Erasmus School of Economics

BACHELOR THESIS ECONOMETRICS AND OPERATIONS RESEARCH

The Convenience of Value-at-Risk Estimates to Construct Volatility Forecasts

Name student: NITIN SOEKHOE

Student ID number: 498200

Supervisor: BRAM VAN OS

Second assessor: DR. RUTGER-JAN LANGE

JULY 4, 2021

Abstract

In this paper, we investigate whether volatility forecasts can be produced by Value-at-Risk (VaR) forecasts. In particular, we examine the employment of VaR estimates obtained from the conditional autoregressive Value-at-Risk (CAViaR) models. An extensive overview of the volatility forecasting performance of various models and methods, including generalized autoregressive conditional heteroskedasticity (GARCH) models, is provided. As an extension, we create volatility forecasts with the incorporation of an uniformly spaced series of estimated quantiles instead of the interval between symmetric quantiles, as in Taylor (2005). The results show an outperformance of the established methods and, on top of that, the extension of involving the whole distributional pattern of returns led to an improvement of the volatility forecasts for all CAViaR models in terms of informational content.

KEYWORDS: VOLATILITY FORECASTING - VALUE-AT-RISK - CAVIAR MODELS

The views stated in this thesis are those of the author and not necessarily those of the supervisor, second assessor, Erasmus School of Economics or Erasmus University Rotterdam.

Contents

1 Introduction

Volatility models have become an essential tool for financial applications, including risk management and option pricing. Formally, volatility is defined as the standard deviation of logarithmic returns and measures the degree of price variation and uncertainty of a financial asset. Since volatility can have a crucial role in financial markets and portfolio management, numerous models have been proposed to account for volatility clustering, after (Engle, 1982) introduced the autoregressive conditional heteroskedasticity (ARCH) processes. This phenomenon represents the fact that great moments of uncertainty and stability alternate in the market. However, (generalized) ARCH models rely on the assumption that the shape of the conditional distribution is fixed over time while the volatility can vary. If the shape of the distribution (for example Student's t) in (G)ARCH models changes over time, the volatility forecasts of these models are harmed. On the contrary, the conditional autoregressive Value-at-Risk (CAViaR) models, proposed by Engle and Manganelli (2004), specify the behaviour of quantiles in an autoregressive framework such that no assumption has to be made about the conditional distribution of the model. Nonetheless, volatility estimates are not obtained directly from the CAViaR models and of interest is whether the CAViaR quantile forecasts can result in superior volatility forecasts. Taylor (2005) investigated this idea by creating volatility forecasts from one pair of symmetric quantile forecasts and showed an outperformance of the moving average and GARCH models. Furthermore, we examine whether the incorporation of the whole distributional pattern, with a wider set of quantiles, may lead to an improvement of the volatility forecasts. The method of Taylor, namely, can suffer from extreme distribution tails and is subject to an arbitrary choice of the quantile pair. The latter findings are summarized in the following two research questions:
RQ1: *'Can volatility forecasts from CAViaR models outperform GARCH and other established methods?'*
RQ2: *'Does the incorporation of the whole distributional pattern, instead of one pair of quantiles, deliver better volatility forecasts?'*

First, using 10 years of daily return data for five stock indices, we estimate and forecast volatility with three different methods: (i) moving average and GARCH models, (ii) the VaR-based methods, which generate volatility forecasts from quantile forecasts from the CAViaR models and (iii) quantile-based methods, where an evenly spaced series of estimated quantiles is used to produce volatility forecasts. Eventually, we produce 500 out-of-sample forecasts for each index for three different holding periods.

After evaluating the absolute and relative performance of the volatility forecasts, the results show that the CAViaR model which accounts for the 'leverage effect' outperforms all moving average, GARCH and other CAViaR models for the five stock indices. Moreover, including an uniformly spaced series of estimated quantiles to generate volatility forecasts led to an improvement of the informational content of the volatility forecasts for all CAViaR models.

The main contribution of this paper is that it provides an approach to forecast volatility of financial assets, which does not assume an explicit form for the underlying distribution. When the left and right tails of the conditional distribution vary over time, the method of Taylor still produces conditional volatilities based on a constant simple function of the interval between symmetric conditional quantiles. Our extension provides volatility forecasts with more complete information when these tails suffer from extreme behaviour or are arbitrarily chosen.

The remainder of this paper is organized as follows: in Section 2, we present a literature review of relevant studies conducted in the past, followed by a data description in Section 3. Next, we introduce the models and methods in Section 4. Afterwards, the results are presented in Section 5. Finally, we conclude and discuss our results in Section 6.

2 Literature

In this section, we provide the main findings of previous research that has been conducted on volatility forecasts as well as how we add to the existing literature.

For decades, volatility forecasting has been a key subject in financial literature. The volatility of a financial asset can be interpreted as its uncertainty and is defined as its standard deviation. Poon and Granger (2003) classify volatility forecast methods in four categories. Firstly, the moving average and exponentially weighted moving average methods incorporate previous standard deviations in determining volatility predictions. However, the driving force in volatility forecasting was already established in the 1980s. Bollerslev (1986), namely, introduced the GARCH class of models where lagged error terms and variances are included. Several modifications of the GARCH models have been proposed in literature, for example: IGARCH (Nelson, 1990), GJRGARCH (Glosten et al., 1993) and Beta-t-EGARCH (Harvey and Chakravarty, 2008). The GJRGARCH model, for instance, allows for asymmetry caused by the fact that negative returns have larger effect on volatility. Besides these two time-series methods, Poon and Granger (2003) thirdly discuss the volatility forecasts from the Black-Scholes option pricing model. Two drawbacks of this method are: (i) by choosing a different series of options, multiple forecasts can be derived for a single financial asset and (ii) financial assets may not have option derivatives. Fourthly and lastly, stochastic volatility forecasts (Ghysels et al., 1996) are found to be computationally hard and hence not widely used.

In this paper, we build on the fundamentals laid down by Pearson and Tukey (1965). They found that the ratio of the standard deviation to the interval between symmetric quantiles in the tails of the distribution is remarkably constant for different distributions. When quantile forecasts are obtained from CAViaR models (Engle and Manganelli, 2004), the results of Pearson and Tukey (1965) form a basis to construct volatility forecasts. The main advantage of our method is that if the left and right tails of the conditional distribution are driven by different forces over time, the conditional volatility can still be approximated by the interval between symmetric conditional quantiles, whereas GARCH models only use an autoregressive model for the variance with a fixed conditional distribution. Taylor (2005), who utilized Value-at-Risk estimates from CAViaR models, indeed found that volatility forecasts from one pair of symmetric quantile forecasts outperformed forecasts from GARCH models and moving average methods for 25 stock indices and individual stocks. The method of Taylor led to approach of Huang (2012) to include an uniformly spaced series of quantiles to construct volatility forecasts. He argues that by doing so, the volatility forecasts will not be subjected to overestimation or underestimation due to extreme distribution tails. Additionally, the arbitrary choice of the pair of quantiles does not have to be considered and more information is encompassed in the forecasting procedure. His results showed an improvement of the volatility forecasts compared to previous methods.

3 Data

The data set consists of the daily close-to-close log returns of five stock indices over a 10 year period ranging from April 1993 up to and including April 2003. The closing prices are obtained from the Yahoo Finance database. The close-to-close log return is defined by $r_t = 100 \times \log(y_t/y_{t-1})$, where y_t is the level of the index or stock at the end of day t. The five stock indices are: the French CAC 40, the German DAX 30, the Hongkongese HSI, the Japanese NIKKEI 225, and the U.S. S&P 500. Eventually, we end up with different numbers of log returns for different indices, as can be seen in the last column of Table 1, since not all markets were closed on the same days. We can furthermore confirm that the mean of the log returns is close to 0 for all indices. Evaluating the other descriptive statistics, we observe that all indices suffer from deviant skewness and large kurtosis. This indicates that the distribution of the returns is likely to be non-normal. The rejection of the null hypothesis of normality is confirmed by the Jarque-Bera test statistic in the second to last column of Table 1. The presence of non-normality and heavy tails advocates for the use of CAViaR models, which do not require distributional assumptions.

