsilon_t \\ h_t &= \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 s_{t-1} \varepsilon_{t-1}^2 + \beta h_{t-1} + \gamma IVOL_{t-1}^2 + \delta RNG_{t-1}^2 \end{aligned} \quad (9)$$

r_t return on day t

h_t conditional volatility on day t

$IVOL_t$ implied volatility at end of index options trading on day t

RNG_t intraday high-low range volatility on day t

In this model, good news ($\varepsilon_{t-1} > 0$), and bad news ($\varepsilon_{t-1} < 0$) have differential impacts on conditional variance. The impact of good news alone is measured by the coefficient α_1 , while the impact of bad news is measured by the sum of coefficients $\alpha_1 + \alpha_2$. A priori we expect α_2 alone as well as the sum $\alpha_1 + \alpha_2$ to be positive. Lagged implied volatility and lagged high-low range volatility measures become additional explanatory variables to augment the basic GJR-GARCH model.

By placing varied restrictions on parameters, we obtain four different volatility models that compare the incremental forecast information of implied volatility and high-low price range volatility. These four models are specified as follows:

1) GJR-GARCH(1,1) model

The GJR-GARCH(1,1) model is implemented by setting the restrictions $\gamma = \delta = 0$. This specification yields a model with no exogenous regressors.

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 s_{t-1} \varepsilon_{t-1}^2 + \beta h_{t-1} \quad (10)$$

2) High-low range volatility excluded

This specification has the single restriction $\delta = 0$ to exclude intraday high-low range volatility. It combines the GJR-GARCH(1,1) model with lagged implied volatility as an additional regressor to assess the incremental information content of implied volatility.

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 s_{t-1} \varepsilon_{t-1}^2 + \beta h_{t-1} + \gamma IVOL_{t-1}^2 \quad (11)$$

3) Implied volatility excluded

This specification has the single restriction $\gamma = 0$ to exclude implied volatility. It combines the GJR(1,1) model with intraday high-low range volatility. Comparison with the basic GJR(1,1) model yields an assessment of the incremental information content of the high-low price range volatility.

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 s_{t-1} \varepsilon_{t-1}^2 + \beta h_{t-1} + \delta RNG_{t-1}^2 \quad (12)$$

4) Unrestricted model

This specification has no restrictions and therefore represents a complete implementation of equation (9), which is reproduced here for convenient reference.

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 s_{t-1} \varepsilon_{t-1}^2 + \beta h_{t-1} + \gamma IVOL_{t-1}^2 + \delta RNG_{t-1}^2$$

Parameter estimates for all four specifications stated above are obtained by a quasi-likelihood methodology, by which covariances and standard errors are computed using methods suggested in Bollerslev and Wooldridge (1992).

3.2 Out-of-sample forecasts

Each of the four GJR-GARCH models specified above provides out-of-sample volatility forecasts over $N = 1, 10, 20$ days. We begin by calibrating each model using parameters estimated over an initial estimation period. For S&P 100 and S&P 500 indexes, we use the first 2,000 days of data as an initial parameter estimation period. Due to a shorter time span of available Nasdaq 100 data, we use only the first 1,000 days of data as an initial parameter estimation period. After the initial parameter estimation, each GJR-GARCH model yields out-of-sample volatility forecasts over the $N = 1, 10, 20$ days immediately subsequent to the estimation period. One-day forecasts are obtained directly from each model specification, while $N = 10, 20$ -day forecasts are generated by multiplying the one-day volatility forecast by N . Along with these $N = 1, 10, 20$ -day volatility forecasts, corresponding realized volatilities for the $N = 1, 10, 20$ -day forecasts are computed as the sum of the squared daily returns through the N -day forecast period.

After initial parameter estimation and forecast construction, the entire procedure is repeated by rolling forward the parameter estimation period and re-estimating GJR-GARCH model parameters. This procedure is repeated through all remaining data. For the S&P 100 and S&P 100 indexes, this ultimately yields 1,541 one-day forecasts, 154

10-day forecasts, and 77 20-day volatility forecasts. For the Nasdaq 100 index, this yields 1,266 one-day forecasts, 126 10-day forecasts, and 63 20-day volatility forecasts. All volatility forecasts are out-of-sample forecasts and all forecast intervals are non-overlapping.

Figure 1 provides a graphical illustration of out-of-sample one-day volatility forecasts. Panel A and Panel B display daily realized and forecast volatility values for the 1,541 one-day forecasts for the S&P 100 and S&P 500 stock indexes, respectively. Panel C displays daily realized and forecast volatility values for the 1,266 one-day forecasts for the Nasdaq 100 index. While Figure 1 suggests that the volatility forecasts capture a large proportion of the variability of predictable volatility, a formal evaluation of forecast efficacy is provided immediately below.

