Implied volatility indexes and daily Value-at-Risk models

Pierre Giot*

April 14, 2003

ABSTRACT

In this paper, we assess the information content of volatility forecasts based on the VIX and VXN implied volatility indexes in a daily market risk evaluation framework. Our empirical application focuses on the S&P100 and NASDAQ100 indexes and we highlight the models' performances in distinct historical time periods which include bull/bear markets and high/low volatility markets. The performance of the VaR models is evaluated using a wide range of tests which span LR, independence, conditional coverage and density forecast tests. Our results show that straightforward volatility forecasts based on the implied volatility indexes provide meaningful results when market risk must be quantified. Furthermore, the models' performances do not deteriorate in challenging trading environments.

*The author is from Department of Business Administration & CEREFIM at University of Namur, Rempart de la Vierge, 8, 5000 Namur, Belgium, Phone: +32 (0) 81 724887, Email: pierre.giot@fundp.ac.be, and Center for Operations Research and Econometrics (CORE) at Université catholique de Louvain, Belgium. We are grateful to Carol Alexander, Luc Bauwens, Peter Christoffersen, Sébastien Laurent and seminar participants at CORE, Université catholique de Louvain, ISMA Center at the University of Reading and University of Namur who offered helpful comments.

Forecasting volatility has been and still is one of the major success story in the quantitative finance and financial econometrics literature. Indeed, volatility forecasting models have enjoyed a tremendous success since the early 1980's.[1] In financial econometrics, the seminal paper by Engle (1982) has spurred considerable research into ARCH-type models, i.e. the attempt to forecast volatility based on the information given by (past) squared returns. More simple techniques rely on the use of 'rolling window estimation' for the variance of the asset returns.[2] On the other hand, there is a growing trend in the applied finance literature to advocate the use of implied volatility as the best estimate of future volatility. In the framework of an option pricing model such as the Black and Scholes (1973) model, the expected volatility of the asset over the life of the option is the volatility embedded in the price of the option. If call or put option prices are available, then the Black and Scholes (1973) pricing formula can be inverted such that the expected volatility over the life of the option is computed from the observed market prices of the call or put options. Indeed, when all the other option parameters are known, there is a one-to-one relationship between the option prices and underlying (expected) asset volatility. This yields the so-called implied volatility. Details are provided in Hull (2000). Because of the growing importance of modelling and predicting asset volatility, the relevance of implied volatility vs volatility forecasts based on historical returns in order to deliver unbiased and efficient forecasts of future realized volatility is an important topic in modern finance. Moreover volatility forecasting has found numerous applications in quantitative finance, such as portfolio management, option pricing or risk management. In these three fields, volatility forecasts are one of the main inputs to the relevant models and are thus of paramount importance in the empirical application.

While early papers (see a review in Figlewski, 1997) had to rely on somewhat crude datasets, more recent studies use improved databases of actively traded options to evaluate the information content of implied volatility vs volatility computed from historical returns. However, the empirical evidence is rather mixed as to which volatility forecast performs best, although a broad survey of recent papers by Poon and Granger (2003) indicates that, broadly speaking, forecasts based on implied volatility beat forecasts based on historical returns. For example, Day and Lewis (1992) compare the information content of implied volatility of call options on the S&P100 index to GARCH type conditional volatility. Their evidence is rather mixed. Xu and Taylor (1995) focus on the informational efficiency of the PHLX currency options market. According to Jorion (1995) who deals with FOREX data, implied volatility is an efficient but biased forecast of future volatility. Canina and Figlewski (1993) (see also Figlewski, 1997) show that there is almost no correlation between implied volatility and future realized volatility. Christensen and Prabhala (1998) argue that the use of overlapping data and the inclusion of the October 1987 market crash in the Canina and Figlewski (1993) paper is one of the main explanation as to why implied volatility was found inefficient and biased and compared so poorly with volatility forecasts based on historical returns. They show that implied volatility indeed outperforms past volatility in forecasting future volatility and features a high information content. For the S&P100 index and VIX implied volatility index, Blair, Poon, and Taylor (2001) show that historical returns do not provide much incremental information compared to the information given by the VIX index of implied volatility. For three class of assets (stock index, exchange rate and oil), Martens and Zein (2002) show that implied volatility measures do provide superior volatility forecasts compared to daily GARCH-type models. However the switch to high-frequency intraday returns and realized volatility modelled using long memory models alters the outcome of the tests as "long memory volatility forecasts can compete with implied volatility". For foreign exchange volatility and using intraday returns, Neely (2002) argues that implied volatility is a biased estimator of future realized volatility and that volatility forecasts from econometric models should be taken into account. Ederington and Guan (2002) examine the relevance of implied volatility forecasts using S&P500 futures options data and conclude that "implied volatility has strong predictive power and generally subsumes the information in historical volatility". Giot (2003) compares the incremental information content of lagged implied volatility to GARCH models of conditional volatility for a collection of agricultural commodities traded on the New York Board of Trade and shows that past squared returns only marginally improve the information content provided by the lagged implied volatility.

The two implied volatility measures that we consider in this paper are the VIX and VXN implied volatility indexes. They are computed and made available by the Chicago Board of Options Exchange (CBOE) and have enjoyed a widespread acceptance in the community of both academics and practitioners.[3] By construction, the VIX (VXN) index is a weighted average of the implied volatilities computed from a total of eight call and put near-the-money, nearby and second nearby American option contracts on the underlying S&P100 (NASDAQ100) index. The weighting method ensures that VIX/VXN give the implied volatility of a hypothetical at-the-money option with a constant maturity of 22 trading days to expiry. In an efficient market where option prices reflect all available information, the level of VIX (VXN) "is the market's best assessment of the expected volatility of the underlying stock index over the remaining life of the option", a life of 22 trading days and the underlying S&P100 (NASDAQ100) index in this case (quote is from Whaley, 2000, p. 1). Details regarding the construction of the VIX index are available in Whaley (1993), Fleming and Whaley (1995) or Whaley (2000). Note that, by construction, these implied volatility indexes take into account early exercise and dividend payments features, and they do not use as inputs market prices from options that are not actively traded (thus avoiding the troublesome problem of stale quotes for deep out-of-the money or in-the-money options). Thus, the VIX and VXN indexes deliver easy-to-use information regarding future volatility and should be less prone to computation errors than previous measures of implied volatility.

Our study thus builds on the previous work highlighted above as we assess the information content of implied volatility indexes in a Value-at-Risk (more precisely daily earning at risk, or DEAR) framework. Thus we extend the analysis of the latter papers by focusing specifically on the quantification of VaR type market risk using the VIX and VXN implied volatility indexes. We also perform an historical analysis by focusing on distinct sub-periods that include the 1/8/1994 - 30/5/1997 (low volatility, bull market), 2/6/1997 - 31/3/2000 (high volatility, bull market) and 3/4/2000 - 31/1/2003 (high volatility, bear market) time periods. This allows us to test the models' performances in challenging trading environments and look at their stability through time. To assess the VaR performances, we use a range of recent statistical tests which span LR tests (i.e. tests based on the proportion of VaR violations), independence and conditional coverage tests and tests set in the density forecast framework (i.e. tests based on the probability integral transformations of the residuals). Prior to the VaR application, we also run encompassing regressions for the realized volatility at the 5- and 10-day horizon versus implied volatility to assess the efficiency and unbiasedness of these latter volatility forecasts.