Table 1: Descriptive Statistics for the Daily Log Returns r_t

Index	Mean	St. Dev.	Skewness	Kurtosis	Jarque-Bera	Observations
CAC 40	0.017	1.467	-0.062	5.13	486**	2512
DAX 30	0.023	1.588	-0.128	5.80	827**	2518
HSI	0.008	1.831	0.102	11.29	7075**	2471
NIKKEI 225	-0.041	1.473	0.102	5.15	478**	2463
S&P 500	0.029	1.128	-0.107	6.46	1257**	2519

Note. ** $p < 0.01$.

4 Methodology

In this section, we provide the time-series methods to model and forecast volatility. More specifically, moving averages, GARCH models and CAViaR models are considered to describe the volatility's behaviour. Finally, we discuss the evaluation tools to compare the models' predictive ability.

Throughout the paper, we assume that the close-to-close log return r_t is generated as follows: $r_t = \mu + \varepsilon_t$ for $1 \leq t \leq T$, where T denotes the sample size, μ is the mean and ε_t can be seen as the "shock" term. If the conditional variance of r_t, $\mathrm{V}[r_t|\mathcal{I}_{t-1}]$, is denoted as σ_t^2, volatility is defined as σ_t, where $\mathcal{I}_{t-1}$ denotes the information set at time $t-1$.

Lastly, the estimation, forecasting and evaluation of all methods and models is conducted in MATLAB (version R2021a).

4.1 Moving Averages

The first procedure to estimate volatility is to consider a moving average (MA) of the past squared shocks:

$$\sigma_t^2 = \frac{1}{T}\sum_{i=1}^{T}(r_{t-i} - \bar{r})^2, \tag{1}$$

where $\bar{r}$ is the average return, $\frac{1}{T}\sum_{i=1}^{T} r_{t-i}$, and a rolling window of 30 days is applied. This rather simple approach includes few drawbacks. Firstly, σ_t^2 sharply increases whenever an extreme return occurs due to the fact that all observations are equally weighted. This phenomenon is known as the 'ghost feature'. Secondly, the length of the window T is subjective; too few observations will cause sampling errors whereas a large window causes rather slow reaction of the estimates to the true volatility. Thirdly, the specification in (1) assumes the variance of the returns in the window is the same.

To obviate these drawbacks, we furthermore examine the exponentially weighted moving average (EWMA) method to estimate volatility:

$$\sigma_t^2 = \lambda\sigma_{t-1}^2 + (1-\lambda)\varepsilon_{t-1}^2. \tag{2}$$

Despite of the fact that setting $\lambda = 0.94$ results in the well known RiskMetrics model (Longerstaey and Spencer, 1996), we choose to optimize the parameter value in (2) by minimizing the in-sample sum of squared deviations between the variance forecasts, $\hat{\sigma}_{t+1|t}^2$, and the squared error, ε_{t+1}^2.

The total volatility of the return from period $t+1$ until $t+h$ (multiperiod forecast), $\hat{\sigma}_{t,h|t}^2 = \sum_{i=1}^{h}\hat{\sigma}_{t+i|t}^2$, is obtained by simply multiplying the one-step ahead forecast, $\hat{\sigma}_{t+1|t}^2$, by h for both MA and EWMA methods (square-root-of-time rule).

4.2 GARCH Models

The GARCH models, introduced by Bollerslev (1986), are nowadays the most widely used statistical volatility models. The conditional variance in this class of models is explained by lagged conditional variance terms and lagged squared errors. The GARCH (1,1) model for example reads

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2, \tag{3}$$

where ω, α and β are parameters with $\omega > 0$, $\alpha > 0$ and $\beta \geq 0$ in order to ensure $\sigma_t^2 > 0$. It follows that variance for the sum of days up to some horizon h (multiperiod forecast), $\sum_{i=1}^{h}\hat{\sigma}_{t+i|t}^2$, is given by

$$\hat{\sigma}_{t,h|t}^2 = \frac{h\omega}{1-\alpha-\beta} + \left(\frac{1-(\alpha+\beta)^h}{1-\alpha-\beta}\right)\left(\hat{\sigma}_{t+1|t}^2 - \frac{\omega}{1-\alpha-\beta}\right), \tag{4}$$

where $\hat{\sigma}_{t+1|t}^2$ is the one-step-ahead volatility forecast.

In practice, it has been found that $\beta \approx 1-\alpha$ and the integrated GARCH (IGARCH) models (Nelson, 1990) incorporate the equality $\beta = 1-\alpha$, such that the multiperiod forecast of the IGARCH(1,1) model can be written as

$$\hat{\sigma}_{t,h|t}^2 = \frac{1}{2}h(h-1)\omega + h\hat{\sigma}_{t+1|t}^2. \tag{5}$$

The IGARCH models are estimated with the help of the MFE Toolbox [1].

[1] https://www.kevinsheppard.com/files/code/matlab/mfe-toolbox-documentation.pdf

Typically, negative shocks ε_t affect volatility more than positive shocks ('leverage effect'). The GJRGARCH model of Glosten et al. (1993) allows for such asymmetry and the GJRGARCH (1,1) model reads as follows:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2(1 - I[\varepsilon_{t-1} > 0]) + \gamma\varepsilon_{t-1}^2 I[\varepsilon_{t-1} > 0] + \beta\sigma_{t-1}^2, \tag{6}$$

where $\omega > 0$, $\alpha > 0$, $\gamma > 0$ and $\beta \geq 0$ in order to ensure $\sigma_t^2 > 0$ and the indicator function $I[A]$ equals one if event A happens and zero otherwise. With the assumption that the median of the distribution of ε_t is zero, which implies that the expectation of the indicator function in (6) is 0.5, the multiperiod forecast is given by

$$\hat{\sigma}_{t,h|t}^2 = \frac{h\omega}{1 - \frac{1}{2}(\alpha + \gamma) - \beta} + \left(\frac{1 - (\frac{1}{2}(\alpha + \gamma) + \beta)^h}{1 - \frac{1}{2}(\alpha + \gamma) - \beta}\right)\left(\hat{\sigma}_{t+1|t}^2 - \frac{\omega}{1 - \frac{1}{2}(\alpha + \gamma) - \beta}\right). \tag{7}$$

4.3 VaR-based Methods

Value-at-Risk (VaR) is the θ-th quantile of the distribution of the h-day return $r_{t+h,h}$ $(= r_{t+1} + r_{t+2} + ... + r_{t+h})$ and measures the down-side risk of $r_{t+h,h}$. VaR is defined as $Q_t(\theta, h)$, such that $\mathrm{P}[r_{t+h,h} \leq Q_t(\theta, h)] = \theta$. For example, if a stock portfolio has Value-at-Risk at 95% for the next 5 days of -2.8%, the probability that the return is lower than -2.8% is 5%.

4.3.1 CAViaR Models

Three different approaches of VaR estimation are mentioned in literature, as found by Manganelli and Engle (2001).

First, VaR is estimated with volatility forecasts obtained with, for example, GARCH estimation with a Student's t-distribution.

Second, nonparametric estimation does not require distributional assumptions for the reason that historical simulation estimates the VaR as the quantile of the empirical distribution of returns from a moving window of the most recent periods. Similar to the EWMA method in Section 4.1, Boudoukh et al. (1998) introduced such approach for quantiles. The so-called BRW method multiplies the past T returns with exponentially decreasing weights:

$$\left(\frac{1-\lambda}{1-\lambda^T}\right), \left(\frac{1-\lambda}{1-\lambda^T}\right)\lambda, \left(\frac{1-\lambda}{1-\lambda^T}\right)\lambda^2, ..., \left(\frac{1-\lambda}{1-\lambda^T}\right)\lambda^{T-1}. \tag{8}$$

Hereafter the returns are ordered in ascending order and the weights are summed until quantile θ is achieved. The θ-th quantile will then be equal to the return corresponding to the final weight in the summation. We work with the (arbitrary) value of 0.99 for λ, as experimented by Boudoukh et al.. Lastly, we used a moving window of one year for both the historical simulation and BRW methods.