[INSERT FIGURE 1]

3.3 Forecast efficiency evaluations

To evaluate the accuracy of the out-of-sample volatility forecasts Blair et al. (2001) suggest the P -statistic specified in equation (13), which measures the proportion of the variance of realized volatilities explained by volatility forecasts. In this P -statistic, $y_{t,N}$ and $\hat{y}_{t,N}$ denote realized and forecast volatility values over a N -day forecast horizon beginning on day t .

$$P = 1 - \frac{\sum_{i=1}^{(T-S)/N} (y_{S+N \times (i-1), N} - \hat{y}_{S+N \times (i-1), N})^2}{\sum_{i=1}^{(T-S)/N} (y_{S+N \times (i-1), N} - \bar{y})^2} \quad (13)$$

For the data used in this study the parameters of equation (13) are $T = 3,544$ and $S = 2,001$ for the S&P 100 and S&P 500 indexes, $T = 2,266$ and $S = 1,001$ for the Nasdaq 100 index, and $N = 1, 10, 20$ for the three different forecast horizons.

Two additional measures of forecast accuracy are the Root Mean Squared Error ($RMSE$) and Mean Absolute Error (MAE); these are computed as shown in equations (14) and (15).

$$RMSE = \frac{1}{(T-S)/N} \sum_{i=1}^{(T-S)/N} \left(y_{S+N \times (i-1), N} - \hat{y}_{S+N \times (i-1), N} \right)^2 \quad (14)$$

$$MAE = \frac{1}{(T-S)/N} \sum_{i=1}^{(T-S)/N} \left| y_{S+N \times (i-1), N} - \hat{y}_{S+N \times (i-1), N} \right| \quad (15)$$

An alternative measure of forecast ability is the R -squared from the regression of an N -day realized volatility $y_{t,N}$ on the corresponding forecast volatility $\hat{y}_{t,N}$, as specified in equation (16).

$$y_{t,N} = \alpha + \beta \hat{y}_{t,N} + e_t \quad (16)$$

To simultaneously compare forecasts made from Model 2 and Model 3, we regress realized volatility on forecasts from Model 2 and Model 3 as specified in equation (17), in which $\hat{y}_{t,N}^{M2}$ denotes an N -day forecast from Model 2 and $\hat{y}_{t,N}^{M3}$ denotes an N -day forecast from Model 3.

$$y_{t,N} = a_1 + a_2 \times \hat{y}_{t,N}^{M2} + a_3 \times \hat{y}_{t,N}^{M3} + e_t \quad (17)$$

The R -squared of this regression is often interpreted as a measure of information content based on the proportion of the variance of volatility explained by the best linear function of the forecasts.

[INSERT TABLE 2]

4. Empirical results

4.1 Full sample GJR-GARCH model results

Table 2 presents GJR-GARCH parameter estimates and related statistics obtained from all available data for all three stock indexes. In Table 3, Panel A reports results obtained from S&P 100 index data, Panel B corresponds to S&P 500 index data, and Panel C reports results from Nasdaq 100 data. Parameter estimates are reported in columns 2 through 7, with robust t -statistics reported in parentheses below each coefficient estimate. Log-likelihood values are listed in column 8, with excess log-likelihoods over Model 1 listed in column 9. Observe that all Durbin-Watson statistics

reported in column 10 are close to a value of two and therefore do not indicate the presence of significant autocorrelations in regression errors.

Panels A and B of Table 2 corresponding to S&P 100 and S&P 500 indexes, respectively, reveal that log-likelihood statistics in column 8 increase monotonically moving from Model 1 to Model 4. For the Nasdaq 100, Panel C reveals a similar pattern, but with a log-likelihood value for Model 2 greater than that for Model 3.

For all models and all indexes, the GJR-GARCH coefficients α_1 measuring the impact of good news are never significantly positive and in some cases are significantly negative. This indicates that good news has either no effect or a small negative effect on the magnitude of conditional volatility. By contrast, the coefficients α_2 are always significantly positive and the coefficient sums $\alpha_1 + \alpha_2$ measuring the full impact of bad news are always positive. This indicates a consistently positive effect of bad news on conditional volatility. Thus, overall we observe a pervasive asymmetric effect of past daily returns on conditional volatility in which bad news ($\varepsilon_{t-1} < 0$) has a positive impact on conditional variance while good news ($\varepsilon_{t-1} > 0$) has a weaker negative impact. This is consistent with empirical findings in Blair et al. (2001) who also find that past volatility has a similar asymmetric impact on conditional volatility.