When VaR models are estimated, our tests show that volatility forecasts based on the VIX/VXN indexes are meaningful inputs in VaR models as the number of VaR violations is correctly modelled in most cases, the null hypotheses of independence (for the VaR violations) and conditional coverage are usually not rejected and the probability integral transformations are not correlated. As such, our results show that implied volatility indexes do provide accurate and meaningful information as to future volatility forecasts for the underlying indexes when set in a risk management application. Furthermore, there are no real differences in the models' performances across the sub-periods although these were characterized by very different market conditions (bull/bear markets, high/low volatility). This is actually good news as it shows that the VaR models do not break down during 'difficult' market conditions. Our study also provides additional evidence that volatility computed from traded derivatives provides valuable inputs to market participants. This also suggests that options and futures exchanges should compute implied volatility indexes for a large set of market indexes.

The rest of the paper is structured as follows. After this introduction, we briefly assess the efficiency and unbiasedness of implied volatility indexes in Section I. In Section II we focus on the main part of the paper, i.e. the application to the computation of Value-at-Risk models. Finally, Section III concludes.

I. Efficiency and unbiasedness of implied volatility

A. Implied volatility and realized volatility

Although the main focus of the paper is on the VaR analysis, we first assess the information content of the implied volatility indexes at a relatively short time-horizon (5 days and then 10 days). While our analysis is similar to the recent literature in that field, we highlight the performance of the implied volatility indexes in three distinct time periods: the 1/8/1994 - 30/5/1997 (low volatility, bull market), 2/6/1997 - 31/3/2000 (high volatility, bull market) and 3/4/2000 - 31/1/2003 (high volatility, bear market) time periods. Because these three time periods feature almost exactly the same number of observations, a direct comparison of the encompassing regression results is meaningful and gives us insight on the forecasting performance in three different trading environments.

Our measure of implied volatility is given by the level of the VIX and VXN implied volatility indexes, for the S&P100 and NASDAQ100 respectively. By definition and as explained in the introduction of the paper, the forward-looking time horizon is equal to 22 trading days and the implied volatility indexes are expressed in annualized terms. Unfortunately, there are no implied volatility term structures given by the CBOE. Therefore, we use the 'square root of time rule' to switch from a time horizon of 22 days to the required 5- or 10-day interval. Hence and for the 5-day forward-looking horizon, the implied volatility forecast on day t for the S&P100 index is equal to:

$$\sigma_{imp,5,t} = \sqrt{\frac{5}{360}} VIX_t. \tag{1}$$

Thus $\sigma_{imp,5,t}$ is the expected volatility over the $[t+1, t+5]$ period. For the NASDAQ100 index, the corresponding expression is $\sigma_{imp,5,t} = \sqrt{\frac{5}{360}} VXN_t$. For the 10-day time horizon, the volatility forecasts are $\sigma_{imp,10,t} = \sqrt{\frac{10}{360}} VIX_t$ and $\sigma_{imp,10,t} = \sqrt{\frac{10}{360}} VXN_t$.

Given daily returns $r_t = ln(P_t) - ln(P_{t-1})$ for the S&P100 or NASDAQ100 index, the forward-looking realized volatility over a time horizon of 5 days is computed by taking the square root of the sum of the (future) squared returns over this 5-day period. At time t, the forward-looking realized volatility $RV_{5,t}$ for the time period $[t+1, t+5]$ is thus computed as:

$$RV_{5,t} = \sqrt{\sum_{j=1}^{5} r_{t+j}^2}. \tag{2}$$

Note that this volatility measure is computed ex-post, i.e. at time $t+5$ when all returns have been observed. Similar expressions can be computed for the 10-day time horizon. For the encompassing regressions estimated below, we define realized volatility computed from non-overlapping data. Indeed, the measure of realized volatility computed using Equation (2) and using all $\{RV_{5,t}\}$ for $t = 1 \ldots T$ yields strongly correlated volatility measures. As pointed out in Christensen and Prabhala (1998), the use of realized volatility computed from overlapping data in regression analysis yields potentially big estimation problems as the regression's residuals will be strongly auto-correlated. Hence we also define realized volatility measures computed from non-overlapping squared returns data. While Equation (2) is still valid, we no longer compute it for all $t = 1 \ldots T$ but for a subset of those times such that the newly defined $\{RV_{5,k}\}$ use unique data. In this case, it is straightforward to see that the sampling times k are $\{1, 5, 10 \ldots\}$ for the 5-day horizon and $\{1, 10, 20 \ldots\}$ for the 10-day horizon.

B. Encompassing regression analysis

On day t, the ex-post observed realized over the forward-looking 5-day horizon is $RV_{5,t}$. The so-called encompassing regression analysis suggests three relevant statistical tests: (a) is the volatility forecast efficient? (b) is the volatility forecast unbiased? (c) is the volatility forecast unbiased and efficient? For the 5-day time horizon, the encompassing regression is thus:

$$ln(RV_{5,t}) = \beta_0 + \beta_1 ln(\sigma_{imp,5,t}) + \beta_2 ln(RV_{5,t-1}) + e_t. \qquad (3)$$

Note that we estimate the regression for the natural logarithm of implied and realized volatility. This is suggested by most recent studies such as Christensen and Prabhala (1998) and Andersen, Bollerslev, Diebold, and Ebens (2001): the log realized volatility and the log implied volatility exhibit density distributions quite close to the normal distribution. The null hypothesis for the first test (efficiency) is $H0 : \beta_2 = 0$. If it cannot be rejected, then the volatility forecast is efficient with respect to the observed realized volatility $RV_{5,t}$. The null hypothesis of the second test (unbiasedness) is $H0 : \beta_0 = 0$ AND $\beta_1 = 1$. Finally, the volatility forecast is both unbiased and efficient if $H0 : \beta_0 = 0$ AND $\beta_1 = 1$ AND $\beta_2 = 0$ cannot be rejected.[4] For all regression models, the error term e_t is the usual error term in regression analysis, i.e. $e_t \sim N(0, \sigma^2)$. All tests are made using the Wald test statistic programmed within the PcGive econometric software.

Estimation results for the S&P100 index and NASDAQ100 index are given in Tables I and II respectively. Each table is split into 4 panels. Each panel shows the estimation and test results for the 5- and 10-day time horizon. For the test results, we report successively the P-values (Wald test) for the tests of efficiency, unbiasedness and then both efficiency and unbiasedness. The top panel gives the results for the global 1/8/1994 - 31/1/2003 time period. Panels B,C and D highlight the results for the 1/8/1994 - 30/5/1997, 2/6/1997 - 31/3/2000 and 3/4/2000 - 31/1/2003 time periods. As indicated in the introduction, these three sub-periods feature (about) the same number of observations and correspond to three well-characterized time periods: low volatility, bull market; high volatility, bull market; high volatility, bear market. Note that for the NASDAQ100 index, the first sub-period is somewhat shorter as the VXN data is only available from 3/1/1995. For the S&P100 index, the coefficients β_0, β_1 and β_2 are individually quite close to 0, 1 and 0 respectively in most cases. However, the outcome of the Wald tests indicate that most null hypothesis are rejected. The joint test of efficiency and unbiasedness is always rejected, while unbiasedness is rejected in all but one case. For the efficiency test, results are mixed as the corresponding H0 is not rejected at the 5-day horizon in all sub-periods and it is not rejected at the 10-day horizon in the 3/4/2000 - 31/1/2003 time period. Results are rather different regarding the NASDAQ100 index. Indeed, most null hypotheses are not rejected for that stock index at the 10-day time horizon: efficiency is never rejected (whatever the sub-period and also for the global period), unbiasedness and even both efficiency and unbiasedness are never rejected. At the 5-day horizon, unbiasedness and both efficiency and unbiasedness are usually rejected however, while efficiency is not rejected. For both horizons, the coefficients β_0, β_1 and β_2 are individually close to 0, 1 and 0 respectively. Note that, for both stock indexes, results are strikingly similar across the sub-periods and for the global 1/8/1994 - 31/1/2003 time period. Hence the outcome of the efficiency/unbiasedness/efficiency and unbiasedness tests do not seem to depend on the sub-period and whether it is a bull or bear market, or a market exhibiting high or low volatility.