Third, semiparametric methods include estimation by extreme value theory and quantile regression. The latter methods resulted in the conditional autoregressive Value-at-Risk (CAViaR) models, developed by Engle and Manganelli (2004). As mentioned before, no distributional assumptions have to be considered and Engle and Manganelli proposed the four following CAViaR models:

Symmetric Absolute Value CAViaR:

$$Q_t(\theta) = \omega + \alpha Q_{t-1}(\theta) + \beta|\varepsilon_{t-1}|. \tag{9}$$

Asymmetric Slope CAViaR:

$$Q_t(\theta) = \omega + \alpha Q_{t-1}(\theta) + \beta_1(\varepsilon_{t-1})^+ + \beta_2(\varepsilon_{t-1})^-. \tag{10}$$

Indirect GARCH (1,1) CAViaR:

$$Q_t(\theta) = (1 - 2I[\theta < 0.5])(\omega + \alpha Q_{t-1}(\theta)^2 + \beta\varepsilon_{t-1}^2)^{\frac{1}{2}}. \tag{11}$$

Adaptive CAViaR:

$$Q_t(\theta) = Q_{t-1}(\theta) + \alpha(\theta - I[\varepsilon_{t-1} \leq Q_{t-1}(\theta)]. \tag{12}$$

where $Q_t(\theta)$ ($= Q_t(\theta, 1)$) is the θ-th quantile, ω, α, β, β_1 and β_2 are parameters and $(x)^+ = \max(x,0)$ and $(x)^+ = -\min(x,0)$ in (10). The VaR estimate in t is reduced by the indicator function in (12) if the quantile estimate in $t-1$ is greater than the error in that period. Whereas the GARCH models involve the lagged squared error, the Symmetric Absolute Value CAViaR and the Asymmetric Slope CAViaR models incorporate the value of the regular error. Similar to the GJRGARCH models, the Asymmetric Slope CAViaR model is designed to accommodate for the 'leverage effect'.

To estimate the parameters of the CAViaR models, we follow Koenker and Bassett Jr (1978). The quantile regression minimization reads

$$\min\left(\sum_{t|r_t \geq Q_t(\theta)} \theta|r_t - Q_t(\theta)| + \sum_{t|r_t < Q_t(\theta)} (1-\theta)|r_t - Q_t(\theta)|\right). \tag{13}$$

We followed Engle and Manganelli (2004) in the CAViaR optimization routine. First, we generated 10^5 vectors of parameters from a uniform random number generator between 0 and 1. Second, we selected the 10 vectors that produced the lowest QR Sum in (13) and ran them in the Nelder-Mead simplex algorithm to acquire initial values for the quasi-Newton algorithm. This procedure was repeated until the convergence criterion was satisfied. Finally, the one producing the lowest QR Sum was chosen as the final parameter vector. We used the MATLAB codes provided by Engle and Manganelli for the estimation of the CAViaR models [2].

[2] http://www.simonemanganelli.org/Simone/Research.html

4.3.2 Volatilty Forecasts from VaR Forecasts

To obtain volatility forecasts from VaR forecasts, we firstly perform Ordinary Least Squares (OLS) (Heij et al., 2004) of the squared error (ε_t^2), which serves as proxy for the true variance, on the squared interval between the symmetric quantile estimates $\hat{Q}_{t+1}(1-\theta)$ and $\hat{Q}_{t+1}(\theta)$ with in-sample data. Once one-step-ahead quantile forecasts are secured from the CAViaR models, one-step-ahead variance forecasts are formed by

$$\hat{\sigma}^2_{t+1|t} = \hat{\alpha}_1 + \hat{\beta}_1(\hat{Q}_{t+1}(1-\theta) - \hat{Q}_{t+1}(\theta))^2, \tag{14}$$

where $\hat{\alpha}_1$ and $\hat{\beta}_1$ are in-sample OLS estimates. Unlike the moving averages method in Section 4.1, multiperiod forecasts cannot be computed by multiplying the one-step-ahead forecasts by horizon h due to the fact that the assumption of constant variance during h is not valid. The Adaptive CAViaR model is an exception on the latter. Analytical solutions for the other CAViaR models do no exist and forecasting by simulation is rather computationally hard. For this purpose, $\sum_{i=1}^{h} \varepsilon^2_{t+i}$ is regressed on $(\hat{Q}_{t+1}(1-\theta) - \hat{Q}_{t+1}(\theta))^2$. By assuming the shocks have constant conditional mean and no autocorrelation is present during horizon h, the realized multiperiod variance is approximated by the sum of h squared errors. Eventually, the expression of the multiperiod variance predictions reads

$$\hat{\sigma}^2_{t,h|t} = \hat{\alpha}_h + \hat{\beta}_h(\hat{Q}_{t+1}(1-\theta) - \hat{Q}_{t+1}(\theta))^2, \tag{15}$$

where $\hat{\alpha}_1$ and $\hat{\beta}_1$ are estimated by in-sample OLS regression of $\sum_{i=1}^{h} \varepsilon^2_{t+i}$ on $(\hat{Q}_{t+1}(1-\theta) - \hat{Q}_{t+1}(\theta))^2$.

Summarizing, after the parameters of the CAViaR models in Section 4.3.1 are estimated and forecasted quantiles are acquired, variance forecasts in (14) and (15) are computed with the OLS estimates $\hat{\alpha}_1$ and $\hat{\beta}_1$.

4.4 Quantile-based Methods

As an extension to the work of Taylor (2005), we propose three alternative ways to compute volatility forecasts from quantile estimates. Whereas Taylor uses one squared pair of quantiles to generate volatility forecasts, Huang (2012) employs an uniformly spaced series of quantiles for this purpose. By doing so, the two following drawbacks are obviated: (i) overestimation from extremely asymmetric tails and (ii) underestimation from highly symmetric tails. Moreover, the arbitrary choice of which quantile interval best represents the width of the distribution is avoided and a much more complete information set of the volatility is available compared to the approach of Taylor (2005).

One can see the forecasts from (15) in the following general form:

$$\hat{\sigma}^2_{t,h|t} = \hat{\alpha}_h + \hat{\beta}_h F(\hat{Q}_{t+1}(\theta, h)), \tag{16}$$

where $F(\cdot)$ is an unspecified function and $\hat{Q}_{t+1}(\theta)$ a vector of quantiles $\{\theta, 2\theta, ..., m\theta\}$, which are estimated at time $t+1$ ($\theta > 0$ and $m\theta < 1$) with one of the CAViaR models and where m controls for the number of quantiles employed. For $F(\cdot)$, Huang (2012) define the following three Standard Deviation (SD) functions:

$$\text{SD} : F(\cdot) = \left(\frac{1}{m-1} \sum_{m=1}^{99} (Q(0.01m) - \bar{Q})^2 \right)^{\frac{1}{2}}, \tag{17}$$

$$\text{Weighted SD (WSD)} : F(\cdot) = \left(\sum_{m=1}^{99} W(Q(0.01m) - \bar{Q})^2 \right)^{\frac{1}{2}}, \tag{18}$$

$$\text{Median SD (MSD)} : F(\cdot) = \left(\frac{1}{m-2} \sum_{m=1}^{99} (Q(0.01m) - Q(0.5))^2 \right)^{\frac{1}{2}}, \tag{19}$$

where $\bar{Q}$ is the conditional mean of all quantiles. By implementing these functions instead of the rather simplistic approach of the squared interval of quantiles, the movements of the quantiles not only reflect the tail behaviours, but also the whole distributional pattern. More specifically, the width of the distribution of the estimated quantiles around their conditional mean is considered. The WSD function in (18), for instance, applies a weight W to each squared deviation in the sum. We set W equal to $\theta/25$ when $\theta < 0.5$, and $(1-\theta)/25$ otherwise, such that that extreme quantiles will have a greater impact on the volatility forecasts and the often found kurtosis in return distributions of financial assets is accommodated for. Finally, the MSD function captures the irregularity of the return distributions as the median is less outlier sensitive than the mean.

