

Norges miljø- og
biovitenskapelige
universitet

Master's Thesis 2016 30 ECTS
Norwegian University of Life Sciences
Faculty of Social Sciences
School of Economics and Business

Value-at-Risk forecasting with different quantile regression models

Øyvind Alvik
Master in Business Administration

Acknowledgments

I would like to thank my supervisor Sjur Westgaard for his time and help in the process of writing this thesis during the spring term, 2016.

I would also like to thank my significant other, Veronica, for being supportive and for taking good care of our one-year-old daughter Emilia at times when the latter has insisted on getting involved in the thesis writing.

Oslo, May, 2016

Abstract

Forecasting volatility and Value-at-Risk (VaR) are popular topics of study in econometrical finance. Their popularity can likely be attributed to the statistical challenges related to producing reliable VaR estimates across a wide array of assets and data series. As many financial assets offer unique statistical properties, it has proven to be a difficult task to find a model reliable enough to be considered accepted as the best method.

This study focuses on the problem of forecasting volatility and one-day-ahead VaR. The thesis has two main purposes. Firstly, we want to further examine the performance of linear quantile regression models in VaR forecasting against more established models as benchmarks. Secondly, we want to compare the performance between each of the three quantile regression models to see which one performs the best. The three quantile regression models in question are HAR-QR, EWMA-QR and GARCH(1,1) QR.

Our findings strongly support the conclusion that quantile regression outperforms the three benchmark models in predicting one-day-ahead VaR for all of the five assets examined. When subjected to coverage tests for both unconditional and conditional coverage each quantile regression delivered perfect unconditional coverage. However, only the HAR-QR model delivered perfect conditional coverage and thus performed the best of the three models.

The benchmarks models RiskMetrics, GARCH(1,1) and Historical Simulation showed particular problems with estimating the left tail quantiles of the distribution. The study shows that compared to the QR approach, these models fail to capture time variant volatility and the negative skewness and leptokurtosis that is present in most of the assets return distributions.

Contents

- Abstract 1
- 1 Introduction 2
- 2 Literature review 4
- 3 Data and Descriptive Statistics..... 6
- 4 Methodology 13
 - 4.1 Value-at-risk models..... 13
 - 4.1.2 RiskMetrics 13
 - 4.1.3 GARCH(1,1) 14
 - 4.2 Historical Simulation (HS) 15
 - 4.3 Quantile regression (QR)..... 15
 - 4.4 HAR-QREG (HeterogenousAutoregressive Quantile Regression Model) 16
 - 4.5 Backtesting procedures 16
- 5 Results 18
 - 5.1 RiskMetrics and GARCH (1,1) 21
 - 5.2 Historical Simulation (HS) 22
 - 5.3 Quantile Regression models 22
- 6 Conclusion..... 24
- References 26
- Appendix 28

1 Introduction

In finance does risk refer to measurement of uncertainty. This uncertainty lies in what the price of an investment or asset (such as equities, bonds, commodities, swaps, etc.) will be in the future, i.e how the price will change. The price fluctuations represent the financial returns, and how returns vary over time is referred to as the assets volatility and is measured by the standard deviation. Investors are often less concerned with achieving higher than expected gains than they are about mitigating unexpectedly large but feasible losses, which forms the basis for risk assessment. Therefore, Value-at-Risk (VaR) modeling is an important technique used to measure and quantify the risk associated with a particular investment.

VaR can be defined as the maximum expected loss expressed in percentage or nominal terms that will occur with a certain degree of certainty over some given time period. VaR can also be expressed as a measure of the loss an investor can be certain will *not* occur with a certain degree of certainty over that same time horizon.

An accurate estimate for VaR is important to investors or financial institutions in order to provide a more reliable assessment of risk. There are however several challenges related to the modelling and forecasting of reliable VaR estimates. The basic parametric VaR models do for example assume the returns are normally distributed and the volatility is constant. However, have numerous studies showed that empirical returns for financial assets more often than not exhibit skewed probability distributions with leptokurtosis, volatility clustering and leverage effects:

Leptokurtosis means a distribution that has fatter tails and higher peak than the normal distribution. This implicates the distribution has more observations further out on the sides, i.e the probability for extreme outcomes is greater than what would be the case if the returns were in fact normal distributed.

Volatility clustering refers to the tendency of volatility in financial time series to be time variant. Periods with large price movements are often followed by periods with more large movements, and vice versa for tranquil periods. Thus markets go through cycles with the implications that volatility is non-constant and found to be clustering together in periods.

Leverage effect is a term that describes how volatility for some assets such as equities has a tendency to increase more in the wake of large price falls compared to the periods after price

increases of the same magnitude. Commodities in general exhibit the opposite asymmetry, where there usually is a positive correlation between returns and volatility (Alexander 2008).

These properties of financial returns do forecasting of VaR more difficult than it would in a world where financial returns are normally distributed and volatility is constant. In an attempt to produce reliable VaR forecasts, several different approaches for VaR estimation have been proposed. In addition to the parametric models, the Historical Simulation (HS) and Monte Carlo Simulation models are common approaches associated with VaR analysis. In more recent quantile regression methods have emerged as a seemingly robust method for VaR forecasting. Quantile regression will be discussed in further detail in the Literature review section. When provided with independent variables that take time variant volatility into account, quantile regression models have delivered very reliable results.

The purpose of this thesis is to further test the ability of quantile regression models to produce reliable daily VaR estimates. In order to perform the test ,three different quantile regression models will be implemented to forecast the one-day-ahead VaR values of two stocks traded at the New York Stock Exchange, two commodities traded at the Chicago Mercantile Exchange and one stock index. We will attempt to forecast the VaR values for 1% and 5% for long positions and the 95% and 99% VaR values for short positions.

In addition to the quantile regression models implemented in Haugom et al. (2014) and Steen et al.(2015) using short, medium and long-term volatility and exponentially weighted moving average volatility as regressor(s) respectively, we have implemented a quantile regression model with volatility from a GARCH(1,1) to test its accuracy against the other models. The choice of using a model with GARCH(1,1) volatility as the independent variable stems from volatility forecasts from GARCH(1,1) in several studies that have proven to model volatility pretty accurately for many assets. Hence, a priori it seems like this will be a reliable model.