II. Volatility forecasts and the quantification of market risk

A. A brief review of VaR

In this section, we now deal with the information content of the VIX/VXN based volatility forecasts in a market risk evaluation framework. More precisely, we wish to assess the added value of the VIX/VXN based volatility forecasts when these forecasts are used to quantify short-term market risk. We consider the familiar Value-at-Risk framework which provides, at a given percentage level, the most likely loss for a financial institution. For example, the VaR at level α for a time horizon of *h* days is the nominal h-day loss that will not be exceeded in $100 \cdot \alpha$ portfolio realizations out of 100. The literature on VaR models has grown remarkably since the middle of the 1990's because of the popularity of the RiskMetrics VaR specification of JP Morgan and the risk-adjusted measures of capital adequacy enforced by the Basel committee. Recent developments are presented in Jorion (2000), Saunders (2000), Dowd (2002) or Berkowitz and O'Brien (2002) who show that simple GARCH-type univariate models can adequately model the VaR of a global financial institution.

The definition "nominal h-day loss that will not be exceeded in $100 \cdot \alpha$ portfolio realizations out of 100" refers to the statistical notion of a quantile of a density distribution. More precisely, the VaR at time t for the forthcoming $[t+1, t+h]$ time horizon is the expected quantile at the α level for the density distributions of the h-day portfolio returns. In a parametric VaR framework, given a collection of daily demeaned returns r_t, a conditional volatility model for g_t such that $r_t = \sqrt{g_t}\,\varepsilon_t$ and a density distribution for the standardized residuals ε_t, the parametric one-day VaR at time t is immediately given by:[5]

$$VaR_t = Z_\alpha \sqrt{g_t} \tag{4}$$

where Z_α is the quantile at $100 \cdot \alpha$ percent of the standardized density distribution. Following Giot and Laurent (2002) we focus in this paper on the skewed Student density distribution which allows for excess kurtosis and skewness in the distribution of returns.[6] According to Lambert and Laurent (2001), the innovation process ε is said to be (standardized) skewed Student distributed if:

$$f(\varepsilon|\xi,\upsilon) = \begin{cases} \frac{2}{\xi+\frac{1}{\xi}} sg\left[\xi\left(s\varepsilon+m\right)|\upsilon\right] & \text{if} \quad \varepsilon < -\frac{m}{s} \\ \frac{2}{\xi+\frac{1}{\xi}} sg\left[\left(s\varepsilon+m\right)/\xi|\upsilon\right] & \text{if} \quad \varepsilon \geq -\frac{m}{s} \end{cases} \tag{5}$$

where $g(.|\upsilon)$ is the symmetric (unit variance) Student density and ξ is the asymmetry coefficient;[7] m and s^2 are respectively the mean and the variance of the non-standardized skewed Student. Note that this density distribution encompasses the Student density distribution ($\xi = 1$) and the normal density distribution ($\xi = 1$ and $\upsilon = \infty$).

Next to the selection of the adequate density distribution, g_t must be specified to fully characterize the VaR model. In our framework, the implied volatility indexes are the key inputs in the g_t specification, i.e. the volatility part of the VaR model is directly specified by the VIX/VXN indexes. More precisely, we use $g_t = \omega + \eta\,\sigma^2_{imp,1,t-1}$: the volatility input is directly proportional to the 1-day scaled implied volatility (for the S&P100 index, $\sigma_{imp,1,t} = \sqrt{\frac{1}{360}VIX_t}$). Furthermore we introduce two possible alternatives, i.e. two competing models that could easily be used by market practitioners. Our first competing model is the now standard RiskMetrics model, while the second alternative is the GJR-GARCH model of Glosten, Jagannathan, and Runkle (1993). Therefore the three VaR models are:

VIX/VXN

$$VaR_t = Z_\alpha \sqrt{g_t} \tag{6}$$

where $g_t = \omega + \eta\,\sigma^2_{imp,1,t-1}$;

RiskMetrics

$$VaR_t = Z_\alpha \sqrt{g_t} \tag{7}$$

where $g_t = 0.04 r^2_{t-1} + 0.94 g_{t-1}$;

GJR-GARCH

$$VaR_t = Z_\alpha \sqrt{g_t} \tag{8}$$

where $g_t = \omega + \alpha_1 r^2_{t-1} + \alpha_n r^2_{t-1} d_{t-1} + \delta_1 g_{t-1}$.

In all cases, Z_α is the relevant quantile at the $100 \cdot \alpha$ % level from the skewed Student distribution as given in Equation (5). Note that we thus consider a RiskMetrics type volatility specification, but enhanced by the use of a skewed fat tailed density distribution.

B. Back-testing the VaR models

In order to the back-test the VaR results, we first use the familiar Kupiec (1995) LR test. Given the ex-post (i.e. at date $t+1$) observed returns $\{r_{t+1}\}$ and ex-ante (i.e. at date t) forecasted $\{VaR_t\}$, the empirical failure rate $\widehat{f}$ is given by the number of returns smaller than the VaR. If the VaR model is correctly specified, this proportion must be equal to α. More precisely and in the binomial framework, Kupiec (1995) shows that the hypothesis $H_0 : f = \alpha$ against $H_1 : f \neq \alpha$, can be tested with the LR statistic $LR = -2\,ln(\alpha^{T-N}(1-\alpha)^N) + 2\,ln\Big((1-(N/T))^{T-N}(N/T)^N\Big)$, where N is the number of VaR violations, T is the total number of observations and f is the theoretical failure rate. Under the null hypothesis that f is the true failure rate, the LR test statistic is asymptotically distributed as a $\chi^2(1)$.

Secondly we use the independence and conditional coverage tests put forward by Christoffersen (1998). Indeed, the Kupiec (1995) LR test assesses only the equality between the proportion of VaR violations and the expected α level (what is refereed to as the unconditional coverage in Christoffersen, 1998). In a risk management framework, it is also of paramount importance that the VaR violations be uncorrelated over time which leads to the independence and conditional coverage tests based on the evaluation of interval forecasts. Using the same notation as Christoffersen (1998), we defined the indicator sequence of VaR violations as $\{I_t\}$ where I_t is a dummy variable that is equal to 1 if there is a VaR violation at time t (i.e. r_t is smaller than VaR_t) and is equal to 0 if there is no VaR violation at time t. If $\pi_{i,j}$ is the transition probability for two successive I_t dummy variables, i.e. $\pi_{i,j} = P(I_t = j | I_{t-1} = i)$, then the approximate likelihood function for the sequence of I_t is equal to:

$$L_1 = \pi_{0,0}^{n_{0,0}}\,\pi_{0,1}^{n_{0,1}}\pi_{1,0}^{n_{1,0}}\pi_{1,1}^{n_{1,1}} \tag{9}$$

where $n_{i,j}$ is the number of observations with value i followed by value j. The approximate maximized likelihood function is then equal to:

$$\widehat{L}_1 = \left[\frac{n_{0,0}}{n_{0,0}+n_{0,1}}\right]^{n_{0,0}} \left[\frac{n_{0,1}}{n_{0,0}+n_{0,1}}\right]^{n_{0,1}} \left[\frac{n_{1,0}}{n_{1,0}+n_{1,1}}\right]^{n_{1,0}} \left[\frac{n_{1,1}}{n_{1,0}+n_{1,1}}\right]^{n_{1,1}} \tag{10}$$