4.5 Forecast Evaluation

To evaluate the volatility forecasting accuracy of the models, we firstly examine the absolute performance and relative performance of the volatility forecasts and secondly the quantile forecasts (which result in volatility forecasts) of the VaR-based and quantile-based methods with three metrics.

Firstly, we evaluate the volatility forecasts by the R^2 of the Mincer-Zarnowitz regression (Mincer and Zarnowitz, 1969), $\sum_{i=1}^{h} \varepsilon^2_{t+i} = \alpha_{f,h} + \beta_{f,h}\hat{\sigma}^2_{t,h|t} + u_{t,h}$. A higher R^2 indicates a larger explanatory power of the forecast relative to the proxy. Note that the estimated parameters in (14) and (15) do not influence the outcome of the Mincer-Zarnowitz regression, since a new constant and slope is estimated. Hence, as well for the VaR-based and the quantile-based methods, the R^2 reflects the informational content of the forecasts. Moreover, we perform an encompassing test, which examines whether a method outperforms another. This test is based on the following model:

$$\hat{\sigma}^2_{t,h|t} = w\hat{\sigma}^2_{At,h|t} + (1-w)\hat{\sigma}^2_{Bt,h|t} + e_{t,h}, \tag{20}$$

where $\hat{\sigma}^2_{t,h|t}$ is the realized multiperiod variance forecast and $\hat{\sigma}^2_{Mt,h|t}$ is variance prediction of method M, with $M = A, B$. Moreover, $e_{t,h}$ is the OLS error term and w is the weight estimated if OLS is performed of $(\hat{\sigma}^2_{t,h|t} - \hat{\sigma}^2_{At,h|t})$ on $(\hat{\sigma}^2_{At,h|t} - \hat{\sigma}^2_{Bt,h|t})$. A method is said to be encompassed by another if the weight of that method is zero. The test in (20) can be seen as if $\hat{\sigma}^2_{t,h|t}$ is a weighted average of two forecasts (Granger and Newbold, 1973; Chong and Hendry, 1986). The use of non-overlapping data is of great importance, for the reason that OLS regression with overlapping data conceivably encounters autocorrelation and hence an invalid test (Christensen and Prabhala, 1998).

Secondly, we follow Engle and Manganelli (2004) by examining the quantile forecasts through three following absolute measures: (i) hit rate, (ii) dynamic quantile (DQ) test statistic and (iii) QR Sum of (13). The hit rate captures the percentage of observations below the estimator. In the best case scenario, the hit rate equals θ. The DQ test for correct conditional coverage of Engle and Manganelli evaluates whether the hit variable, defined as $Hit_t = I[\varepsilon_t \leq \hat{Q}_t(\theta)] - \theta$, follows i.i.d. Bernoulli distribution with probability θ and furthermore is independent of the quantile estimator, $\hat{Q}_t(\theta)$. Under the null hypothesis of perfect conditional coverage, the DQ test statistic is $\chi^2(p+2)$ distributed, where p is the number of lags of Hit_t included. At last, the model with the lowest out-of-sample QR Sum is preferred.

5 Results

This section is split in two. First, we shortly discuss the outcomes of the estimation of the CAViaR models. Thereafter, the evaluation of the volatility and quantile forecasts is presented. For each stock index, we subtracted the in-sample unconditional mean, μ, from the daily log return, r_t and estimation and forecasting of all methods and models were performed on the resultant ε_t.

5.1 Estimation of CAViaR Models

In Table 2, for all five indices, the OLS parameters of the regression of ε_t^2 on the squared interval between symmetric quantiles, which are estimated by the Asymmetric Slope CAViaR model for the in-sample period for the three different intervals, are shown. These parameter are used to construct the one-step-ahead volatility forecasts as in (14). In only three out of 15 cases, the constant, α_1, is significantly different from 0 on a 5% significance level. Finally, the estimates of β_1 differ significantly from the values proposed by Pearson and Tukey (1965), which may indicate on a difference in volatility forecast accuracy between the approaches.

Table 2: OLS Regression Results of ε_t^2 on the Squared Interval Between Symmetric Quantiles Estimated by the Asymmetric Slope CAViaR Model for In-Sample Data

Interval	Parameter	CAC 40	DAX 30	HSI	NIKKEI 225	S&P 500	Pearson and Tukey values
0.98	α_1	-0.038	-0.182	-0.248	0.129	-0.034	0
		(0.123)	(0.120)	(0.302)	(0.188)	(0.085)	
	β_1	0.058**	0.068**	0.056**	0.068**	0.062**	$4.65^{-2} = 0.046$
		(0.004)	(0.003)	(0.003)	(0.005)	(0.004)	
0.95	α_1	-0.018	-0.046	-0.778*	0.113	0.036	0
		(0.122)	(0.114)	(0.318)	(0.168)	(0.083)	
	β_1	0.097**	0.099**	0.121**	0.092**	0.090**	$3.92^{-2} = 0.065$
		(0.006)	(0.005)	(0.006)	(0.007)	(0.006)	
0.90	α_1	0.064	0.223*	-0.736*	0.170	-0.042	0
		(0.114)	(0.103)	(0.295)	(0.165)	(0.089)	
	β_1	0.150**	0.137**	0.205**	0.155**	0.184 **	$3.25^{-2} = 0.095$
		(0.009)	(0.007)	(0.008)	(0.012)	(0.012)	

Note. Standard errors are in parentheses; * $p < 0.05$, ** $p < 0.01$.

The News Impact Curves (NIC) of Engle and Ng (1993) for the four different CAViaR models applied to the S&P 500 are shown in Figure 1. For estimated parameters and an (arbitrary) lagged Q_{t-1} of 1.645, the NIC quantifies the impact of changes in the past shock, ε_t, on Q_t. The past shocks have a symmetric impact on Q_t for both the Symmetric Absolute Value and the Indirect GARCH (1,1) CAViaR models. We furthermore notice the strong asymmetry in the NIC of the Asymmetric Slope CAViaR model, which suggests that negative returns have greater impact on the future VaR estimate. Finally, the NIC of the Adaptive CAViaR model shows how the indicator function in (12) also creates the incorporation of the 'leverage effect' as the curve falls sharply when the past shock increases in magnitude.

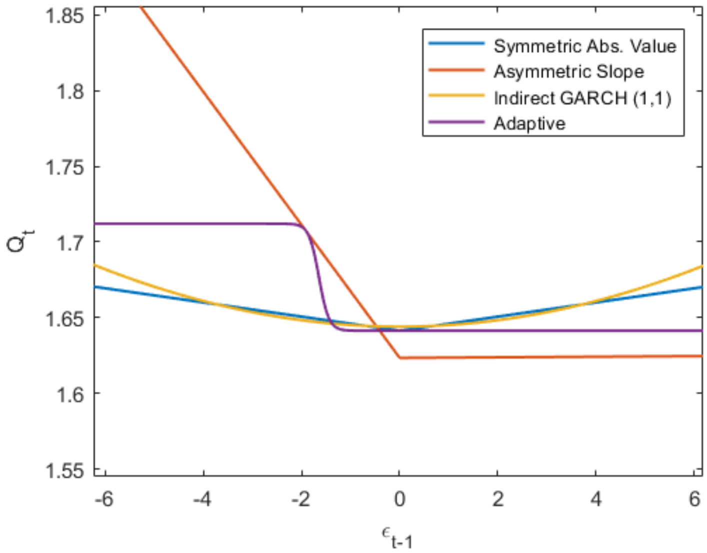

Figure 1: NIC for the S&P 500

The (normal and absolute) S&P 500 returns and the one-day-ahead volatility forecasts acquired from the Asymmetric Slope and the Adaptive CAViaR models for the 90% interval are plotted in Figure 2. The forecasts for the last 500 observations are post-sample and determined with in-sample estimated parameters. We observe that the Asymmetric Slope CAViaR model captures the volatility of the S&P 500 more adequately, especially in periods with large fluctuations and negative returns in the market.