VaR from RiskMetrics, GARCH(1,1) and Historical Simulation will also be estimated to serve as benchmarks for the three quantile regression models.

This paper is organised as follows. Section 2 describes some of the relevant literature and previous studies of the models. Section 3 describes the dataset and provides descriptive statistics of the data. Section 4 summarizes the methodology used in this study in greater detail, the models performance criteria and how we test them. Section 5 presents the results, while section 6 concludes our findings and suggests further research on the topic.

2 Literature review

Studies in risk management in general and Value-at-Risk in particular are numerous and the literature has become quite extensive. Below, we have highlighted and summarized some of the most relevant research used in this particular study.

With volatility in terms of the standard deviation being such a significant component of risk management the primary focus of a lot of earlier studies is how to best model stochastic volatility. In one such study Akgiray (1989) found the GARCH(1,1) model to be superior over ARCH and exponentially weighted moving average (EWMA) when forecasting monthly US stock index volatility. Bollerslev et al (1992) summarized a lot of the studies conducted in the area of modeling volatility and concluded the extensions of the ARCH model. i.e the GARCH family of models, have the ability to be effective tools for reliable volatility forecasting.

In *Does anything beat a Garch(1,1)?* Hansen and Lunde (2005) compared 330 models from the ARCH family to test who best predicts conditional variance for the DM-\$ exchange rate and the IMB stock. With GARCH(1,1) used as the benchmark, the study found no evidence the GARCH(1,1) was outperformed by more sophisticated models when applied to exchange rates. However, the GARCH(1,1) was inferior to GARCH models that took leverage effects into account when applied to the IMB stock.

RiskMetrics was developed in the mid-nineties by J.P Morgan and contributed to making VaR the industry-wide standard as a risk management measure (RiskMetrics Group, 1996). The method consists of modeling volatility dynamically by an exponentially weighted moving average process where volatility is a weighted function of previous observed volatility. McMillan and Kambouroudis (2009) conducted a study where RiskMetrics and various GARCH models were tested in forecasting VaR for 31 different stock indices. The results demonstrated RiskMetrics provide reliable VaR estimates at the 5% level. However, RiskMetrics was the worst performing model when forecasting the 1% VaR.

The Historical Simulation (HS) method for estimating VaR was suggested by Boudoukh et al.(1998) and Barone-Adesi et al.(1998, 1999). This approach fast gained popularity as a survey among financial institutions conducted by Perigon and Smith (2006) reported that nearly 3 out of 4 of the respondents prefer HS as a method for predicting VaR. Several studies involving the HS method have been conducted and Sharma (2012) summarizes the findings of 38 of these papers. The conclusion of the survey is that HS provides better unconditional

coverage than both simple and sophisticated GARCH models. With respect to conditional coverage, the HS method yields inferior performances compared to the models with dynamically modeled volatility. However, the filtered HS method demonstrates better conditional coverage results than unfiltered HS method. The Filtered HS approach involves making the VaR estimation based on volatility adjusted return rather than ordinary empirical returns (for applications see for example Alexander (2008)).

The quantile regression (QR) was first proposed by Koenker and Basset (1978). Where the basic ordinary least square method provides an estimate of the conditional mean of the endogenous variable, the quantile regression estimates the various conditional quantiles of interest directly. Hence the method seems very well suited for VaR predictions. Taylor (1999) demonstrated the ability of the QR model in VaR forecasting when he showed that the QR technique was found to perform well compared to exponential smoothing and GARCH volatility for estimating VaR.

Steen et al. (2015) compares QR to RiskMetrics and Historical Simulation in VaR estimation for 20 different commodities and one commodity index. Conditional volatility modeled as exponentially weighted moving average (EWMA) was used as the independent variable in the quantile regression, similar to the volatility of that used in RiskMetrics. The study found that over the sample the QR risk model performed better than both RiskMetrics and HS, both at the more extreme 1% and 99% levels as well as the 5% and 95% levels.

Another version of the QR model to use for VaR predictions is the HAR-QR model and was introduced in Haugom et al. (2014). This model predicts the conditional quantiles of interest directly based on measures of daily, weekly and monthly volatility estimated from observed returns. The approach is motivated by the heterogeneous market hypothesis, which claims the asymmetric behavior of the observable volatility is due to traders' different assessment of risk horizon. More specifically, the hypothesis states that short-term traders are mainly acting on short-term volatility, while long-term traders tend to disregard short-term volatility in favor of long-term volatility when deciding whether to buy or sell an asset. The HAR-QR model is a modified version of the approximating long-term memory model of Corsi (2009), which was developed to capture volatility for different horizons in accordance with the heterogeneous market hypothesis.

Compared to HS and RiskMetrics, as well as the more complex methods Symmetric Absolute Value, Assymetric Grinding, adaptive and indirect Garch(1,1), and the skewed t-student t-APARCH model, the HAR-QREG was found to perform well both regarding conditional as well as unconditional coverage.

We conclude this section with a general note that results of empirical studies often are sensitive to particular data samples. Different sample sizes may yield different results. For example, Alexander (2008) demonstrates how normal parametric VaR and HS estimation over the same data sample yields more similar results than what each one of the methods do when applied to two different data samples. This serves as an illustration of how sample size and data input may be a larger determinant of VaR estimates than the features of the particular method applied. We will attempt to address this issue by utilizing a large sample size with qualitative variation, so that the overall results reflect the models abilities rather than the data sample.

3 Data and Descriptive Statistics

The data used for this study consists of five financial assets; two stocks ExxonMobile (NYSE:XOM) and Freeport McMoRan Copper & Gold (NYSE:FCX), two commodities; WTI crude oil (CL1) and copper (HG1), and the stock index; S&P500. Logarithmic returns from daily closing prices for the period 01.12.1999 – 31.12.2015 yield five financial time series with 4045 observations each. For both commodities we use front month future prices from the CME Group. In order to avoid the jumps in returns that are typically generated when a front contract is rolled over to the next front contract, we use price adjusted contracts where the roll-over returns are smoothed out over the last four days before delivery. All of the data is downloaded from Quandl. (See www.quandl.com for more information.)

Both stocks are S&P500 components and the two companies whose stock prices are being examined trade in one each of the selected commodities (Freeport in copper, ExxonMobile in oil). The relationships will (at least perhaps for some) serve only as a quietly interesting background, as closer attempts to determine any empirical relationships constitute sufficiently work to be topics for studies on their own.