If the $\{I_t\}$ sequence, i.e. the sequence of VaR violations, is (first-order) independent, then $\pi_{0,0} = 1-\pi$, $\pi_{0,1} = \pi$, $\pi_{1,0} = 1-\pi$ and $\pi_{1,1} = \pi$. This gives the likelihood under the null of first-order independence:

$$L_2 = (1-\pi)^{n_{0,0}+n_{1,0}} \pi^{n_{0,1}+n_{1,1}} \tag{11}$$

which is then estimated as:

$$\widehat{L}_2 = \left[1 - \frac{n_{0,1}+n_{1,1}}{n_{0,0}+n_{1,0}+n_{0,1}+n_{1,1}}\right]^{n_{0,0}+n_{1,0}} \left[\frac{n_{0,1}+n_{1,1}}{n_{0,0}+n_{1,0}+n_{0,1}+n_{1,1}}\right]^{n_{0,1}+n_{1,1}} \tag{12}$$

The LR test statistic for (first-order) independence in the VaR violations is equal to:

$$LR_{ind} = -2(ln(\widehat{L}_2) - ln(\widehat{L}_1)) \sim \chi^2(1). \tag{13}$$

Provided that we condition on the first observation in the test for unconditional coverage, the LR statistic for conditional coverage (i.e. the joint hypothesis of unconditional coverage and independence) is equal to:

$$LR_{cc} = LR_{uc} + LR_{ind} \sim \chi^2(2) \tag{14}$$

where LR_{uc} is the LR statistic for unconditional coverage (computed above for the Kupiec, 1995 test). See additional details and proofs in Christoffersen (1998).

Thirdly we use some of the tests that rely on the density forecast framework such as discussed in Crnkovic and Drachman (1995), Crnkovic and Drachman (1996), Diebold, Gunther, and Tay (1998), Bauwens, Giot, Grammig, and Veredas (2000) and Berkowitz (2001). While we refer the reader to these papers for a full discussion of the density forecast approach (a summary is given in Dowd, 2002), we can summarize the chosen methodology as:

- For all ex-post observed returns $\{r_{t+1}\}$ we first compute the probability integral transformation:

$$x_{t+1} = \int_{-\infty}^{r_{t+1}/\sqrt{g_{t+1}}} f(u)\,du \tag{15}$$

where $f(u)$ is the ex-ante (i.e. computed at time t) standardized skewed Student density distribution as given by Equation (5). Note that g_{t+1} is also computed at time t using one of the three specifications defined above.

- If the model is correctly specified, then it can be shown that x_t should be IID and distributed uniformly on $[0,1]$. Diebold, Gunther, and Tay (1998) suggest a graphical methodology to check this (the histogram should be flat, the ACF should indicate the absence of autocorrelation,…).

- Berkowitz (2001) advocates a likelihood-ratio testing framework as $z_t = \Phi^{-1}(x_t)$ should be IID $N(0,1)$ when x_t is IID uniform over $[0,1]$, i.e. when the VaR model is correctly specified.[8] In this paper, we estimate the following regression:

$$z_t = \beta_0 + \rho_1 z_{t-1} + u_t \tag{16}$$

where $u_t \sim N(0,\sigma^2)$. We first test H0: $\rho_1 = 0$, i.e. that the z_t are uncorrelated and then the much more restrictive H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$, i.e. that the z_t are uncorrelated and that the error term u_t is indeed $N(0,1)$. Last we also look at whether past squared returns are relevant in the extended regression $z_t = \beta_0 + \rho_1 z_{t-1} + \eta_1 r_{t-1}^2 + u_t$ by focusing on the t-statistic of coefficient η_1.

C. Empirical results

We apply the VaR methodology and the tests detailed in Section B to the return series for the S&P100 and NASDAQ100 indexes. The global time period and sub-periods are the same as in Section I which allows us to perform an historical analysis. Because we first need one-day out-of-sample volatility forecasts based successively on the implied volatility, RiskMetrics and GJR-GARCH models, we need rolling estimations of the volatility specifications and skewed Student parameters for the VaR models given in Equations 6, 7 and 8. Note that all models are evaluated in a purely out-of-sample framework. Thus all models require that some initial returns be put aside to estimate the first set of parameters. This implies that, for the NASDAQ100 index, the back-testing global period is somewhat shorter than in Section I as the first two years are used to get the first estimates of the parameters. Therefore the global back-testing period for the NASDAQ100 index ranges from 2/6/1997 to 31/1/2003 and we have two distinct sub-periods: 2/6/1997 - 31/3/2000 and 3/4/2000 - 31/1/2003. For the S&P100 index, the CBOE provides VIX data from 2/1/1986 so that we can keep the same global period and sub-periods as in Section I. After this initialization, all models are re-estimated on a weekly basis as the analysis is rolled forward. At the end of this procedure, we thus have one-day ahead out-of-sample VaR forecasts for the 1/8/1994 - 31/1/2003 (S&P100 index) and for the 2/6/1997 - 31/1/2003 time period (NASDAQ100 index). These VaR forecasts $\{VaR_t\}$, pertaining to the returns defined on $[t, t+1]$, can then be back-tested against the observed returns $\{r_{t+1}\}$. For the back-testing of the VaR forecasts, we first compute the empirical failure rates and Kupiec (1995) LR tests for the right and left quantiles at 10%, 5% and 1%. Note that the left quantile is relevant for a trader who is long the index (trading losses occur on the left side of the returns density), while the right quantile is relevant for a trader who is short the index (trading losses occur on the right side of the returns density). Secondly we compute the independence and conditional coverage tests detailed above (see Equations (13) and (14)). Thirdly, we compute the x_t as in Equation (15), then the $z_t = \Phi^{-1}(x_t)$ and estimate Equation (16).

The empirical results for the three models are given in Tables III, IV, V and VI (S&P100 index, respectively global 1/8/1994 - 31/1/2003 time period and the three sub-periods) and Tables VII, VIII and IX (NASDAQ100 index, respectively global 2/6/1997 - 31/1/2003 and the two sub-periods). Thereafter and for the two indexes, sub-period 1 refers to 1/8/1994 - 30/5/1997, sub-period 2 refers to 2/6/1997 - 31/3/2000 while sub-period 3 is for 3/4/2000 - 31/1/2003 (thus, results in sub-periods 1, 2 and 3 are available for the S&P100 index, results in sub-periods 2 and 3 are available for the NASDAQ100 index). Looking at failure rates and the outcome of the Kupiec (1995) LR tests, we see that there are few differences between the competing models, although the VIX/VXN based model seems to be somewhat better while the performance of the GJR-GARCH model is somewhat disappointing for the NASDAQ100 index. Quite interestingly, there are no real differences across the sub-periods. For example, the VIX/VXN based models perform equally well in volatile/quiet markets and in bull/bear markets. Furthermore, rather simple models such as the RiskMetrics or implied volatility models do not underperform with respect to the more sophisticated GJR-GARCH model.[9] The P-values for the independence test (rows labelled by I in the tables) show that the null hypothesis of first-order independence is not rejected for all models except for the RiskMetrics model with the left tail at 5% (second sub-period). Note that when $\alpha = 1\%$ there are many cases where the test cannot be used because one has $n11 = 0$ (i.e. there are no successive VaR violations). The P-values for the conditional coverage test (rows labelled by CC in the table) show that the null hypothesis of conditional coverage is usually not rejected. This indicates that the models' performances are quite stable across time and do not deteriorate in turbulent markets. For example, all three models perform well in the third sub-period (3/4/2000 - 31/1/2003) although these were not easy times for market participants. While the failure rates tests are mandatory in the Basel committee requirements, we also have the interesting result (according to the independence tests) that VaR violations are not correlated. In the density forecast framework, the $z_t = \Phi^{-1}(\int_{-\infty}^{r_t/\sqrt{g_t}} f(u)\,du)$ are always uncorrelated, except in sub-period 1 (S&P100 index). In sub-period 2 (S&P100 index), η_1 is significant in the $z_t = \beta_0 + \rho_1 z_{t-1} + \eta_1 r_{t-1}^2 + u_t$ regression which suggests that some volatility clustering has not been correctly taken into account. Quite surprisingly and although the skewed Student density distribution is very flexible, the null hypothesis that H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$ is rejected for all models and for both indexes (except in sub-period 3 for the S&P100).[10] Note however that we focus in that case on the whole density distribution, and not on the left or right tails. Regarding the volatility forecasts made from the VIX/VXN indexes, these tests thus show that they are meaningful inputs in VaR models as the number of VaR violations is correctly modelled in most cases, the null hypotheses of independence and conditional coverage are usually not rejected and the probability integral transformations are not correlated. This is however also true in most cases for the simple RiskMetrics specification.