The augmented GJR-GARCH specification for Model 2 excludes only intraday high-low range volatility and that for Model 3 excludes only implied volatility. In all three panels of Table 2, Model 2 yields significant regression coefficients for implied volatility and Model 3 yields significant regression coefficients for high-low range volatility. For the S&P 100 and S&P 500 indexes, Panels A and B reveal that log-likelihood values for Model 3 are smaller than those for Model 2, thereby suggesting that high-low range volatility has greater information content than does implied volatility. However for the Nasdaq 100 index, Panel C reveals that the log-likelihood value for Model 2 is smaller than that of Model 3. This suggests that implied volatility has greater information content than intraday high-low range volatility for that index. Notwithstanding these differences, the results obtained from Model 4 for all three indexes indicate significant slope coefficients for both intraday high-low range volatility and

implied volatility, thereby suggesting that both volatility measures provide incremental information content not entirely subsumed by the other.

[INSERT TABLE 3]

4.2 Out-of-sample forecast evaluation

Table 3 summarizes the out-of-sample accuracy of volatility forecasts from the various GJR-GARCH model specifications by reporting P -statistics, root mean squared errors (RMSE), and mean absolute errors (MAE) of volatility forecasts for the S&P 100, S&P 500, and Nasdaq 100 indexes. Results for forecast intervals of $N = 1, 10, 20$ days are reported.

Perhaps the most notable aspect of the out-of-sample volatility forecast results reported in Table 3 is the contrasting performance between Model 1 and Model 4. Model 1 represents the basic GJR-GARCH model, while Model 4 is the GJR-GARCH model augmented by both intraday high-low range volatility and implied volatility. Model 4 displays markedly improved performance over Model 1 across all three stock indexes and all forecast horizons. For example, one-day volatility forecasts for the S&P 100, S&P 500, and Nasdaq 100 indexes from Model 1 yield P -statistic values of 0.117, 0.121, and 0.141, respectively. By contrast, Model 4 yields P -statistic values of 0.147, 0.145, and 0.172 for corresponding one-day forecasts. Similarly for 20-day volatility forecasts, Model 1 yields P -statistic values of 0.243, 0.294, and 0.407, while Model 4 yields the corresponding values of 0.367, 0.425, and 0.534. These results indicate that volatility forecasts formed from the GJR-GARCH model are appreciably improved with the additional information contained in implied volatility and high-low range volatility. Notable improvement occurs for all three stock indexes and all forecast horizons. Essentially the same results are reflected in the RMSE and MAE statistics.

Looking more closely at Table 3, a specific comparison of results obtained from Model 2 and Model 3 holds some interest. Model 2 excludes only high-low range volatility and Model 3 excludes only implied volatility. P -statistics in Panel A of Table 4 shows that Model 3 yields superior volatility forecasts to Model 2 at all forecast horizons for the S&P 100 index. By contrast, Model 2 yields P -statistics indicating superior

forecasts to Model 3 for Nasdaq 100 volatility forecasts at all forecast horizons. Results are mixed for the S&P 500 index, for which Model 3 yields P -statistics indicating a superior volatility forecast to Model 2 at the one-day horizon, an inferior forecast at the 10-day horizon, and nearly equivalent forecast performance at the 20-day horizon. Thus evidence presented here does not indicate uniformly superior forecast performance for either the high-low range volatility or implied volatility. However, both Model 2 and Model 3 both provide uniformly superior forecast performance over Model 1 indicating that both intraday high-low range volatility and implied volatility contain information not captured by Model 1.

[INSERT TABLE 4]

Table 4 reports R -squared values and coefficient values from regressions of realized volatility on out-of-sample volatility forecasts across the four models, three indexes, and three forecast horizons. Panel A reports results from S&P 100 volatility forecasts, Panel B reports results for S&P 500 volatility forecasts, and Panel C reports results for Nasdaq 100 volatility forecasts.

The R -squared values shown in Panel A of Table 4 indicate that Model 3 yields out-of-sample S&P 100 volatility forecasts superior to those obtained from Model 2 at forecast horizons of 10 and 20 days. However for one-day S&P 100 volatility forecasts, the R -squared value of 0.148 is the same for Model 3 and Model 2. Model 3 excludes only implied volatility and Model 2 excludes only intraday high-low range volatility. Thus at the 10- and 20-day horizon, it appears that nearly all incremental information contained in implied volatility is subsumed in intraday high-low range volatility. At the one-day horizon, both implied volatility and intraday high-low range volatility appear to provide similar information content for the S&P 100 index. The relative weakness of implied volatility at the 10-day and 20-day forecast horizons is surprising since VXO volatility index represents a volatility forecast over a 22-day horizon.

The R -squared values shown in Panel B of Table 4 indicate that Model 3 yields out-of-sample S&P 500 volatility forecasts superior to those obtained from Model 2 at the one-day forecast horizon, but inferior volatility forecasts at the 10-day and 20-day

horizons. Thus at the one-day horizon, the intraday high-low range volatility appears to yield superior volatility forecasts while implied volatility appears to yield superior forecasts at the 10-day and 20-day horizons. In this case, the relative strength of implied volatility at the 10-day and 20-day forecast horizons is as expected since the *VIX* volatility index represents a 22-day volatility forecast.