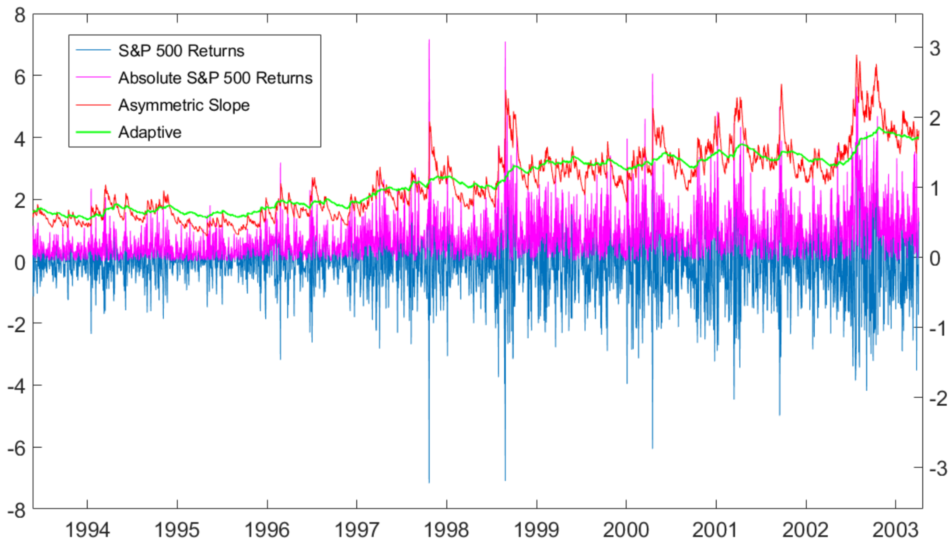

Figure 2: S&P 500 Returns Plotted on the Left Axis and the One-Day-Ahead Volatility Forecasts from two CAViaR Models for the 90% Interval Plotted on the Right Axis

5.2 Forecast Evaluation

For all methods and models, we generated 500 out-of-sample volatility forecasts for three different horizons h: one day, 10 days and 20 days.

5.2.1 Mincer-Zarnowitz Regression

The R^2 values of the Mincer-Zarnowitz regressions for all models and methods for the one-day and 10-day holding periods are presented in Tables 3 to 6. The results for the 20-day forecast horizon are summarized in Appendix A. In each column, the R^2 value of the best performing method is in bold and the last column shows the mean of the R^2 for all five stock indices.

First of all, the GJRGARCH (1,1) model outperforms the MA methods and the GARCH (1,1) and IGARCH (1,1) models for all indices and horizons. To account for non-normal, heavy tailed distributions, we moreover evaluated the GJRGARCH (1,1) model with two so-called winsorized data sets (Hoaglin et al., 2000). In the first winsorized data set with a rather simplistic approach, we replaced the largest (lowest) 1% of the in-sample observations with the 0.99-th (0.01-th) unconditional quantile. For the second data set, all in-sample observations larger (lower) than their corresponding in-sample conditional 0.99-th (0.01-th) quantile were set equal to this value (we refer to this method as CAViaR winsorized GJRGARCH (1,1) in the remainder of the paper). The conditional quantiles were obtained from the Asymmetric Slope CAViaR model. We find an outperformance of the CAViaR winsorized GJRGARCH (1,1) model for all indices and horizons compared to the GJRGARCH (1,1) model, hence proving the robustness of the incorporation of conditional quantiles.

Turning our attention to the VaR-based methods, we again notice that the model which accounts for the 'leverage effect' performs best, namely the Asymmetric Slope CAViaR model. For the 10-day and 20-day holding periods (Tables 4 and 9), the results show that the dominance of the Asymmetric Slope CAViaR model becomes even more visible. Although the Symmetric Absolute Value and the Indirect GARCH CAViaR models deliver decent R^2 values, they are not able to beat the GJRGARCH (1,1) model. Moreover, forecasts from historical simulation, BRW and the Adaptive CAViaR model do not seem promising regarding their relative low R^2 values. The rather low R^2 values for the Indirect GARCH (1,1) CAViaR model for the 98% interval for the S&P 500 can be explained by the fact that the VaR time series 'exploded', which even was not prevented by increasing the initial number of parameters or the amount of replications.

For the reason that the R^2 is not affected by the parameters in (14) and (15), we can state that, overall, the volatility forecasts from the Asymmetric Slope CAViaR model possess more informational content for all indices and horizons than any other model or method in our study.

Table 3: R^2 of the Mincer-Zarnowitz Regression of ε_{t+1}^2 on $\hat{\sigma}_{t+1|t}^2$ for One-Day-Ahead Volatility Forecasts

Model	**CAC 40**	**DAX 30**	**HSI**	**NIKKEI 225**	**S&P 500**	**Mean**
MA and GARCH methods						
Historical Volatility	10.55	8.76	3.33	0.89	4.86	5.68
EWMA	12.50	13.06	3.34	1.86	8.75	7.90
GARCH (1,1)	11.42	12.39	4.66	2.12	8.34	7.79
IGARCH (1,1)	10.79	12.15	4.61	1.98	8.08	7.52
GJRGARCH (1,1)	14.24	14.13	6.01	2.89	16.58	10.77
Simplistic winsorized GJRGARCH (1,1)	12.75	13.49	5.79	2.72	15.81	10.11
CAViaR winsorized GJRGARCH (1,1)	14.56	14.57	5.98	**2.95**	17.25	11.06
VaR-based methods						
Historical Simulation 98%	5.22	6.50	1.73	0.03	0.29	2.75
Historical Simulation 95%	3.55	5.09	2.14	0.04	0.18	2.20
Historical Simulation 90%	3.00	4.62	2.04	0.45	0.14	2.05
BRW 98%	4.27	3.09	0.73	1.65	1.74	2.30
BRW 95%	7.73	2.95	2.89	0.08	0.05	2.74
BRW 90%	5.80	4.69	3.13	0.17	0.03	2.76
Sym. Abs. Value 98%	11.10	12.51	4.71	1.58	9.51	7.88
Sym. Abs. Value 95%	11.29	12.59	5.01	1.61	8.97	7.89
Sym. Abs. Value 90%	10.39	12.77	4.94	1.41	8.73	7.65
Asym. Slope 98%	14.34	15.56	5.85	1.41	**20.24**	11.37
Asym. Slope 95%	14.34	16.62	**6.41**	2.26	20.13	**11.95**
Asym. Slope 90%	**15.75**	**17.06**	5.47	2.29	16.84	11.48
Indirect GARCH (1,1) 98%	12.20	13.09	4.09	2.30	0.56	6.45
Indirect GARCH (1,1) 95%	12.32	12.99	4.64	2.18	9.93	8.41
Indirect GARCH (1,1) 90%	11.46	13.04	4.59	1.92	9.90	8.18
Adaptive 98%	8.30	7.71	1.38	0.41	1.70	3.90
Adaptive 95%	7.40	8.76	2.58	0.62	6.52	5.18
Adaptive 90%	7.55	10.51	4.09	0.00	2.70	4.97

Note. R^2 in percentages.