III. Conclusion

Our study dealt with the information content of the VIX and VXN implied volatility indexes when these are taken as volatility inputs in a daily Value-at-Risk model. Our empirical analysis focused on the 1/8/1994 - 31/3/2003 time period but we also focused on distinct sub-periods that include the 1/8/1994 - 30/5/1997 (low volatility, bull market), 2/6/1997 - 31/3/2000 (high volatility, bull market) and 3/4/2000 - 31/1/2003 (high volatility, bear market) time periods. This allowed us to test the models' performances in challenging trading environments and look at their stability through time. To assess the VaR performances, we used a wide range of statistical tests such as LR tests, independence and conditional coverage tests and tests set in the density forecast framework. Prior to the VaR application, we also assessed the efficiency and unbiasedness of the implied volatility indexes with respect to the realized volatility during the global period and three sub-periods.

Regarding the VaR application, the statistical tests show that implied volatility indexes provide meaningful volatility information in VaR models as the number of VaR violations is correctly modelled in most cases, the null hypotheses of independence and conditional coverage are usually not rejected and the probability integral transformations are not correlated. This is however also true in most cases for the RiskMetrics and GJR-GARCH specifications (which are the competing models in our VaR analysis). Our study also shows that the models' performances do not deteriorate during the challenging environments (i.e. the second and third sub-periods) and are stable through time. This is also the case for the outcome of the efficiency and unbiasedness tests. Broadly speaking, this analysis supports the hypothesis that rather simple market-based volatility forecasts combined with a flexible density distribution provides adequate inputs to be used in risk quantification models. Moreover, these models do not break down during challenging trading environments.

References

Alexander, C., 2001, *Market models*. (Wiley).

Andersen, T.G., T. Bollerslev, F.X. Diebold, and H. Ebens, 2001, The distribution of stock return volatility, *Journal of Financial Economics* 61, 43–76.

Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys, 2002, Modelling and forecasting realized volatility, Mimeo.

Bauwens, L., P. Giot, J. Grammig, and D. Veredas, 2000, A Comparison of Financial Duration Models Via Density Forecasts, CORE DP 2060.

Berkowitz, J., 2001, Testing density forecasts, with applications to risk management, *Journal of Business and Economic Statistics* 19, 465–474.

Berkowitz, J., and J. O'Brien, 2002, How accurate are Value-at-Risk models at commercial banks?, *Journal of Finance* 57, 1093–1112.

Black, F., and M. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659.

Blair, B.J., S. Poon, and S.J. Taylor, 2001, Forecasting S&P100 volatility: the incremental information content of implied volatilities and high-frequency index returns, *Journal of Econometrics* 105, 5–26.

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

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

Christoffersen, P.F., 1998, Evaluating interval forecasts, *International Economic Review* 39, 841–862.

Claessen, H., and S. Mittnik, 2002, Forecasting stock market volatility and the informational efficiency of the DAX-index options market, CFS Working Paper 2002/04.

Crnkovic, C., and J. Drachman, 1995, A universal tool to discriminate among risk measurement techniques, Mimeo.

Crnkovic, C., and J. Drachman, 1996, Quality control, *Risk* 9, 139–143.

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

Diebold, F. X., T. A. Gunther, and A.S. Tay, 1998, Evaluating Density Forecasts, with Applications to Financial Risk Management, *International Economic Review* 39, 863–883.

Dowd, K., 2002, *Measuring market risk*. (Wiley).

Ederington, L., and W. Guan, 2002, Is implied volatility an informationally efficient and effective predictor of future volatility?, *Journal of Risk* 4, 3.

Engle, R.F., 1982, Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* 50, 987–1007.

Figlewski, S., 1997, Forecasting volatility, *Financial Markets, Institutions and Instruments* 6, 2–87.

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

Giot, P., 2003, The information content of implied volatility in agricultural commodity markets, *Journal of Futures Markets* 23, 441–454.

Giot, P., and S. Laurent, 2001, Modelling daily Value-at-Risk using realized volatility and ARCH type models, Maastricht University METEOR RM/01/026.

Giot, P., and S. Laurent, 2002, Value-at-Risk for long and short trading positions, *Forthcoming in the Journal of Applied Econometrics*.

Glosten, L.R., R. Jagannathan, and D.E. Runkle, 1993, On the relation between expected excess return on stocks, *Journal of Finance* 48, 1779–1801.

Hull, J.C., 2000, *Options, futures and other derivatives*. (Prentice-Hall).

Jorion, P., 1995, Predicting volatility in the Foreign Exchange market, *Journal of Finance* 50, 507–528.

Jorion, P., 2000, *Value-at-Risk*. (McGraw-Hill).

Kupiec, P., 1995, Techniques for verifying the accuracy of risk measurement models, *Journal of Derivatives* 2, 173–84.

Lambert, P., and S. Laurent, 2001, Modelling financial time series using GARCH-type models and a skewed Student density, Mimeo, Université de Liège.

Martens, M., and J. Zein, 2002, Predicting financial volatility: high-frequency time-series forecasts vis-a-vis implied volatility, Mimeo, Erasmus University Rotterdam.

Neely, C.J., 2002, Forecasting foreign exchange volatility: is implied volatility the best we can do?, Federal Reserve Bank of St.Louis, WP 2002-017.

Pagan, A., and G. Schwert, 1990, Alternative models for conditional stock volatility, *Journal of Econometrics* 45, 267–290.

Poon, S., and C. Granger, 2003, Forecasting volatility in financial markets: a review, *Forthcoming in the Journal of Economic Literature*.

Saunders, A., 2000, *Financial institutions management*. (McGraw-Hill).

Whaley, R.E., 1993, Derivatives on market volatility: hedging tools long overdue, *Journal of Derivatives* 1, 71–84.

Whaley, R.E., 2000, The investor fear gauge, Mimeo, The Fuqua School of Business, Duke University.

Xu, X., and S.J. Taylor, 1995, Conditional volatility and the informational efficiency of the PHLX currency options market, *Journal of Banking and Finance* 19, 803–821.