In contrast to results reported from Panel A, *R*-squared values shown in Panel C of Table 4 indicate that Model 2 provides superior Nasdaq 100 volatility forecasts at all forecast horizons compared to those obtained from Model 3. For one-day Nasdaq 100 volatility forecasts, the *R*-squared value of 0.183 for Model 2 is almost identical to the value of 0.186 for Model 4. Thus at the one-day horizon, it appears that nearly all incremental information contained in intraday high-low range volatility is already subsumed in implied volatility. However, at the 10-day forecast horizon the *R*-squared value for Model 2 is 0.496, while the corresponding value for Model 4 is 0.546. This indicates that intraday high-low range volatility provides incremental information not entirely subsumed by implied volatility. At the 20-day forecast horizon, Model 2 and Model 4 yield the same *R*-squared value of 0.395, suggesting that information contained in intraday high-low range volatility is already subsumed by implied volatility at this forecast horizon.

Results shown in Table 4 mirror and reinforce those shown in Table 3. These all indicate that intraday high-low range volatility contains incremental information beyond that contained in implied volatility for out-of-sample S&P 100 volatility forecasts at all forecast horizons. However, implied volatility provides incremental information above that contained in intraday high-low range volatility for Nasdaq 100 volatility forecasts at all forecast horizons. For out-of-sample S&P 500 volatility forecasts, the results are mixed. The intraday high-low range appears to provide minimal incremental information over implied volatility at one-day and 20-day horizons, but results favor implied volatility at a 10-day horizon.

In all cases examined, both intraday high-low range volatility and implied volatility bring significant improvements over the basic GJR-GARCH model. Model 2 and Model 3 clearly outperform Model 1 with all indexes across all forecast horizons.

Thus intraday high-low range volatility and implied volatility both provide incremental information for forecasting conditional volatility.

[INSERT TABLE 5]

To assess the relative contributions of high-low range volatility and implied volatility to improved conditional volatility forecasts, Table 5 reports results from regressions of realized volatility against competing out-of-sample forecasts. In these regressions, the coefficients α_2 and α_3 represent slope coefficients for forecasts from Model 2 and Model 3, respectively. In general, slope coefficients for both Model 2 forecasts and Model 3 forecasts are significant across all indexes and forecast horizons. Two exceptions occur for Nasdaq 100 regressions with one-day and 20-day forecast horizons for which the coefficient α_3 has *t*-statistic values of 0.959 and 0.186, respectively. The coefficient α_3 is uniformly larger than the coefficient α_2 at all forecast horizons for the S&P 100 index, indicating relatively greater information content for intraday high-low range volatility. However, the coefficient α_2 is uniformly larger than the coefficient α_3 at all forecast horizons for the S&P 500 index and the Nasdaq 100 index, indicating relatively greater information content for implied volatility.

5. Summary and conclusions

This is the first study to compare the incremental information content of the intraday high-low price range and implied volatility when used to augment ARCH model forecasts of stock market volatility. We examine conditional volatility forecasts for three broad market indexes: the S&P 100 and the S&P 500 over the period January 1990 through December 2003, and the Nasdaq 100 over the period January 1995 through December 2003. We find that the intraday high-low range provides incremental information beyond that already contained in an ARCH model augmented by implied volatility. This finding is demonstrated by comparing conditional volatility forecasts based on various configurations of the GJR-GARCH model augmented by implied volatility and intraday high-low range volatility.

For the S&P 100 index, we find evidence that intraday high-low range volatility provides greater incremental information than implied volatility over one-day, 10-day, and 20-day forecast horizons. By contrast, results obtained from the Nasdaq 100 index indicate that implied volatility provides greater information content than intraday high-low range volatility over all forecast horizons. For the S&P 500 index, our results also favor implied volatility over the intraday high-low range, albeit less dramatically than for the Nasdaq 100.

However for all three major stock indexes examined, we find no evidence to suggest that either intraday high-low range volatility or implied volatility subsumes entirely the information content of the other. Indeed, we find evidence that intraday high-low range volatility might improve ARCH models already augmented by implied volatility. Our empirical results strongly suggest that an ARCH model augmented by intraday high-low volatility and/or implied volatility improves conditional volatility forecasts provided by an ARCH model alone.

Appendix

Theorem

For *iid* returns, $Kurt(r_h) - 3 = n \times (Kurt(r_d) - 3)$.