Table 4: R^2 of the Mincer-Zarnowitz Regression of $\sum_{i=1}^{10} \varepsilon_{t+i}^2$ on $\hat{\sigma}_{t,10|t}^2$ for 10-Day Horizon Volatility Forecasts

Model	CAC 40	DAX 30	HSI	NIKKEI 225	S&P 500	Mean
MA and GARCH methods						
Historical Volatility	34.52	24.68	12.37	3.00	11.59	17.32
EWMA	44.53	38.09	26.37	7.64	22.44	27.81
GARCH (1,1)	39.45	36.09	23.06	8.79	21.20	25.32
IGARCH (1,1)	36.54	35.21	22.51	8.10	20.54	24.58
GJRGARCH (1,1)	50.40	41.13	33.77	12.13	42.47	35.98
Simplistic winsorized GJRGARCH (1,1)	43.85	39.22	30.21	11.68	40.98	33.19
CAViaR winsorized GJRGARCH (1,1)	52.91	42.73	34.37	12.64	45.58	37.65
VaR-based methods						
Historical Simulation 98%	18.31	22.20	8.36	0.62	0.18	9.93
Historical Simulation 95%	13.15	17.60	10.00	0.07	0.06	8.18
Historical Simulation 90%	11.39	16.87	9.04	0.72	0.12	7.63
BRW 98%	16.46	11.34	3.12	2.90	3.38	7.78
BRW 95%	35.07	10.27	16.54	1.29	1.47	12.93
BRW 90%	23.46	17.14	17.80	0.07	0.51	11.80
Sym. Abs. Value 98%	39.41	38.94	29.11	10.13	25.81	28.68
Sym. Abs. Value 95%	40.51	37.93	28.84	10.11	23.61	28.20
Sym. Abs. Value 90%	35.99	38.75	25.81	8.27	23.02	26.37
Asym. Slope 98%	52.92	46.36	39.52	11.23	**53.57**	39.51
Asym. Slope 95%	51.52	49.93	**41.43**	15.63	53.02	**42.31**
Asym. Slope 90%	**56.60**	**51.62**	40.45	**15.67**	44.78	41.82
Indirect GARCH (1,1) 98%	42.77	38.20	27.93	11.04	2.32	24.45
Indirect GARCH (1,1) 95%	43.65	37.86	26.27	10.42	27.26	29.09
Indirect GARCH (1,1) 90%	39.18	38.01	23.68	8.02	27.19	27.22
Adaptive 98%	27.36	22.28	5.86	2.32	4.04	12.37
Adaptive 95%	24.01	26.40	10.98	2.71	17.39	16.30
Adaptive 90%	25.57	32.17	16.37	0.62	6.23	16.19

Note. R^2 in percentages.

On top of that, the extension of the work of Taylor (2005) by generating forecasts as in Section 4.4, increases the explaining power for all of the CAViaR models. Tables 5, 6 and 10 summarize the R^2 values of the proposed Mincer-Zarnowitz regressions for the quantile-based methods for the three horizons. For each CAViaR model, volatility forecasts are generated as in (16) with the three functions in (17), (18) and (19). We indeed conclude that the incorporation of the whole distribution pattern improves the volatility forecasts, since all models experience an enlargement of the explaining power of the forecasts. Once more, we find that the volatility forecasts with VaR forecasts of the Asymmetric Slope CAViaR model have the highest R^2 values for all indices and horizons. Hence, throughout the paper, we have seen multiple times that the 'leverage effect' is present in the stock indices. Lastly, the difference in volatility forecast performance between the SD, WSD and MSD functions is minimal. To take into consideration is the fact that for the quantile-based methods, we encountered that for some indices and CAViaR models a few VaR time series 'exploded'. In Appendix B, the exact number of 'exploded' quantile time series can be found for all indices and CAViaR models. We observed that for especially the Indirect GARCH (1,1) CAViaR model and the S&P 500 index incorrect time series were delivered. However, taking into account the majority of the quantiles were estimated correctly and still containing substantially more information than the single interval between symmetric quantiles of Taylor, we were convinced the volatility forecasts, produced with the remaining quantiles, to possess enough distributional information.

Table 5: R^2 of the Mincer-Zarnowitz Regression ε^2_{t+1} on $\hat{\sigma}^2_{t+1|t}$ for One-Day-Ahead Volatility Forecasts

Model	**CAC 40**	**DAX 30**	**HSI**	**NIKKEI 225**	**S&P 500**	**Mean**
Quantile-based methods						
SD Sym. Abs. Value	11.74	12.49	5.07	1.65	9.28	8.05
WSD Sym. Abs. Value	11.75	13.23	5.08	1.66	9.31	8.21
MSD Sym. Abs. Value	11.77	12.45	5.06	1.65	9.38	8.06
SD Asym. Slope	16.06	**17.21**	**6.45**	**2.26**	17.46	**11.89**
WSD Asym. Slope	15.98	17.17	6.43	2.25	**17.53**	11.87
MSD Asym. Slope	**16.10**	**17.21**	6.43	2.23	17.37	11.87
SD Indirect GARCH (1,1)	12.57	13.77	4.86	0.88	9.95	8.41
WSD Indirect GARCH (1,1)	12.58	13.76	4.86	0.72	9.90	8.36
MSD Indirect GARCH (1,1)	12.77	13.76	4.85	0.89	10.09	8.47
SD Adaptive	7.93	9.37	3.51	0.48	3.47	4.95
WSD Adaptive	7.90	9.37	3.37	0.48	3.50	4.92
MSD Adaptive	8.30	10.25	3.79	0.52	3.86	5.34

Note. R^2 in percentages.

Table 6: R^2 of the Mincer-Zarnowitz Regression of $\sum_{i=1}^{10} \varepsilon_{t+i}^2$ on $\hat{\sigma}_{t,10|t}^2$ for 10-Day Horizon Volatility Forecasts

Model	CAC 40	DAX 30	HSI	NIKKEI 225	S&P 500	Mean
Quantile-based methods						
SD Sym. Abs. Value	41.17	41.19	28.39	9.90	26.14	29.36
WSD Sym. Abs. Value	41.18	42.67	28.52	9.98	26.16	29.70
MSD Sym. Abs. Value	41.21	41.17	28.44	9.86	26.13	29.36
SD Asym. Slope	58.23	**52.89**	41.35	**16.51**	50.62	43.92
WSD Asym. Slope	57.94	52.72	41.27	16.41	**50.94**	43.86
MSD Asym. Slope	**58.33**	52.88	**41.38**	16.32	50.80	**43.94**
SD Indirect GARCH (1,1)	44.12	41.39	26.86	4.81	27.75	28.99
WSD Indirect GARCH (1,1)	44.13	41.37	27.01	4.06	27.58	28.83
MSD Indirect GARCH (1,1)	45.34	41.38	26.89	4.83	27.95	29.28
SD Adaptive	26.68	28.71	14.95	3.45	8.26	16.41
WSD Adaptive	26.58	28.50	14.33	3.35	8.31	16.21
MSD Adaptive	27.92	31.29	16.18	3.61	9.28	17.66

Note. R^2 in percentages.

5.2.2 Encompassing Test

In Table 7, the results of the encompassing test are shown, where we chose the 90% interval Asymmetric Slope CAViaR model for A and the GJRGARCH (1,1) model (with non-winsorized data set) for B in (20). The estimate of w and the p-values for the null hypotheses H_0: $w = 1$ and H_0: $w = 0$ are presented for each index and forecast horizon. In five out of 15 cases, $w = 1$ is rejected at 10% significance level, indicating that we can reject that the Asymmetric Slope CAViaR model encompasses the GJRGARCH (1,1) model. Especially for the HSI, we find that the GJRGARCH (1,1) model outperforms the Asymmetric Slope CAViaR model, which is only in line for the one-day-ahead horizon in Table 3. In addition, the encompassing of the Asymmetric Slope CAViaR model by the GJRGARCH (1,1) model for the one-step-ahead forecasts of the NIKKEI 225 corresponds to the higher R^2 value of the GJRGARCH (1,1) model compared to the VaR-based method ($2.89 > 2.29$). At last, in six out of 15 cases, we cannot reject that the GJRGARCH (1,1) model encompasses the Asymmetric Slope CAViaR model, due to the inability to reject $w = 0$ at 10% significance level.