In table 1 is descriptive statistics for the five assets presented.

	S&P500	Freeport	ExxonMobile	WTI crude oil	Copper
Observations	4045	4045	4045	4045	4045
Mean	0.01%	0.00%	0.02%	0.01%	0.02%
Maximum	10.96%	25.20%	15.86%	13.34%	10.36%
Minimum	-9.47%	-21.20%	-15.03%	-16.54%	-11.71%
Standard deviation	1.27%	3.21%	1.59%	2.29%	1.75%
Excess kurtosis	8.04	5.27	10.21	3.32	3.78
Skewness	-0.19	-0.21	0.04	-0.24	-0.20
Jarque-Bera	10929	4718	17584	1894	2440
Jarque-Bera p-value	0.00	0.00	0.00	0.00	0.00
LM test for ARCH(1)	214.88	134.91	419.00	119.81	179.23
LM test p-value	0.00	0.00	0.00	0.00	0.00
Breusch-Godfrey, 2.lags	34.21	7.12	93.63	6.32	19.94
Breusch-Godfrey p-value	0.00	0.02	0.00	0.04	0.00
ADF test, 2 lags	-43.10	-41.37	-44.75	-41.18	-41.47

Table 1: Descriptive statistics for the data sample for the period 01.12.1999 – 31.12.2015

The upper panel shows the assets daily values of mean, maximum, minimum, standard deviation, skewness, and excess kurtosis. While all the assets mean returns are close to 0 as one would expect when examining daily financial data, we see the standard deviations differ substantially. In effect being a diversified portfolio (with market capitalization ratio as weights), the S&P500 has the lowest volatility in terms of standard deviation (1.27% on a daily basis) which is consistent with the teachings of Markowitz (1952). Freeport is in percentage points roughly twice as volatile as ExxonMobile over the sample period (3.21% and 1.59% respectively), while WTI is more volatile than Copper (2.29% versus 1.75% or 36.21% versus 27.67% on an annual basis when we multiply the standard deviations by the square root of 250 trading days).

Normally distributed data are assumed to have a symmetrical distribution around its mean which implies a skew of 0. Datasets with skew deviating from 0 thus deviates from the normal distribution. Skew is negative for all assets except the ExxonMobile stock which has a small positive skewness (0.04). WTI exhibits the most negative skewness (-0.21) followed by Freeport and Copper (-0.21 and -0.20, respectively). Negative skew means the distribution has a longer left tail, which means greater probability of negative returns to occur than if the distribution was symmetric (see also A.1). This means numerical parametric models such as RiskMetrics and GARCH(1,1) may be less reliable.

Excess kurtosis is the amount of kurtosis greater than 3 which is the kurtosis of the normal distribution (see also A.2). Kurtosis of a dataset provides information about how concentrated

the returns are around their mean. The values are positive for all assets, with the largest being 10.21 (ExxonMobile) and smallest 3.32 (WTI). The lowest excess kurtosis for WTI contributes to the lowest Jarque-Bera value for test of normality in return distribution for each of the assets. However, it is well above the critical level and all five null hypothesis of normality is rejected in favor of non-normality as all p-values are (approximately) zero. The kurtosis values well above 3 indicate leptokurtosis in the assets return distributions which mean the distributions have higher peak and fatter tails than the normal distribution. Further, this implies that more of the variance in the data is due to extreme deviations from the mean than would be the case if the distributions were normally distributed.

The lower panel displays some other statistics for each of the time series. Lags have been set ad hoc as we only aim to investigate for presence of different properties and not to determine the exact specification of lags for each series.

The Augmented Dickey-Fueller statistic well exceeds the critical value of -3.43 at 1% significance for all assets. Hence we conclude the five time series of returns are stationary, which gives persistence to the return series through stability in mean, volatility and autocorrelation over time and thus rules out random walk processes.

The Breusch-Godfrey test checks for the presence of autocorrelation in the return series, i.e the method tests whether the return series is independently distributed or not. The null hypothesis in our case states there is no autocorrelation for 2 autocorrelation lags, which is rejected for all five assets at 5% significance. We do however note WTI is close to not being rejected at 5%.

Presence of ARCH effects is further emphasized by the Engel's Lagrange Multiplier test for heteroscedasticity in the variance term conditional upon previous returns, as the null hypothesis of no ARCH effects from the immediate previous return on the present returns is rejected for all assets. The presence of ARCH effects suggest the time series exhibit time variant volatility which in turn imply the distributional properties are time variant. This again leads to changes in VaR quantiles over time in accordance with changing market conditions.

Figures 1-5 shows price versus returns, distributional plots and estimates of skewness and kurtosis using a rolling window of 250 trading days. We see from the development in asset prices that each of the five has experienced multiple periods of rising prices as well as periods dominated by price declines. The corresponding returns illustrate the volatility of prices over the sample. We can see that volatility increased substantially for each returns series during the financial crisis period in 2008, where all five assets experienced steep drops in prices.