Table I

Realized volatility: encompassing regressions (S&P100 index)

Horizon	β_0	β_1	β_2	E	U	UE
Panel A: 1/8/1994 - 31/1/2003						
5-day	-0.371 (0.064)	1.074 (0.087)	0.106 (0.051)	0.004	0	0
10-day	-0.193 (0.091)	0.808 (0.112)	0.247 (0.077)	0.001	0	0
Panel B: 1/8/1994 - 30/5/1997						
Horizon	β_0	β_1	β_2	E	U	UE
5-day	-0.267 (0.106)	0.869 (0.182)	0.118 (0.086)	0.171	0	0
10-day	0.032 (0.155)	0.450 (0.201)	0.361 (0.128)	0.005	0	0
Panel C: 2/6/1997 - 31/3/2000						
Horizon	β_0	β_1	β_2	E	U	UE
5-day	-0.268 (0.219)	1.033 (0.224)	0.017 (0.089)	0.848	0.007	0
10-day	-0.110 (0.336)	0.918 (0.293)	0.055 (0.140)	0.696	0.367	0
Panel D: 3/4/2000 - 31/1/2003						
Horizon	β_0	β_1	β_2	E	U	UE
5-day	-0.336 (0.182)	1.029 (0.190)	0.162 (0.089)	0.069	0	0
10-day	0.042 (0.239)	0.594 (0.216)	0.333 (0.129)	0.010	0.008	0.001

Results for the encompassing regression of the log realized volatility at the 5- and 10- horizon vs the log implied volatility forecasts and lagged log realized volatility, i.e. $ln(RV_{h,t}) = \beta_0 + \beta_1 ln(\sigma_{imp,h,t}) + \beta_2 ln(RV_{h,t-1}) + e_t$ where $RV_{h,t}$ is the realized volatility at the h-day horizon and $\sigma_{imp,h,t}$ is the VIX-based volatility forecast. E gives the P-value for the Wald test of efficiency, i.e. H0: $\beta_2 = 0$; U gives the P-value for the Wald test of unbiasedness, i.e. H0: $\beta_0 = 0$ AND $\beta_1 = 1$; UE gives the P-value for the Wald test of unbiasedness and efficiency, i.e. H0: $\beta_0 = 0$ AND $\beta_1 = 1$ AND $\beta_2 = 0$. Results are for the S&P100 index (non-overlapping realized volatility data).

Table II

Realized volatility: encompassing regressions (NASDAQ100 index)

Horizon	β_0	β_1	β_2	E	U	UE
Panel A: 3/1/1995 - 31/1/2003						
5-day	-0.257 (0.074)	1.015 (0.076)	0.096 (0.052)	0.064	0	0
10-day	-0.140 (0.105)	1.073 (0.107)	-0.025 (0.081)	0.764	0.397	0.069
Panel B: 3/1/1995 - 30/5/1997						
Horizon	β_0	β_1	β_2	E	U	UE
5-day	-0.251 (0.180)	1.017 (0.186)	0.110 (0.091)	0.228	0.016	0
10-day	-0.001 (0.267)	0.918 (0.238)	0.044 (0.140)	0.751	0.823	0.333
Panel C: 2/6/1997 - 31/3/2000						
Horizon	β_0	β_1	β_2	E	U	UE
5-day	-0.370 (0.208)	1.151 (0.175)	0.022 (0.089)	0.803	0.070	0.001
10-day	-0.103 (0.283)	1.173 (0.212)	-0.156 (0.134)	0.244	0.642	0.201
Panel D: 3/4/2000 - 31/1/2003						
Horizon	β_0	β_1	β_2	E	U	UE
5-day	-0.579 (0.276)	1.144 (0.193)	0.133 (0.092)	0.147	0.005	0.003
10-day	-0.090 (0.426)	0.958 (0.291)	0.069 (0.161)	0.670	0.803	0.759

Results for the encompassing regression of the log realized volatility at the 5- and 10- horizon vs the log implied volatility forecasts and lagged log realized volatility, i.e. $ln(RV_{h,t}) = \beta_0 + \beta_1 ln(\sigma_{imp,h,t}) + \beta_2 ln(RV_{h,t-1}) + e_t$ where $RV_{h,t}$ is the realized volatility at the h-day horizon and $\sigma_{imp,h,t}$ is the VXN-based volatility forecast. E gives the P-value for the Wald test of efficiency, i.e. H0: $\beta_2 = 0$; U gives the P-value for the Wald test of unbiasedness, i.e. H0: $\beta_0 = 0$ AND $\beta_1 = 1$; UE gives the P-value for the Wald test of unbiasedness and efficiency, i.e. H0: $\beta_0 = 0$ AND $\beta_1 = 1$ AND $\beta_2 = 0$. Results are for the NASDAQ100 index (non-overlapping realized volatility data).

Table III
VaR results for the S&P100 index (1/8/1994 - 31/1/2003)

Lagged implied volatility (VIX)						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**11.54**	5.37	1.17	10.84	**6.26**	1.03
I	0.59	0.32	-	0.51	0.82	-
CC	0.06	0.45	-	0.39	0.03	-
H0: $\rho_1 = 0$; P-value = 0.224						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0						
H0: $\eta_1 = 0$; t-stat = 2.742						

RiskMetrics						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	11.30	**6.17**	**1.54**	10.93	**6.26**	1.03
I	0.33	0.31	0.10	0.44	0.88	-
CC	0.09	0.03	0.02	0.30	0.04	-
H0: $\rho_1 = 0$; P-value = 0.088						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0						
H0: $\eta_1 = 0$; t-stat = 2.922						

GJR-GARCH						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**11.86**	5.65	1.21	**11.63**	5.74	**0.56**
I	0.56	0.73	-	0.53	0.19	-
CC	0.02	0.37	-	0.04	0.13	-
H0: $\rho_1 = 0$; P-value = 0.114						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0.002						
H0: $\eta_1 = 0$; t-stat = 2.334						

Summary results for the back-testing of the out-of-sample VaR models on the S&P100 index, 1/8/1994 - 31/1/2003. The first volatility specification (top panel) is $g_t = \omega + \eta \sigma^2_{imp,1,t-1}$, the second volatility specification (middle panel) is $g_t = 0.06 r^2_{t-1} + 0.94 g_{t-1}$ while the third volatility specification (bottom panel) is $g_t = \omega + \alpha_1 r^2_{t-1} + \alpha_n r^2_{t-1} d_{t-1} + \delta_1 g_{t-1}$, all three with a skewed Student density distribution for the error term. We first give the empirical failure rates for the left (LQ) and right (RQ) quantiles at 10, 5 and 1%. A bold figure indicates that the empirical failure rate is significantly different (LR test or unconditional coverage) from the theoretical value. Secondly, *I* and *CC* give the P-values for the independence and conditional coverage tests respectively. Thirdly, the statistical tests (density forecast framework) $H0: \ldots$ are for $z_t = \beta_0 + \rho_1 z_{t-1} + \eta_1 r^2_{t-1} + u_t$, $u_t \sim N(0, \sigma^2)$, where z_t is equal to $\Phi^{-1}(\int_{-\infty}^{r_t/\sqrt{g_t}} f(u)\, du)$.