Proof

By the definitions of variance and kurtosis:

$$Kurt(r_d) = \frac{E(r_d^4)}{(Var(r_d))^2} = \frac{E\left(\sum_{h=1}^n r_h\right)^4}{(nVar(r_h))^2}$$

It is then sufficient to show that for *iid* intraday returns,

$$E\left(\sum_{h=1}^n r_h\right)^4 = (Var(r_h))^2 (nKurt(r_h) + 3n(n-1))$$

We make use of the following identity:

$$\begin{aligned} (Var(r_d))^2 \times Kurt(r_d) &= E(r_d)^4 \\ &= E\left(\sum_{h=1}^n r_h\right)^4 \\ &= E\left(\sum_{h=1}^n \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n r_h r_i r_j r_k\right) \end{aligned}$$

For *iid* returns with $E(r_h)=0$, we have that $E(r_h r_i)=0$ and $E(r_h^2 r_i^2)=(Var(r_h))^2$ for $h \neq i$, and $E(r_h^4)=(Var(r_h))^2 \times Kurt(r_h)$. There are $n(n-1)$ cases in which $h=i$, $j=k$, and $h \neq j$; another $n(n-1)$ disjoint cases in which $h=j$, $i=k$, and $h \neq i$; as well as another $n(n-1)$ cases in which $h=k$, $i=j$, and $h \neq i$. Thus there are $3n(n-1)$ cases of $(Var(r_h))^2$. Also there are n cases where $h=i=j=k$. Thus we obtain,

$$\begin{aligned} E\left(\sum_{h=1}^n r_h\right)^4 &= nE(r_h^4) + 3n(n-1)(Var(r_h))^2 \\ &= (Var(r_h))^2 (nKurt(r_h) + 3n(n-1)) \end{aligned}$$

Substitution into $Kurt(r_d)$ above finishes the proof.

Mixing normals

The kurtosis of a normals mixture is,

$$Kurtosis = \frac{3(p\sigma^4 + 1 - p)}{(p\sigma^2 + 1 - p)^2}$$

If we set $p = 1/Kurtosis$ then σ^2 solves as,

$$\sigma^2 = \frac{Kurtosis - 1 + \sqrt{3((Kurtosis - 2)^2 - 1)}}{2}$$

References

- Alizadeh, S., Brandt, M.W., and Diebold, F.X. (2002). Range-based estimation of stochastic volatility models. *Journal of Finance*, 57, 1047-1091.
- Andersen, T.G., Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economics Review*, 39, 885-905.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., and Ebens, H. (2001a). The distribution of realized stock return volatility. *Journal of Financial Economics*, 61, 43-76.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., and Labys, P. (2001b). The distribution of exchange rate volatility. *Journal of the American Statistical Association*, 96, 42-55.
- Andersen, T.G., Bollerslev, T., and Lange, S. (1999). Forecasting financial market volatility: Sampling frequency vis-à-vis forecast horizon. *Journal of Empirical Finance*, 6, 457-477.
- Areal, N.M.P.C. and Taylor, S.J. (2002). The realized volatility of FTSE-100 futures prices. *Journal Futures Markets*, 22, 627-648.
- Bai, X., Russell, J.R. and Tiao, J.C. (2001). Beyond Merton's Utopia: effects of non-normality and dependence on the precision of variance estimates using high-frequency financial data. *Working Paper*, University of Chicago.
- Ball, C.A. and Torous, W.N. (1984). The maximum likelihood estimation of security price volatility: Theory, evidence, and application to option pricing. *Journal of Business*, 57, 97-112.
- Blair, B.J., Poon, S.H., and Taylor, S.J. (2001). Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns. *Journal of Econometrics*, 105, 5-26.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity . *Journal of Econometrics*, 31, 307-327.
- Bollerslev, T. (1987). A conditional heteroskedastic time series model for speculative prices and rates of returns. *Review of Economics and Statistics*, 69, 542-547.
- Bollerslev, T., Chou, R.Y., and Kroner, K.F. (1992). ARCH modeling in finance: a review of the theory and empirical evidence. *Journal of Econometrics* 52, 5-59.
- Bollerslev, T., Engle, R.F., and Nelson, D.B. (1994). ARCH Models, Chapter 49 of *Handbook of Econometrics*, Vol. 4, North-Holland.

Bollerslev, T., and Wooldridge, M.J. (1992). Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time Varying Covariances. *Econometric Reviews*, 11, 143–172.

Britten-Jones, M. and Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. *Journal of Finance* 45, 839–866.

Canina, L., and Figlewski, S. (1993). The information content of implied volatility. *Review of Financial Studies*, 6, 659-681.

Chiras, D., and Manaster, S. (1978). The information content of option prices and a test of market efficiency. *Journal of Financial Econometrics*, 6, 213-214.

Christensen, B.J., and Prabhala, N.R. (1998). The relation between realized and implied volatility. *Journal of Financial Economics*, 50, 125-150.

Corrado, C.J., and Miller, T.W. (2004). The forecast quality of CBOE implied volatility indexes, Working Paper, Massey University, New Zealand.

Day, T., and Lewis, C. (1992). Stock market volatility and the information content of stock index options. *Journal of Econometrics*, 52, 267-287.