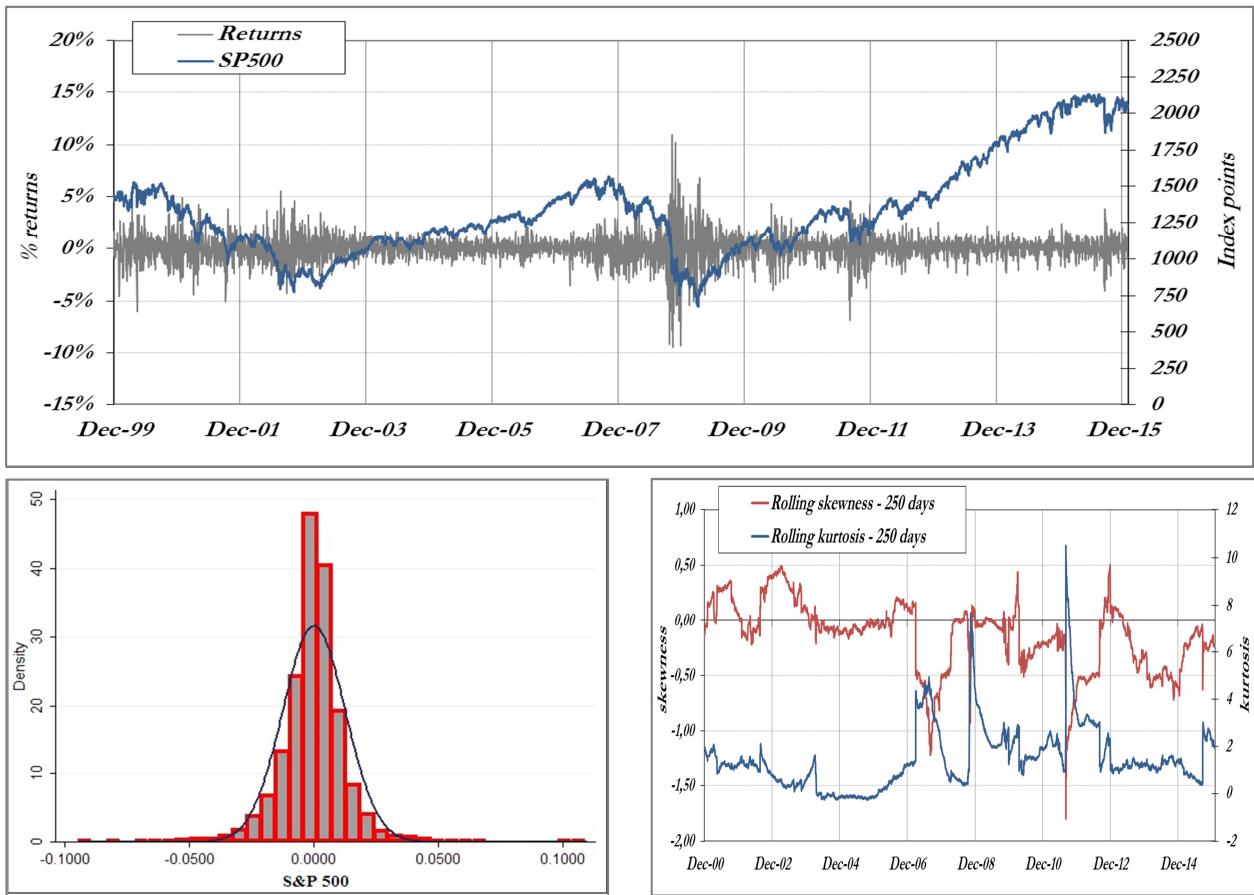


Figure 1: Upper graph: Plots the SP500 index against the right vertical axis and returns against the left vertical axis for 01.12.1999 – 31.12.2015. Lower left: Distributional plot of SP500 returns against the normal distribution. Lower right: Rolling kurtosis and skewness with a window of 250 days.




Figure 2: Upper graph: Plots the stock price of Freeport against the right vertical axis and returns against the left vertical axis for 01.12.1999 – 31.12.2015. Lower left: Distributional plot of Freeport returns against the normal distribution. Lower right: Rolling kurtosis and skewness with a window of 250 days.

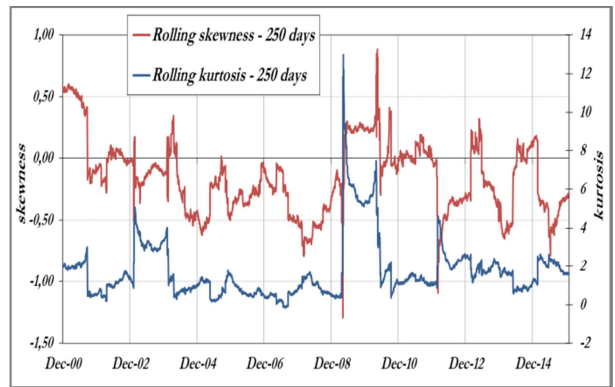
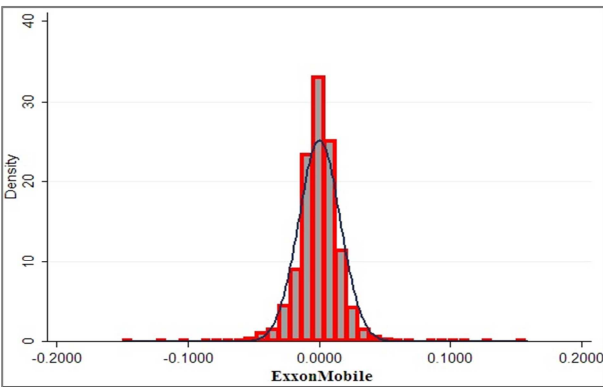



Figure 3: Upper graph: Plots the ExxonMobile stock against the right vertical axis and returns against the left vertical axis for 01.12.1999 – 31.12.2015. Lower left: Distributional plot of ExxonMobile returns against the normal distribution. Lower right: Rolling kurtosis and skewness with a window of 250 days.

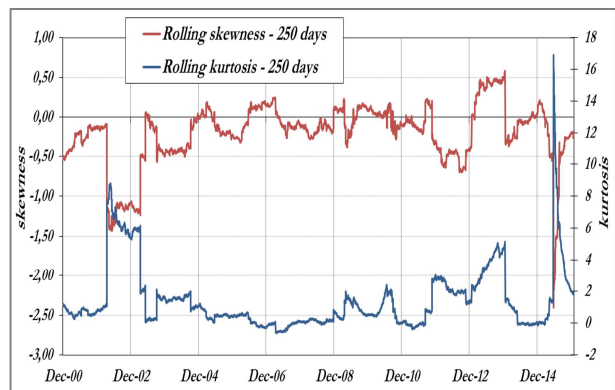




Figure 4: Upper graph: Plots the WTI contract price against the right vertical axis and returns against the left vertical axis for 01.12.1999 – 31.12.2015. Lower left: Distributional plot of WTI returns against the normal distribution. Lower right: Rolling kurtosis and skewness with a window of 250 days.