Table IV
VaR results for the S&P100 index (1/8/1994 - 30/5/1997)

Lagged implied volatility (VIX)

	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	9.36	4.89	1.26	12.15	6.42	0.70
I	0.90	0.82	-	0.39	0.98	-
CC	0.84	0.97	-	0.15	0.24	-

H0: $\rho_1 = 0$; P-value = 0.006
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0.004
H0: $\eta_1 = 0$; t-stat = 0.746

RiskMetrics

	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	10.20	6.42	**1.96**	**13.41**	**7.82**	1.26
I	0.54	0.54	0.03	0.31	0.42	-
CC	0.81	0.20	0.01	0.01	0	-

H0: $\rho_1 = 0$; P-value = 0.005
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0.001
H0: $\eta_1 = 0$; t-stat = 0.520

GJR-GARCH

	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	9.08	5.45	1.40	**12.99**	6.56	0.42
I	0.36	0.93	-	0.97	0.96	-
CC	0.47	0.86	-	0.04	0.18	-

H0: $\rho_1 = 0$; P-value = 0.002
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0
H0: $\eta_1 = 0$; t-stat = 0.652

Summary results for the back-testing of the out-of-sample VaR models on the S&P100 index, 1/8/1994 - 30/5/1997. The first volatility specification (top panel) is $g_t = \omega + \eta\sigma^2_{imp,1,t-1}$, the second volatility specification (middle panel) is $g_t = 0.06r^2_{t-1} + 0.94g_{t-1}$ while the third volatility specification (bottom panel) is $g_t = \omega + \alpha_1 r^2_{t-1} + \alpha_n r^2_{t-1} d_{t-1} + \delta_1 g_{t-1}$, all three with a skewed Student density distribution for the error term. We first give the empirical failure rates for the left (LQ) and right (RQ) quantiles at 10, 5 and 1%. A bold figure indicates that the empirical failure rate is significantly different (LR test or unconditional coverage) from the theoretical value. Secondly, *I* and *CC* give the P-values for the independence and conditional coverage tests respectively. Thirdly, the statistical tests (density forecast framework) $H0 : \ldots$ are for $z_t = \beta_0 + \rho_1 z_{t-1} + \eta_1 r^2_{t-1} + u_t$, $u_t \sim N(0, \sigma^2)$, where z_t is equal to $\Phi^{-1}(\int_{-\infty}^{r_t/\sqrt{g_t}} f(u)\,du)$.

Table V
VaR results for the S&P100 index (2/6/1997 - 31/3/2000)

Lagged implied volatility (VIX)						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	11.75	**6.99**	1.15	11.75	6.43	1.40
I	0.14	-	-	0.75	0.53	-
CC	0.10	-	-	0.29	0.19	-
H0: $\rho_1 = 0$; P-value = 0.919						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0						
H0: $\eta_1 = 0$; t-stat = 2.809						
RiskMetrics						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	10.49	6.57	1.68	10.63	6.15	0.70
I	0.96	0.14	-	0.57	0.21	-
CC	0.90	0.06	-	0.85	0.18	-
H0: $\rho_1 = 0$; P-value = 1						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0						
H0: $\eta_1 = 0$; t-stat = 2.878						
GJR-GARCH						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	11.75	**6.85**	1.82	**13.29**	6.29	1.12
I	0.69	0.39	-	0.91	0.19	-
CC	0.29	0.07	-	0.02	0.13	-
H0: $\rho_1 = 0$; P-value = 0.977						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0						
H0: $\eta_1 = 0$; t-stat = 2.380						

Summary results for the back-testing of the out-of-sample VaR models on the S&P100 index, 2/6/1997 - 31/3/2000. The first volatility specification (top panel) is $g_t = \omega + \eta\sigma^2_{imp,1,t-1}$, the second volatility specification (middle panel) is $g_t = 0.06r^2_{t-1} + 0.94g_{t-1}$ while the third volatility specification (bottom panel) is $g_t = \omega + \alpha_1 r^2_{t-1} + \alpha_n r^2_{t-1} d_{t-1} + \delta_1 g_{t-1}$, all three with a skewed Student density distribution for the error term. We first give the empirical failure rates for the left (LQ) and right (RQ) quantiles at 10, 5 and 1%. A bold figure indicates that the empirical failure rate is significantly different (LR test or unconditional coverage) from the theoretical value. Secondly, *I* and *CC* give the P-values for the independence and conditional coverage tests respectively. Thirdly, the statistical tests (density forecast framework) $H0 : \ldots$ are for $z_t = \beta_0 + \rho_1 z_{t-1} + \eta_1 r^2_{t-1} + u_t$, $u_t \sim N(0, \sigma^2)$, where z_t is equal to $\Phi^{-1}(\int_{-\infty}^{r_t/\sqrt{g_t}} f(u)\,du)$.

Table VI
VaR results for the S&P100 index (3/4/2000 - 31/1/2003)

Lagged implied volatility (VIX)						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**13.66**	4.51	0.85	8.73	5.77	1.27
I	0.93	0.23	-	0.79	0.30	-
CC	0.01	0.41	-	0.51	0.38	-
H0: $\rho_1 = 0$; P-value = 0.525						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0.331						
H0: $\eta_1 = 0$; t-stat = 1.039						
RiskMetrics						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**13.24**	5.77	0.85	8.73	4.79	1.27
I	0.42	0.03	-	0.84	0.58	-
CC	0.02	0.06	-	0.51	0.83	-
H0: $\rho_1 = 0$; P-value = 0.902						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0.200						
H0: $\eta_1 = 0$; t-stat = 1.110						
GJR-GARCH						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**12.53**	3.94	0.99	8.87	5.77	1.13
I	0.453	0.91	-	0.85	0.68	-
CC	0.07	0.41	-	0.59	0.59	-
H0: $\rho_1 = 0$; P-value = 0.331						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0.051						
H0: $\eta_1 = 0$; t-stat = 0.869						

Summary results for the back-testing of the out-of-sample VaR models on the S&P100 index, 3/4/2000 - 31/1/2003. The first volatility specification (top panel) is $g_t = \omega + \eta\sigma^2_{imp,1,t-1}$, the second volatility specification (middle panel) is $g_t = 0.06r^2_{t-1} + 0.94g_{t-1}$ while the third volatility specification (bottom panel) is $g_t = \omega + \alpha_1 r^2_{t-1} + \alpha_n r^2_{t-1}d_{t-1} + \delta_1 g_{t-1}$, all three with a skewed Student density distribution for the error term. We first give the empirical failure rates for the left (LQ) and right (RQ) quantiles at 10, 5 and 1%. A bold figure indicates that the empirical failure rate is significantly different (LR test or unconditional coverage) from the theoretical value. Secondly, *I* and *CC* give the P-values for the independence and conditional coverage tests respectively. Thirdly, the statistical tests (density forecast framework) $H0:\ldots$ are for $z_t = \beta_0 + \rho_1 z_{t-1} + \eta_1 r^2_{t-1} + u_t$, $u_t \sim N(0,\sigma^2)$, where z_t is equal to $\Phi^{-1}(\int_{-\infty}^{r_t/\sqrt{g_t}} f(u)\,du)$.