Engle, R.F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica*, 50, 987-1008.

Fleming, J., Kirby, C. and Ostdiek, B. (2003). The economic value of volatility timing using “realized” volatility. *Journal of Financial Economics*, 67, 473-509.

Fleming, J., Ostdiek, B., and Whaley, R.E. (1995). Predicting stock market volatility: A new measure. *Journal of Futures Markets*, 15, 265-302.

Garman, M.B., and Klass, M.J. (1980). On the estimation of price volatility from historical data. *Journal of Business*, 53, 67-78.

Glosten, L.R., Jagannathan, R., and Runkle, D.E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 48, 1779-1801.

Kunitomo, N. (1992). Improving the Parkinson method of estimating security price volatilities. *Journal of Business*, 65, 295-302.

Lamoureux, C.G., and Lastrapes, W.D. (1993). Forecasting stock-return variances: Toward an understanding of stochastic implied volatilities. *Review of Financial Studies*, 6, 293-326.

- Latane, H., and Rendleman, R. (1976). Standard deviation of stock price ratios implied in option prices. *Journal of Finance*, 31, 369-381.
- Matacz, A. (2000). Financial modeling and option theory with the truncated Levy process. *International Journal of Theoretical and Applied Finance*, 3, 143-160.
- Martens, C. (2001). Forecasting daily exchange rate volatility using intraday returns. *Journal of International Money and Finance*, 20, 1-23.
- Martens, C. (2002). Measuring and forecasting S&P 500 index-futures volatility using high-frequency data. *Journal of Futures Markets*, 22, 497-518.
- Mayhew, S., and Stivers, C. (2003). Stock return dynamics, option volume, and the information content of implied volatility. *Journal of Futures Markets*, 23, 615-646.
- Merton, R.C. (2002). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics*, 8, 323-361.
- Neely, C.J. (2002). Forecasting foreign exchange volatility: Is implied volatility the best that we can do?. *Federal Reserve Bank of St. Louis working paper*, 2002-017.
- 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.
- Pong, S., Shackleton, M.B., Taylor, S.J., and Xu, X. (2003). Forecasting currency volatility: a comparison of implied volatilities and AR(FI)MA models. forthcoming *Journal of Banking and Finance*.
- Rogers, L.C.G., and Satchell, S.E. (1991). Estimating variance from high, low and closing prices. *Annals of Applied Probability*, 1, 504-512.
- Taylor, S.J., and Xu, X. (1997). The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance*, 4, 317-340.
- Whaley, E. (1993). Derivatives on market volatility: Hedging tools long overdue. *Journal of Derivatives*, 1, 71-84.
- Yang, D., and Zhang, Q. (2000). Drift-independent volatility estimation based on high, low, open, and close prices. *Journal of Business*, 73, 477-491.
- Zakoian, J.M. (1990). Threshold heteroskedastic models. Manuscript, CREST, INSEE, Paris.

Table 1: Descriptive statistics for volatility data

The data include 3,544 daily observations for the S&P 100 and S&P 500 indexes over the period January 1990 through December 2003, and 2,266 daily observations for the Nasdaq 100 index over the period January 1995 through December 2003. Volatility measures include squared daily returns, implied variance, and squared intraday high-low price range.

	Mean	Max	Min	Standard Deviation	Skewness	Kurtosis
<i>Panel A: S&P 100 index</i>						
Squared daily returns	1.227	56.497	0.000	2.906	7.540	93.640
Implied variance (<i>VXO</i>)	1.308	6.088	0.252	0.872	1.697	6.830
Squared intraday range	0.959	28.094	0.023	1.532	6.866	82.948
<i>Panel B: S&P 500 index</i>						
Squared daily returns	1.111	50.551	0.000	2.615	7.838	103.568
Implied variance (<i>VIX</i>)	1.938	8.302	0.344	1.189	1.697	6.831
Squared intraday range	1.783	59.708	0.026	3.219	7.271	90.352
<i>Panel C: Nasdaq 500 index</i>						
Squared daily returns	5.597	295.94	0.000	12.345	8.911	155.278
Implied variance (<i>VXN</i>)	7.435	34.447	1.195	5.414	1.336	4.422
Squared intraday range	3.786	133.71	0.074	5.912	8.675	142.667

Table 2: GJR-GARCH regressions for daily S&P 100, S&P 500, and Nasdaq 100 index volatility

Parameter estimates for the augmented GJR-GARCH model specified immediately below. The data span the period 1990-2003 for S&P 100 and S&P 500 indexes, and 1995-2003 for the Nasdaq 100 index.

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \alpha_2 s_{t-1} \varepsilon_{t-1}^2 + \beta h_{t-1} + \gamma IVOL_{t-1}^2 + \delta RANGE_{t-1}^2$$

Log-L and D-W indicate maximum likelihood values and Durbin-Watson statistics, respectively. Robust *t*-statistics are reported in parentheses below coefficient values.