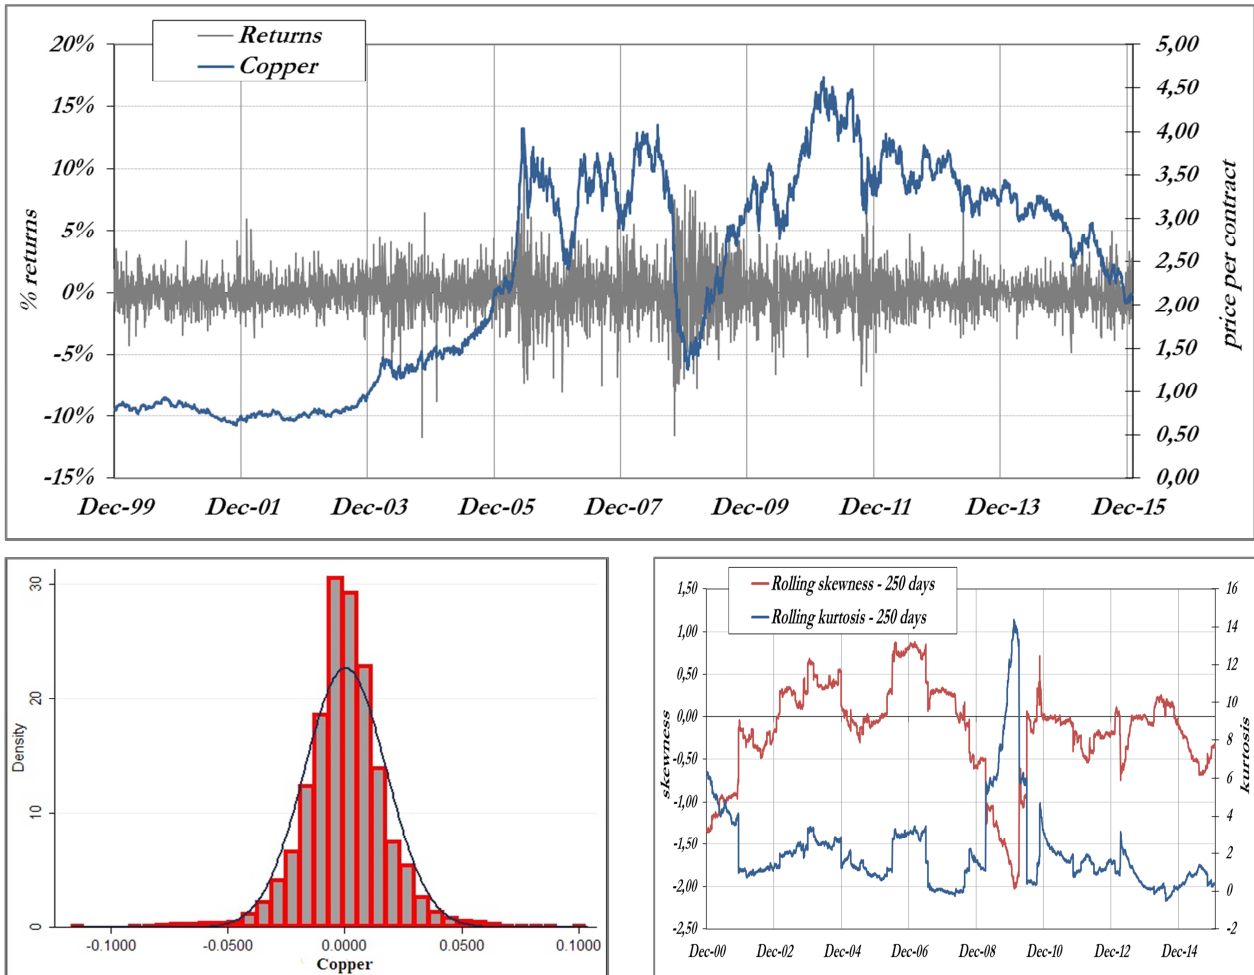


Figure 5: Upper graph: plots the Copper contract price against the right vertical axis and returns against the left vertical axis for 01.12.1999 – 31.12.2015. Lower left: Distributional plot of Copper returns against the normal distribution. Lower right: Rolling kurtosis and skewness with a window of 250 days.

A period of increased volatility can also be seen for the two stocks and the stock index around 2000-2002 when the dot com bubble burst. Even though the crisis was restricted to the technology sector initially (which also affect the SP500 index), it also evidently affected the prices for the ExxonMobile and Freeport stocks, while no particular strong effects of these market events can be seen on the two commodities.

We also notice the volatility clustering through leverage effects in some of the plots. ExxonMobile, for example, demonstrates stable volatility for the period where the stock price rose steadily from roughly December 2003 to before December 2005 and an immediate increase in volatility when the price started to meet resistance just after. That volatility is more stable in periods with rise in prices than in periods with price declines, is also easily seen throughout the SP500 series. However, the positive correlation between volatility and returns that Alexander (2008) suggests apply to commodities in general, cannot be visually detected in the WTI series. For Copper, we do see volatility increasing during the spike in the copper price from around December 2003 to the temporarily peak that was reached in 2005-2006.

Further we see that the volatility was fairly small and stable during the steady price decline from just after December 2011 onwards through the end of the sample period. However, do we also see the volatility increase at several occasions after sharp drops in the copper price, so it seems the selected commodities exhibit no (visually) clear correlation between volatility and returns throughout this data sample.

The distributional density plots of the empirical return series against the normal distribution in figures 1-5 highlight the leptokurtic properties present in the return distributions. All exhibit higher peaks and more extreme observations that constitute fatter tails than what would be the case if the returns were normally distributed. Table 2 illustrates this fact by comparing the empirical quantiles versus the normal quantiles.

	1 %		5 %		95 %		99 %	
	Empir.	Norm.	Empir.	Norm.	Empir.	Norm.	Empir.	Norm.
SP500	-3,51 %	-2,94 %	-1,97 %	-2,08 %	1,80 %	2,08 %	3,60 %	2,94 %
Freeport	-9,96 %	-7,47 %	-5,08 %	-5,28 %	4,73 %	5,28 %	8,14 %	7,47 %
ExxonM	-4,40 %	-3,70 %	-2,44 %	-2,62 %	2,29 %	2,62 %	3,94 %	3,70 %
WTI	-6,10 %	-5,34 %	-3,66 %	-3,77 %	3,53 %	3,77 %	5,80 %	5,34 %
Copper	-5,23 %	-4,08 %	-2,70 %	-2,89 %	2,75 %	2,89 %	4,84 %	4,08 %

Table 2: Comparison of empirical and normal quantiles of daily returns for the five assets. St.deviation for the entire data sample is used for calculating normal quantiles.