Table 7: Encompassing Test Results for the 90% Interval Asymmetric Slope CAViaR Model and the GJRGARCH (1,1) Model

	CAC 40	**DAX 30**	**HSI**	**NIKKEI 225**	**S&P 500**
One-step-ahead					
$\hat{w}$	1.17	1.23	0.09	0.46	0.37
p-value for H_0: $w = 1$	0.78	0.83	0.02	0.16	0.03
H_1: $w < 1$					
p-value for H_0: $w = 0$	0.01	0.01	0.31	0.20	0.07
H_1: $w > 0$					
10-day holding period					
$\hat{w}$	1.56	1.63	0.09	1.26	0.62
p-value for H_0: $w = 1$	0.93	0.94	0.02	0.67	0.08
H_1: $w < 1$					
p-value for H_0: $w = 0$	0.01	0.01	0.34	0.07	0.16
H_1: $w > 0$					
20-day holding period					
$\hat{w}$	1.93	1.43	0.10	1.23	0.60
p-value for H_0: $w = 1$	0.93	0.77	0.02	0.63	0.20
H_1: $w < 1$					
p-value for H_0: $w = 0$	0.02	0.04	0.31	0.09	0.13
H_1: $w > 0$					

Note. Test uses non-overlapping multiperiod forecasts in the post-sample period.

5.2.3 Quantile Forecast Evaluation

The evaluation of the quantile forecasts, constructed from the GJRGARCH (1,1) and Asymmetric Slope CAViaR models, for the 0.05 and 0.95-th quantiles and the 0.01 and 0.99-th quantiles is reported in Tables 8 and 12, respectively (Table 12 can be found in Appendix C). The values in bold point out the best model for each evaluation metric. For the DQ test, we included five lags of Hit_t.

The results in Table 8 display great outperformance of the Asymmetric Slope CAViaR model compared to the GJRGARCH (1,1) model. Both hit rate and correct conditional coverage are in favour of the the VaR-based method. These conclusions are in line with the results in Tables 3, 4 and 9. For the 0.01 and 0.99-th quantiles in Table 12, the two models perform similarly in terms of hit rate and correct conditional coverage. Nonetheless, the GJRGARCH (1,1) model shows more appealing QR Sums. Looking back at Table 3, we indeed see that forecasts of the GJRGARCH (1,1) model deliver higher R^2 values for the HSI and the NIKKEI 225.

Table 8: 500 One-Day-Ahead 0.05 and 0.95-th Quantile Forecasts Evaluation

	CAC 40		**DAX 30**		**HSI**		**NIKKEI 225**		**S&P 500**	
	0.05	0.95	0.05	0.95	0.05	0.95	0.05	0.95	0.05	0.95
Hit %										
GJRGARCH (1,1)	10.00	93.20	10.40	92.20	**6.20**	**94.20**	7.20	**93.60**	6.60	94.60
Asym. Slope	**4.40**	**95.00**	**6.80**	**94.40**	3.20	96.00	**5.40**	93.20	**4.40**	**95.20**
DQ										
GJRGARCH (1,1)	0.00	0.28	0.00	0.11	**0.74**	0.39	0.13	**0.40**	0.20	0.62
Asym. Slope	**0.98**	**0.66**	**0.48**	**0.14**	0.53	**0.83**	**0.46**	0.09	**0.51**	**0.95**
QR										
GJRGARCH (1,1)	98.40	106.64	112.56	122.17	**73.02**	**73.72**	**85.01**	**95.04**	**71.31**	76.04
Asym. Slope	**97.65**	**100.61**	**111.40**	**109.17**	76.04	74.92	85.77	95.07	72.77	**74.95**

In the end, we may say that our results are in line with those of Taylor (2005). For instance, the Asymmetric Slope CAViaR model achieved best results in both our papers. At the same time, one may encounter some dissimilarities when reading both papers, for example the intercept values in Table 2, the scale of the axes in Figures 1 and 2 and the order of magnitude of the QR Sums in Tables 8 and 12. Nonetheless, this is explainable due to the fact that we conducted our research with log returns in percentages, whereas Taylor worked with 'standard' log returns. Moreover, for the DAX 30, the results of the encompassing test in Table 7 are different. In view of the fact that the values for the skewness of the DAX 30 in Table 1 and in the paper of Taylor show great dissimilarity, there might be a difference in data.

6 Conclusion and Discussion

In the last decades, the surge of the importance of volatility of financial assets in numerous applications has not slowed down. Consequently, volatility forecasting methods have drawn great attention from both academic researchers and industrial practitioners. However, many models (such as GARCH) rely on distributional assumptions, which may harm the quality of the volatility forecasts. Therefore, this paper attempts to answer the following two research questions:

RQ1: *'Can volatility forecasts from CAViaR models outperform GARCH and other established methods?'*

RQ2: *'Does the incorporation of the whole distributional pattern, instead of one pair of quantiles, deliver better volatility forecasts?'*

This study analyzes these research questions by the following methods. Firstly, we consider volatility forecasts from the established moving average and GARCH models. Next, volatility forecasts are obtained from the interval between symmetric quantiles estimates, from the CAViaR models of Engle and Manganelli (2004). Lastly, an evenly spaced series of estimated quantiles is used to construct volatility forecasts.

We concluded that, after performing the Mincer-Zarnowitz regression and an encompassing test to evaluate the forecasts, firstly the Asymmetric Slope CAViaR model, which was designed to account for the fact that negative returns have a greater impact on volatility, overall, outperforms all established methods and other CAViaR models.

Regarding **RQ2**, we indeed found that, for all CAViaR models, including an evenly spaced series of estimated quantiles to produce volatility forecasts led to greater informational content of the forecasts compared to the case of one pair of symmetric quantiles .

Nonetheless, our study has several limitations and interesting areas to consider in future research. In particular, we encountered that, during the optimization routine, some Value-at-Risk (VaR) time series 'exploded'. Next, the realized variance was computed by summing squared returns, however this approach delivers a rather 'noisy' proxy. Alternatively, the realized volatility, which is based on the high-frequency data of intra-day returns, may be used as a more reliable proxy. Nevertheless, when consulting the Oxford library, the realized volatility measure was not available fully for all indices for the concerned period of our data set. In addition, the VaR-based and quantile-based methods to generate volatility forecasts were two-stage procedures. First, the parameters for the CAViaR models were estimated and afterwards the parameters to construct the volatility forecasts. In future research, one may extend the CAViaR models to the bivariate case where these processes are modelled together in order to capture their potential interactions. Finally, of great interest would be to estimate higher moments of the returns distribution, such as skewness, kurtosis or the covariance, with the help of VaR series (Brandt and Diebold, 2003).

References

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3):307–327.

Boudoukh, J., Richardson, M., and Whitelaw, R. (1998). The best of both worlds. *Risk*, 11(5):64–67.

Brandt, M. W. and Diebold, F. X. (2003). A no-arbitrage approach to range-based estimation of return covariances and correlations. Technical report, National Bureau of Economic Research.

Chong, Y. Y. and Hendry, D. F. (1986). Econometric evaluation of linear macro-economic models. *The Review of Economic Studies*, 53(4):671–690.

Christensen, B. J. and Prabhala, N. R. (1998). The relation between implied and realized volatility. *Journal of financial economics*, 50(2):125–150.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the econometric society*, pages 987–1007.

Engle, R. F. and Manganelli, S. (2004). CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of business & economic statistics*, 22(4):367–381.

Engle, R. F. and Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *The journal of finance*, 48(5):1749–1778.

Ghysels, E., Harvey, A. C., and Renault, E. (1996). Stochastic volatility. *Handbook of statistics*, 14:119–191.

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. *The journal of finance*, 48(5):1779–1801.