	α_0	α_1	α_2	β	γ	δ	Log-L	D-W
<i>Panel A: S&P 100 index</i>								
Model 1	0.013 (4.80)	0.007 (0.72)	0.114 (5.78)	0.927 (98.70)			-4849.4	2.04
Model 2	-0.001 (-0.09)	-0.039 (-2.74)	0.184 (5.59)	0.678 (10.32)	0.215 (4.24)		-4805.9	2.04
Model 3	0.012 (2.49)	-0.084 (-4.19)	0.102 (5.35)	0.872 (44.96)		0.194 (4.55)	-4788.9	2.04
Model 4	0.002 (0.15)	-0.105 (-4.28)	0.146 (5.48)	0.693 (11.81)	0.156 (3.93)	0.174 (3.16)	-4779.1	2.04
<i>Panel B: S&P 500 index</i>								
Model 1	0.012 (4.68)	0.008 (0.80)	0.110 (6.00)	0.929 (93.91)			-4689.5	1.99
Model 2	-0.070 (-2.97)	-0.041 (-5.04)	0.161 (5.16)	0.443 (3.24)	0.474 (3.78)		-4643.1	1.99
Model 3	0.011 (2.74)	-0.092 (-3.96)	0.096 (5.10)	0.881 (45.65)		0.088 (4.22)	-4599.9	1.99
Model 4	-0.055 (-2.63)	-0.091 (-3.39)	0.139 (4.76)	0.479 (3.55)	0.060 (2.06)	0.393 (3.39)	-4595.7	1.99

Table 2 continued

Log-L and D-W indicate maximum likelihood values and Durbin-Watson statistics, respectively. Robust *t*-statistics are reported in parentheses below coefficient values.

	α_0	α_1	α_2	β	γ	δ	Log-L	D-W
<i>Panel C: Nasdaq 100 index</i>								
Model 1	0.057 (4.45)	0.018 (1.05)	0.097 (3.28)	0.924 (76.90)			-4833.5	2.08
Model 2	-0.048 (-0.85)	-0.038 (-1.56)	0.189 (4.17)	0.570 (5.14)	0.415 (3.63)		-4783.5	2.08
Model 3	0.035 (1.75)	-0.055 (-2.88)	0.102 (3.87)	0.860 (43.64)		0.205 (4.96)	-4806.4	2.08
Model 4	-0.049 (-0.91)	-0.053 (-1.27)	0.149 (3.68)	0.607 (6.14)	0.324 (3.49)	0.116 (1.71)	-4780.7	2.08

Table 3: Volatility forecast statistics for S&P 100, S&P 500, and Nasdaq 100 indexes

Evaluations of $N = 1, 10, 20$ day volatility forecasts based on P -statistics, root mean squared errors ($RMSE$), and mean absolute errors (MAE), as specified in the text. The data span the period 1990-2003 for S&P 100 and S&P 500 indexes, and 1995-2003 for the Nasdaq 100 index.

	$N = 1$			$N = 10$			$N = 20$		
	P	$RMSE$	MAE	P	$RMSE$	MAE	P	$RMSE$	MAE
<i>Panel A: S&P 100 index</i>									
Model 1	0.117	3.511	1.905	0.182	14.979	9.367	0.243	25.445	17.114
Model 2	0.139	3.470	1.891	0.244	14.877	9.318	0.278	25.271	16.841
Model 3	0.147	3.449	1.879	0.322	13.952	8.785	0.367	23.956	16.118
Model 4	0.147	3.453	1.877	0.308	14.039	8.872	0.367	24.104	16.101
<i>Panel B: S&P 500 index</i>									
Model 1	0.121	3.157	1.733	0.214	13.145	8.235	0.294	22.523	14.905
Model 2	0.128	3.125	1.715	0.352	12.383	7.966	0.389	21.608	14.418
Model 3	0.145	3.113	1.712	0.272	12.364	7.840	0.387	21.770	14.482
Model 4	0.145	3.109	1.708	0.379	11.980	7.748	0.425	21.170	14.096
<i>Panel C: Nasdaq 100 index</i>									
Model 1	0.141	14.166	7.505	0.417	56.806	34.612	0.407	97.776	63.743
Model 2	0.168	13.808	7.421	0.513	51.623	31.199	0.534	88.881	54.603
Model 3	0.158	14.009	7.48	0.458	54.459	32.258	0.414	95.356	59.428
Model 4	0.172	13.783	7.415	0.531	50.643	30.444	0.534	88.881	54.603

Table 4: Regressions of realized volatility against volatility forecasts for S&P 100, S&P 500, and Nasdaq 100 indexes

Forecast regression equation specified immediately below, in which $y_{t,N}$ denotes an N -day realized volatility and $\hat{y}_{t,N}$ is the corresponding volatility forecast.

$$y_{t,N} = \alpha + \beta \hat{y}_{t,N} + e_t$$

The data span the period 1990-2003 for S&P 100 and S&P 500 indexes, and 1995-2003 for the Nasdaq 100 index. Robust t -statistics are reported in parentheses below coefficient values.