The comparisons of quantiles clearly demonstrates the fat tails in each of the empirical distributions compared to the normal distributions, as the absolute value of every empirical quantile is greater than the corresponding normal quantile. We see the difference in percentage points is greater in the left tail; that is, the absolute difference between empirical and normal quantile is greater than for the differences in the right tail. This is also evident in the density plots for all five assets, where we see more extreme return values located in the left tail versus the right.

Overall we can conclude that the distributional properties (volatility, skewness, kurtosis and quantiles) vary across the five selected assets. Though the skewness was not that different for four of the five assets, we see from the lower right graph in figures 1-5 that skewness and kurtosis do fluctuate over time. Together with changing and clustering volatility as elaborated above, we also can conclude from the empirical data that the distributional properties for the five assets change over time independently of each other. For a risk model to provide reliable predictions of quantiles it is essential to capture these factors.

4 Methodology

Here is a more in-depth explanation of the methods applied in this study.

4.1 Value-at-risk models

Value-at-Risk is a risk measure that is defined by Alexander (2008) as the loss that we are $100 \times \alpha\%$ confident will not be exceeded if the asset is held over a certain period of time. Alternatively VaR can be expressed as the associated loss that will occur $\alpha\%$ of the time over that particular holding period (see also A.3 for an illustration).

Assuming a mean return of zero for daily data, which the descriptive data section confirmed is suitable for our sample of daily returns, the VaR is given by

$$VaR_{\alpha} = Z_{\alpha}\sigma_t \quad , \quad (4.1)$$

where α denotes the significance level and Z_{α} denotes the quantile corresponding with the normal distribution of this particular significance level. The most basic parametric VaR models use sample standard deviation as input, in the same way the normal quantiles reported in table 2 were derived. However, did we see evidence of time variant volatility in section 3, which suggest dynamically volatility estimations as input in formula (4.1) should improve the estimation of VaR compared to a static estimate. In this study we will use RiskMetrics and GARCH(1,1) as methods for such dynamic volatility estimates.

4.1.2 RiskMetrics

RiskMetrics was first proposed by JP Morgan in 1994. The model is conditional upon normally distributed returns, which ignores the presence of fat tails which often is the case in financial data. The model does however take volatility clustering into account and is given by

$$r_t = \varepsilon_t \sigma_t \quad , \quad \text{where } \varepsilon_t \sim N(0,1), \quad (4.2)$$

$$\sigma^2 = (1 - \lambda)r_{t-1} + \lambda\sigma_{t-1}^2 \quad , \quad (4.3)$$

where λ is set equal to 0.94 for daily data, and hence is equal to an exponentially weighted moving average (EWMA). Since the persistent parameters sum up to 1 is this also classifies as a version of an Integrated GARCH model. The first term in (4.3) determines the intensity

of reaction of volatility to markets events, while the last term denotes the persistence in volatility. That is, the part of the volatility estimate that is insensitive to what happens in the market today but determined by past market events.

4.1.3 GARCH(1,1)

The Generalized AutoRegressive Conditional Heteroscedasticity (Bollerslev, 1986) is a generalization of the autoregressive conditional heteroscedasticity (ARCH) model that was developed by Engle(1982) and can capture the time varying volatility as section 3 revealed is present in our time series. By including a term of lagged conditional variance of returns, σ_{t-1}^2 the GARCH model is able to capture features such as volatility clustering and serial correlation in the returns series. As defined by Alexander (2008), a GARCH(1,1) is given by

$$\sigma_t^2 = \omega + \alpha_1 r_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \quad (4.4)$$

where r_{t-1}^2 is the lagged squared returns and ω is a constant. α_1 will denote how fast the variance will react to market shocks (with squared returns as proxy for market shock or unexpected returns (Alexander 2008)), while β_1 as mentioned above will denote the magnitude of the variance from the previous period of time.

The GARCH model's parameters are estimated by maximizing the value of the log likelihood function as defined by Alexander (2008) is given by

$$\ln L(\theta) = -\frac{1}{2} \sum_{t=1}^T \left(\ln(\sigma_t^2) + \left(\frac{\varepsilon_t}{\sigma_t} \right)^2 \right), \quad (4.5)$$

where ε_t is another notation for returns, r_t . With the restrictions that $\omega > 0$, $\alpha, \beta \geq 0$ and $\alpha + \beta < 1$, ensures the unconditional variance is finite and positive and that the conditional variance will always be positive.

The resulting volatility estimates from either the RiskMetrics or the GARCH procedure is then used as input in (4.1). VaR value for one-day-ahead is calculated by multiplying the square root of the one-day-ahead variance with Z_α .

4.2 Historical Simulation (HS)

Historical Simulation is a method that does not make use of conditional information, and thus estimates the VaR unconditional on volatility. HS is implemented by creating a database of

daily returns and then using a rolling window to find the upper and lower quantiles of interest. This value, which for the 5% VaR simply will be the 50th lowest value of the observed returns in the window of 1000 observations, is the unconditional one-day-ahead VaR:

$$VaR_{\alpha,t+1} = Q^{\alpha}(\{r_r\}_{t=1}^n), \quad (4.6)$$

where Q^{α} is the α quantile and $\{r_r\}_{t=1}^n$ is the series of returns from 1 to n , which represents the rolling window. A window of 1000 observations will be applied since this is suggested as the optimal in Alexander (2008).

4.3 Quantile regression (QR)

Quantile regression introduced by Koenker and Basset (1978) seeks to estimate the quantile of the dependent variable conditional on the independent variable. In forecasting VaR, which simply put is a specific quantile of future returns conditional on current information, QR seems as an ideal method to use. As elaborated in the literature review section studies have shown QR to deliver reliable VaR estimates. The simple linear QR model can be written as

$$r_t^q = \alpha^q + \beta^q \sigma_{t-1} + \varepsilon_t^q, \quad (4.7)$$

where the error terms distribution function is unspecified. As proposed by Koenker and Basset (1978) any solution to the following minimization problem defines the conditional q th quantile, $0 < q < 1$;

$$\min_{\alpha, \beta} \sum_{t=1}^T (q - 1_{r_t \leq \alpha + \beta \sigma_{t-1}}) (Y_t - (\alpha + \beta \sigma_{t-1})), \quad (4.8)$$

where

$$1_{r_t \leq \alpha + \beta \sigma_{t-1}} = \begin{cases} 1 & \text{if } r_t \leq \alpha + \beta \sigma_{t-1} \\ 0 & \text{otherwise} \end{cases} \quad (4.9)$$

Equivalently, but more precisely defined by (Steen et al 2015) for our particular context, the conditional quantile function can be expressed as

$$VaR_t^q | \sigma_{t-1} = \hat{\alpha}_t^q + \hat{\beta}_t^q \sigma_{t-1} + \varepsilon_t^q | \sigma_{t-1}, \quad (4.10)$$

Where sets of unique parameters for α and β can be estimated for each quantile and we can obtain the entire return distribution for a given value for the conditional volatility.