Table VII
VaR results for the NASDAQ100 index (2/6/1997 - 31/1/2003)

Lagged implied volatility (VXN)

	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**12.69**	5.05	0.77	10.03	5.47	1.26
I	0.14	0.33	-	0.63	0.20	-
CC	0	0.62	-	0.89	0.31	-

H0: $\rho_1 = 0$; P-value = 0.226
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0
H0: $\eta_1 = 0$; t-stat = 1.450

RiskMetrics

	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**13.11**	6.03	0.63	10.24	5.40	0.91
I	0.74	0.56	0.10	0.78	0.67	-
CC	0	0.19	0.02	0.92	0.72	-

H0: $\rho_1 = 0$; P-value = 0.655
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0
H0: $\eta_1 = 0$; t-stat = 1.238

GJR-GARCH

	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**13.46**	6.10	0.91	**11.92**	**7.01**	**1.61**
I	0.17	0.75	-	0.40	0.25	-
CC	0	0.17	-	0.04	0	-

H0: $\rho_1 = 0$; P-value = 0.401
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0
H0: $\eta_1 = 0$; t-stat = 0.966

Summary results for the back-testing of the out-of-sample VaR models on the NASDAQ100 index, 2/6/1997 - 31/1/2003. The first volatility specification (top panel) is $g_t = \omega + \eta\sigma^2_{imp,1,t-1}$, the second volatility specification (middle panel) is $g_t = 0.06r^2_{t-1} + 0.94g_{t-1}$ while the third volatility specification (bottom panel) is $g_t = \omega + \alpha_1 r^2_{t-1} + \alpha_n r^2_{t-1} d_{t-1} + \delta_1 g_{t-1}$, all three with a skewed Student density distribution for the error term. We first give the empirical failure rates for the left (LQ) and right (RQ) quantiles at 10, 5 and 1%. A bold figure indicates that the empirical failure rate is significantly different (LR test or unconditional coverage) from the theoretical value. Secondly, *I* and *CC* give the P-values for the independence and conditional coverage tests respectively. Thirdly, the statistical tests (density forecast framework) $H0: \ldots$ are for $z_t = \beta_0 + \rho_1 z_{t-1} + \eta_1 r^2_{t-1} + u_t$, $u_t \sim N(0, \sigma^2)$, where z_t is equal to $\Phi^{-1}(\int_{-\infty}^{r_t/\sqrt{g_t}} f(u)\,du)$.

Table VIII
VaR results for the NASDAQ100 index (2/6/1997 - 31/3/2000)

Lagged implied volatility (VXN)						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	11.45	4.19	1.12	10.75	5.73	1.12
I	0.36	-	-	0.91	0.67	-
CC	0.29	-	-	0.79	0.62	-
H0: $\rho_1 = 0$; P-value = 0.071						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0						
H0: $\eta_1 = 0$; t-stat = 2.337						
RiskMetrics						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	11.73	5.73	0.84	11.31	6.15	0.70
I	0.96	0.67	-	0.40	0.63	-
CC	0.31	0.62	-	0.36	0.35	-
H0: $\rho_1 = 0$; P-value = 0.393						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0						
H0: $\eta_1 = 0$; t-stat = 1.899						
GJR-GARCH						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**13.55**	6.28	1.26	**13.83**	**8.24**	1.82
I	0.48	0.92	-	0.23	0.66	-
CC	0.01	0.31	-	0	0	-
H0: $\rho_1 = 0$; P-value = 0.181						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0						
H0: $\eta_1 = 0$; t-stat = 1.372						

Summary results for the back-testing of the out-of-sample VaR models on the NASDAQ100 index, 2/6/1997 - 31/3/2000. The first volatility specification (top panel) is $g_t = \omega + \eta\sigma^2_{imp,1,t-1}$, the second volatility specification (middle panel) is $g_t = 0.06r^2_{t-1} + 0.94g_{t-1}$ while the third volatility specification (bottom panel) is $g_t = \omega + \alpha_1 r^2_{t-1} + \alpha_n r^2_{t-1} d_{t-1} + \delta_1 g_{t-1}$, all three with a skewed Student density distribution for the error term. We first give the empirical failure rates for the left (LQ) and right (RQ) quantiles at 10, 5 and 1%. A bold figure indicates that the empirical failure rate is significantly different (LR test or unconditional coverage) from the theoretical value. Secondly, *I* and *CC* give the P-values for the independence and conditional coverage tests respectively. Thirdly, the statistical tests (density forecast framework) $H0: \ldots$ are for $z_t = \beta_0 + \rho_1 z_{t-1} + \eta_1 r^2_{t-1} + u_t$, $u_t \sim N(0,\sigma^2)$, where z_t is equal to $\Phi^{-1}(\int_{-\infty}^{r_t/\sqrt{g_t}} f(u)\,du)$.

Table IX
VaR results for the NASDAQ100 index (3/4/2000 - 31/1/2003)

Lagged implied volatility (VXN)						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**14.08**	5.92	**0.28**	9.44	5.35	1.27
I	0.34	0.76	-	0.48	0.19	-
CC	0	0.62	-	0.69	0.39	-
H0: $\rho_1 = 0$; P-value = 0.748						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0.034						
H0: $\eta_1 = 0$; t-stat = 1.367						

RiskMetrics						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**14.51**	6.20	0.42	9.01	4.65	1.13
I	0.72	0.22	-	0.59	0.27	-
CC	0	0.21	-	0.59	0.49	-
H0: $\rho_1 = 0$; P-value = 0.879						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0.385						
H0: $\eta_1 = 0$; t-stat = 1.445						

GJR-GARCH						
	LQ=10%	LQ=5%	LQ=1%	RQ=10%	RQ=5%	RQ=1%
$\widehat{f}$	**13.38**	5.92	0.56	10.14	5.63	1.41
I	0.22	0.71	-	0.90	0.02	-
CC	0.01	0.60	-	0.98	0.06	-
H0: $\rho_1 = 0$; P-value = 0.887						
H0: $\beta_0 = 0$ AND $\rho_1 = 0$ AND $\sigma^2 = 1$; P-value = 0.016						
H0: $\eta_1 = 0$; t-stat = 1.308						

Summary results for the back-testing of the out-of-sample VaR models on the NASDAQ100 index, 3/4/2000 - 31/1/2003. The first volatility specification (top panel) is $g_t = \omega + \eta\sigma^2_{imp,1,t-1}$, the second volatility specification (middle panel) is $g_t = 0.06r^2_{t-1} + 0.94g_{t-1}$ while the third volatility specification (bottom panel) is $g_t = \omega + \alpha_1 r^2_{t-1} + \alpha_n r^2_{t-1} d_{t-1} + \delta_1 g_{t-1}$, all three with a skewed Student density distribution for the error term. We first give the empirical failure rates for the left (LQ) and right (RQ) quantiles at 10, 5 and 1%. A bold figure indicates that the empirical failure rate is significantly different (LR test or unconditional coverage) from the theoretical value. Secondly, *I* and *CC* give the P-values for the independence and conditional coverage tests respectively. Thirdly, the statistical tests (density forecast framework) $H0 : \ldots$ are for $z_t = \beta_0 + \rho_1 z_{t-1} + \eta_1 r^2_{t-1} + u_t$, $u_t \sim N(0, \sigma^2)$, where z_t is equal to $\Phi^{-1}(\int_{-\infty}^{r_t/\sqrt{g_t}} f(u)\, du)$.

Notes

[1]See Poon and Granger (2003) for a recent and thorough review.

[2]See Alexander (2001) for a recent review and discussion of these models.

[3]See the CBOE website at http://www.cboe.com.

[4]See also Pagan and Schwert (1990) or Christensen and Prabhala (1998).

[5]In practice, daily returns can exhibit some weak autocorrelation. This usually involves fitting an AR, MA or ARMA structure on the r_t prior to the VaR analysis, along with a constant. The rest of the paper thus proceeds with r_t, which are the demeaned and uncorrelated returns (or residuals from the ARMA analysis).

[6]Combined with a GARCH-type volatility specification, Giot and Laurent (2002) show that the resulting VaR model provides excellent out-of-sample performance when back-tested against other models.

[7]The asymmetry coefficient $\xi > 0$ is defined such that the ratio of probability masses above and below the mean is $\frac{\Pr(\varepsilon \geq 0|\xi)}{\Pr(\varepsilon < 0|\xi)} = \xi^2$.

[8]$\Phi^{-1}(.)$ is the inverse of the standard normal distribution function.

[9]Note however that all models are used within the framework of a density distribution that allows skewness and fat tails.

[10]Histograms for the z_t (not reported in this paper) tell the same story.