	$N = 1$			$N = 10$			$N = 20$		
	α	β	R^2	α	β	R^2	α	β	R^2
<i>Panel A: S&P 100 index</i>									
Model 1	0.156 (0.87)	0.947 (8.04)	0.117	7.432 (3.28)	0.802 (6.85)	0.232	15.51 (2.90)	0.690 (4.89)	0.241
Model 2	-0.346 (-1.51)	1.285 (8.15)	0.148	6.348 (3.02)	0.703 (7.53)	0.284	12.98 (2.08)	0.894 (4.89)	0.298
Model 3	-0.049 (-0.24)	1.077 (7.99)	0.148	4.662 (2.39)	0.894 (9.30)	0.379	10.98 (2.01)	0.797 (6.77)	0.395
Model 4	-0.407 (-1.72)	1.325 (8.12)	0.158	4.466 (2.23)	0.856 (8.45)	0.328	7.153 (1.35)	0.968 (6.58)	0.373
<i>Panel B: S&P 500 index</i>									
Model 1	0.101 (0.58)	0.964 (7.77)	0.120	3.436 (1.89)	0.827 (6.12)	0.381	11.627 (2.73)	0.679 (4.67)	0.372
Model 2	-0.365 (-1.65)	1.338 (7.84)	0.137	-2.358 (-0.97)	1.308 (7.01)	0.534	1.509 (0.31)	1.080 (5.72)	0.451
Model 3	-0.006 (-0.33)	1.038 (7.53)	0.144	3.512 (2.20)	0.838 (6.37)	0.423	1.423 (2.61)	1.108 (3.75)	0.411
Model 4	-0.414 (-1.86)	1.368 (7.94)	0.147	-2.773 (-1.07)	1.329 (6.67)	0.568	3.547 (0.63)	1.024 (4.61)	0.466

Table 4 continued									
Robust <i>t</i> -statistics reported in parentheses below coefficient values.									
	<i>N</i> = 1			<i>N</i> = 10			<i>N</i> = 20		
	α	β	R^2	α	β	R^2	α	β	R^2
<i>Panel C: Nasdaq 100 index</i>									
Model 1	-0.089 (-0.08)	1.043 -5.58	0.141	15.50 (1.10)	0.834 (3.54)	0.393	67.07 (2.41)	0.579 (2.55)	0.301
Model 2	-2.679 (-1.53)	1.399 (5.20)	0.183	-11.34 (-0.79)	1.207 (5.11)	0.592	32.26 (1.18)	0.835 (3.52)	0.395
Model 3	0.147 (0.14)	1.030 (5.92)	0.159	15.78 (1.23)	0.829 (3.77)	0.456	76.97 (2.14)	0.521 (3.70)	0.350
Model 4	-2.491 (-1.54)	1.383 (5.45)	0.186	-9.560 (-0.62)	1.182 (4.70)	0.592	39.44 (1.30)	0.788 (3.79)	0.395

Table 5: Regression of realized volatility on Model 2 and Model 3 forecasts for S&P 100, S&P 500, and Nasdaq 100 indexes

Forecast regression equation specified immediately below, in which $y_{t,N}$ denotes an N -day realized volatility, and $\hat{y}_{t,N}^{M2}$ and $\hat{y}_{t,N}^{M3}$ represent corresponding volatility forecasts from Model 2 and Model 3, respectively.

$$y_{t,N} = \alpha_1 + \alpha_2 \hat{y}_{t,N}^{M2} + \alpha_3 \hat{y}_{t,N}^{M3} + e_t$$

The data span the period 1990-2003 for S&P 100 and S&P 500 indexes, and 1995-2003 for the Nasdaq 100 index. Robust t -statistics reported in parentheses below coefficient values.

	$N = 1$			$N = 10$			$N = 20$		
	α_2	α_3	R^2	α_2	α_3	R^2	α_2	α_3	R^2
S&P 100	0.571 (2.47)	0.689 (2.49)	0.158	0.592 (2.39)	1.351 (5.20)	0.392	0.747 (2.12)	1.557 (4.63)	0.416
S&P 500	0.689 (2.48)	0.632 (2.84)	0.158	0.706 (2.92)	0.572 (1.95)	0.524	0.743 (2.87)	0.313 (1.56)	0.466
Nasdaq 100	1.111 (2.23)	0.258 (0.96)	0.184	0.647 (2.70)	0.464 (2.69)	0.522	0.981 (2.85)	0.047 (0.19)	0.383

Figure 1: Realized and forecast volatility from Model 4