4.4 HAR-QREG (HeterogenousAutoregressive Quantile Regression Model)

The HAR-QREG model is a variation of the ordinary QR that use the realized log returns to make measures of daily, weekly and monthly historical volatility respectively and predicts the conditional quantiles directly. When

$$\sigma_{day,t} = \sqrt{r^2} \quad (4.11)$$

$$\sigma_{week,t} = \sqrt{\frac{1}{5}(r_{t-4}^2 + r_{t-3}^2 + \dots + r_t^2)} \quad (4.12)$$

$$\sigma_{month,t} = \sqrt{\frac{1}{20}(r_{t-19}^2 + r_{t-18}^2 + \dots + r_t^2)}, \quad (4.13)$$

the HAR-QREG model is defined as

$$r_t^q = \alpha^q + \beta_1^q \sigma_{day,t-1} + \beta_2^q \sigma_{week,t-1} + \beta_3^q \sigma_{month,t-1} + \varepsilon_t^q, \quad (4.14)$$

where the analogous minimization problem as formulated in (4.8) but with two extra parameters to estimate defines the q th quantile of interest.

4.5 Backtesting procedures

With backtesting we mean testing the VaR models accuracy over a historical period when the true outcome is known. For backtesting purposes we implement the Kupiec (1995) and Christoffersen (1998) tests. Kupiec is an unconditional coverage test, which statistically examines whether the frequency of exceedances (or “hits” or “violations”) over the sample is statistically close enough to the selected confidence level. For a confidence level of 99% for instance, we expect an exceedance to occur once every hundred days on average, and five exceedances to occur over that same time interval at a confidence level 95% and so on for other levels of confidence.

In order to execute the test firstly an indicator variable is defined, I_t that gets the value of 1 if the VaR estimate at time t is exceeded and the value of 0 if it is not. For quantiles, q , more than 50% at time t :

$$I_t = \begin{cases} 1 & \text{if } r_t \leq VaR_q \\ 0 & \text{if } r_t > VaR_q \end{cases} \quad (4.15)$$

In the case of short position, i.e estimation of the right tails at 95% and 99%, and generally for any quantile above 50% at time t :

$$I_t = \begin{cases} 1 & \text{if } r_t \geq VaR_q \\ 0 & \text{if } r_t < VaR_q \end{cases} \quad (4.16)$$

When n_0 is the number of violations, n_1 the number of non-violations, π_{exp} is the expected proportions of exceedances, and π_{obs} is the proportion of observed violations, the test statistic under the null hypothesis of correct unconditional coverage is:

$$\begin{aligned} -2 \ln(LR_{cc}) = 2 - [n_0 \ln(1 - \pi_{exp}) + n_1 \ln(\pi_{exp}) - n_0 \ln(1 - \pi_{obs}) \\ - n_1 \ln(1 - \pi_{obs}) - n_1 \ln(\pi_{obs})] \sim \chi_1^2 \end{aligned} \quad (4.17)$$

where n_1 is the number of violations and n_0 is the number of non-violation, π_{exp} is the expected ratio of exceedances and $\pi_{obs} = n_1 / (n_0 + n_1)$ is the observed ratio of exceedances. H_0 states we have a correctly specified model. Critical level of the test is 3.84 (5% level).

However, a reliable VaR model does also require the exceedances to be independent of each other in addition to totaling to the right amount corresponding to the confidence level. We do not want the model to produce clusters of over –or underpredictions. For example is it not desirable that all the violations are in successive order at the end of the data sample.

Christoffersen (1998) proposes a conditional coverage test which examines the models ability to produce estimates that are independent on the VaR estimate in the previous period. The test value of LR_{cc} is however only sensitive to clustering where one hit is immediately followed by another. Should there for example be a pattern where there is an exceedance every third day during a period, the test will not capture this dependency. The Christoffersen statistic is defined:

$$\begin{aligned} -2 \ln(LR_{cc}) = 2 - [n_0 \ln(1 - \pi_{exp}) + n_1 \ln(\pi_{exp}) - n_{00} \ln(1 - \pi_{01}) \\ - n_{01} \ln(\pi_{01}) - n_{10} \ln(1 - \pi_{10}) - n_{11} \ln(\pi_{11})] \sim \chi_2^2 \end{aligned} \quad (4.18)$$

n_{ij} denotes the number of observations with value i is succeeded by an observation with value j , where 1 is a hit and 0 is a no hit, $\pi_{11} = n_{11}/(n_{11} + n_{10})$ and $\pi_{01} = n_{00}/(n_{00} + n_{01})$. The critical value (5% level) from the chi distribution is 5.99. The null hypothesis, H_0 , which states the model is correctly specified statistically speaking with regards to conditional coverage, is rejected if the test statistic exceeds the critical values.

5 Results

From the sample of 4045 observation for the two stocks, two commodities and one index we lose the 20 first observations in order to get the inputs necessary for running the HAR-QREG model through the sample. With respect to the remaining 4025 observations we calculate one-day in-sample VaR estimates and compare them with the empirical returns for the various models at the 1%, 5%, 95% and 99% levels. The results are evaluated using the Kupiec (1995) and Christoffersen (1998) as described in the methodology section.

HS use the 1000 previous daily returns as the database from which the window starts rolling. This gives us the same amount of returns to backtest against as the other models.

The GARCH(1,1) and EWMA volatilities are estimated using all the 4045 observation so that the time period for which volatility input is estimated is the same and comparable for all three quantile regression models. The parameters for the quantile regression models are estimated based the remaining sample of 4025 observations.

The estimated quantile regression models and GARCH series of volatility are estimated in Stata and then implemented in Excel for running the Kupiec and Christoffersen tests (see A.4 for the estimated GARCH parameters for each assets volatility).