Granger, C. W. and Newbold, P. (1973). Some comments on the evaluation of economic forecasts. *Applied Economics*, 5(1):35–47.

Harvey, A. C. and Chakravarty, T. (2008). Beta-t-(E)GARCH. *Unpublished manuscript.*

Heij, C., de Boer, P., Franses, P. H., Kloek, T., van Dijk, H. K., et al. (2004). *Econometric methods with applications in business and economics.* Oxford University Press.

Hoaglin, D. C., Mosteller, F., and Tukey, J. W. (2000). *Understanding robust and exploratory data analysis*. Number Sirsi) i9780471384915.

Huang, A. Y. (2012). Volatility forecasting by quantile regression. *Applied Economics*, 44(4):423–433.

Koenker, R. and Bassett Jr, G. (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, pages 33–50.

Longerstaey, J. and Spencer, M. (1996). Riskmetricstm - technical document. *Morgan Guaranty Trust Company of New York: New York*, 51:54.

Manganelli, S. and Engle, R. F. (2001). Value at risk models in finance.

Mincer, J. A. and Zarnowitz, V. (1969). The evaluation of economic forecasts. In *Economic forecasts and expectations: Analysis of forecasting behavior and performance*, pages 3–46. NBER.

Nelson, D. B. (1990). Stationarity and persistence in the GARCH (1, 1) model. *Econometric theory*, pages 318–334.

Pearson, E. S. and Tukey, J. W. (1965). Approximate means and standard deviations based on distances between percentage points of frequency curves. *Biometrika*, 52(3/4):533–546.

Poon, S.-H. and Granger, C. W. (2003). Forecasting volatility in financial markets: A review. *Journal of economic literature*, 41(2):478–539.

Taylor, J. W. (2005). Generating volatility forecasts from value at risk estimates. *Management Science*, 51(5):712–725.

Appendix A Mincer-Zarnowitz Regression for 20-Day Horizon

Table 9: R^2 of the Mincer-Zarnowitz Regression of $\sum_{i=1}^{20} \varepsilon_{t+i}^2$ on $\hat{\sigma}_{t,20|t}^2$ for 20-Day Horizon Volatility Forecasts

Model	**CAC 40**	**DAX 30**	**HSI**	**NIKKEI 225**	**S&P 500**	**Mean**
MA and GARCH methods						
Historical Volatility	28.78	22.28	7.44	1.21	9.68	13.88
EWMA	37.08	31.77	23.93	5.32	17.05	23.03
GARCH (1,1)	32.47	29.91	17.44	6.54	15.85	20.44
IGARCH (1,1)	29.90	29.57	16.89	5.79	15.50	19.53
GJRGARCH (1,1)	42.99	33.36	29.08	11.45	33.64	30.10
Simplistic winsorized GJRGARCH (1,1)	36.67	31.97	24.95	11.09	32.45	27.43
CAViaR winsorized GJRGARCH (1,1)	45.43	34.48	29.93	12.29	36.22	31.67
VaR-based methods						
Historical Simulation 98%	16.03	23.63	9.75	1.29	0.02	10.13
Historical Simulation 95%	12.77	18.05	11.42	0.28	0.11	8.53
Historical Simulation 90%	10.14	17.68	9.10	1.06	0.00	7.60
BRW 98%	20.29	14.93	2.66	1.43	2.87	8.44
BRW 95%	38.64	11.15	17.42	0.59	3.26	14.21
BRW 90%	23.01	17.34	20.89	0.08	1.05	12.47
Sym. Abs. Value 98%	33.55	33.19	25.34	8.66	17.99	23.75
Sym. Abs. Value 95%	34.59	32.61	24.67	8.51	16.08	23.29
Sym. Abs. Value 90%	30.30	32.81	21.44	6.29	16.36	21.44
Asym. Slope 98%	46.62	40.60	37.99	10.61	**41.37**	35.06
Asym. Slope 95%	44.81	42.88	39.95	16.97	40.89	37.10
Asym. Slope 90%	**50.24**	**44.16**	**41.44**	**17.21**	34.48	**37.51**
Indirect GARCH (1,1) 98%	35.52	31.85	24.33	9.67	2.64	20.80
Indirect GARCH (1,1) 95%	36.27	31.75	21.28	8.81	18.98	23.42
Indirect GARCH (1,1) 90%	32.33	31.79	18.60	5.77	20.06	21.71
Adaptive 98%	22.66	21.25	5.15	2.45	2.35	10.77
Adaptive 95%	20.63	25.06	8.57	2.24	14.32	14.16
Adaptive 90%	22.54	30.13	14.70	0.30	4.23	14.38

Note. R^2 in percentages.

Table 10: R^2 of the Mincer-Zarnowitz Regression of $\sum_{i=1}^{20} \varepsilon_{t+i}^2$ on $\hat{\sigma}_{t,20|t}^2$ for 20-Day Horizon Volatility Forecasts

Model	**CAC 40**	**DAX 30**	**HSI**	**NIKKEI 225**	**S&P 500**	**Mean**
Quantile-based methods						
SD Sym. Abs. Value	35.31	39.73	25.32	7.46	19.55	25.47
WSD Sym. Abs. Value	35.31	40.02	25.48	7.54	19.47	25.56
MSD Sym. Abs. Value	35.39	39.74	25.28	7.43	19.39	25.44
SD Asym. Slope	52.32	47.54	41.91	**17.06**	41.58	40.08
WSD Asym. Slope	52.02	47.40	41.83	16.88	41.92	40.01
MSD Asym. Slope	**52.45**	**47.56**	**41.96**	16.80	**42.00**	**40.15**
SD Indirect GARCH (1,1)	37.12	35.96	22.72	9.00	21.10	25.18
WSD Indirect GARCH (1,1)	37.22	35.94	22.90	8.13	20.96	25.03
MSD Indirect GARCH (1,1)	38.57	35.94	22.78	9.02	21.07	25.48
SD Adaptive	22.69	27.37	13.09	2.68	5.39	14.24
WSD Adaptive	22.59	26.94	12.55	2.61	5.41	14.02
MSD Adaptive	23.84	29.78	14.15	2.82	6.20	15.36

Note. R^2 in percentages.

Appendix B Number of 'Exploded' VaR Time Series

Table 11: Number of 'Exploded' VaR Time Series for each Index and CAViaR Model for the Quantile-based Methods

	CAC 40	**DAX 30**	**HSI**	**NIKKEI 225**	**S&P 500**
Sym. Abs Value	0	0	0	0	3
Asym. Slope	2	0	0	3	23
Indirect GARCH (1,1)	21	6	4	16	17
Adaptive	0	0	0	0	0

Appendix C Quantile Forecast Evaluation

Table 12: 500 One-Day-Ahead 0.01 and 0.99-th Quantile Forecasts Evaluation

	CAC 40		**DAX 30**		**HSI**		**NIKKEI 225**		**S&P 500**	
	0.01	0.99	0.01	0.99	0.01	0.99	0.01	0.99	0.01	0.99
Hit %										
GJRGARCH (1,1)	2.00	**97.20**	2.60	97.80	**1.00**	**99.00**	**1.20**	98.60	**1.40**	98.80
Asym. Slope	**0.60**	97.00	**1.00**	**98.20**	0.40	99.40	0.20	**98.80**	2.00	**99.00**
DQ (p-value)										
GJRGARCH (1,1)	**0.05**	**0.00**	0.00	0.16	**0.99**	**0.95**	**0.99**	0.00	**0.84**	0.02
Asym. Slope	0.00	**0.00**	**0.01**	**0.32**	0.97	**0.95**	0.83	**0.01**	0.02	**1.00**
QR										
GJRGARCH (1,1)	**26.08**	31.25	**27.82**	34.58	**22.52**	**18.33**	**21.57**	**26.59**	**19.87**	**18.36**
Asym. Slope	29.34	**25.20**	28.63	**28.27**	23.60	20.49	28.84	26.78	20.13	24.01