6 models, 5 assets and 4 VaR levels give us a total of 240 test statistics to evaluate. The results are reported in table 3 and 4.

SP500

<i>Model</i>	1 %			5 %			95 %			99 %		
	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.
EWMA QR	1,02 %	0,01	3,26	4,97 %	0,01	0,23	95,03 %	0,01	0,11	99,01 %	0,00	3,41
HAR QR	1,02 %	0,01	3,26	4,97 %	0,01	0,56	95,06 %	0,06	1,98	99,03 %	0,13	0,94
GARCH QR	1,02 %	0,01	3,26	5,04 %	0,02	0,18	95,03 %	0,01	0,21	98,98 %	0,01	3,26
HS	1,54 %	10,09	19,16	5,45 %	1,70	17,88	95,24 %	0,50	6,68	98,61 %	5,48	9,35
RiskMetrics	2,09 %	36,58	40,87	5,99 %	7,79	9,42	95,03 %	0,01	0,11	98,76 %	2,22	4,24
Garch(1,1)	1,84 %	22,91	26,49	5,37 %	1,11	1,41	95,63 %	3,48	3,58	99,16 %	1,03	0,00

Freeport

<i>Model</i>	1 %			5 %			95 %			99 %		
	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.
EWMA QR	0,99 %	0,00	3,41	4,97 %	0,01	2,61	95,03 %	0,01	2,10	99,03 %	0,04	0,78
HAR QR	0,92 %	0,27	0,00	4,94 %	0,03	0,21	94,93 %	0,02	3,72	98,93 %	0,08	0,00
GARCH QR	1,02 %	0,01	3,26	5,04 %	0,02	2,29	94,98 %	0,00	2,27	98,98 %	0,01	0,63
HS	1,39 %	5,48	12,55	5,68 %	3,73	29,26	94,99 %	0,00	12,06	98,83 %	1,05	1,39
RiskMetrics	1,99 %	32,20	35,75	5,71 %	4,14	12,58	95,06 %	0,03	1,13	98,71 %	3,17	3,34
GARCH(1,1)	1,76 %	19,33	26,13	5,04 %	0,02	2,29	95,80 %	5,74	6,02	99,13 %	0,72	1,76

ExxonMobile

<i>Model</i>	1 %			5 %			95 %			99 %		
	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.
EWMA QR	1,02 %	0,01	3,26	4,97 %	0,01	7,51	95,03 %	0,01	0,56	98,98 %	0,01	0,63
HAR QR	1,02 %	0,01	0,63	4,94 %	0,03	0,62	95,01 %	0,01	0,21	98,93 %	0,08	0,00
GARCH QR	0,99 %	0,00	3,41	5,02 %	0,00	7,10	95,03 %	0,01	0,23	99,01 %	0,00	0,68
HS	1,29 %	3,12	15,34	4,91 %	0,07	19,46	95,22 %	0,40	6,40	98,81 %	1,39	3,65
RiskMetrics	1,99 %	30,80	33,55	5,64 %	3,34	7,40	95,35 %	1,09	1,39	98,68 %	3,71	5,41
GARCH(1,1)	1,81 %	21,69	23,36	5,27 %	0,59	7,15	95,78 %	5,38	5,93	98,98 %	0,01	0,63

Table 3: Shows backtesting results for VaR models from tests applied to SP500, Freeport and ExxonMobile. The columns show number of hits, Kupiec and Christoffersen statistics for the 1%,5%, 95% and 99% quantiles. Backtesting period is 30.12.1999 – 31.12.2015. Critical values for Kupiec and Christoffersen statistics at the 5% significance level is 3.84 and 5.99 respectively.

<i>Model</i>	WTI															
	1 %				5 %				95 %				99 %			
	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	
EWMA QR	1,04 %	0,08	3,16	5,02 %	0,00	5,72	95,03 %	0,01	0,56	99,03 %	0,04	3,62	98,93 %	0,08	0,64	
HAR QR	0,94 %	0,13	0,94	4,94 %	0,03	1,87	95,03 %	0,03	1,13	98,93 %	0,08	0,64	98,93 %	0,08	0,64	
GARCH QR	0,99 %	0,00	3,41	5,02 %	0,00	10,43	94,96 %	0,02	0,89	98,98 %	0,01	3,26	98,98 %	0,01	3,26	
HS	1,46 %	7,63	17,56	5,23 %	0,44	10,23	94,94 %	0,03	6,78	98,81 %	1,39	10,63	98,81 %	1,39	10,63	
RiskMetrics	1,91 %	26,74	32,28	5,44 %	1,60	12,63	95,06 %	0,03	0,27	98,78 %	1,80	6,95	98,78 %	1,80	6,95	
GARCH(1,1)	1,52 %	9,33	12,42	4,89 %	0,10	5,36	95,45 %	1,79	2,24	99,08 %	0,27	4,21	99,08 %	0,27	4,21	

<i>Model</i>	Copper															
	1 %				5 %				95 %				99 %			
	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	Hits	Kupiec	Christoff.	
EWMA QR	1,04 %	0,08	6,84	4,99 %	0,00	3,48	95,03 %	0,01	2,10	98,98 %	0,01	0,00	98,98 %	0,01	0,00	
HAR QR	1,02 %	0,01	0,63	5,07 %	0,04	0,85	94,93 %	0,02	2,38	99,03 %	0,13	0,00	99,03 %	0,13	0,00	
GARCH QR	0,97 %	0,04	7,61	4,97 %	0,01	3,62	95,01 %	0,00	1,23	98,98 %	0,01	0,68	98,98 %	0,01	0,68	
HS	0,97 %	0,05	0,79	4,88 %	0,12	5,41	94,70 %	0,77	1,53	98,86 %	0,77	0,95	98,86 %	0,77	0,95	
RiskMetrics	1,81 %	21,69	34,62	5,42 %	1,43	5,75	94,63 %	1,11	1,48	98,34 %	14,96	15,59	98,34 %	14,96	15,59	
GARCH(1,1)	1,37 %	4,90	8,92	4,82 %	0,28	4,76	95,33 %	0,94	3,13	98,66 %	4,29	5,89	98,66 %	4,29	5,89	

Table 4: Shows backtesting results for VaR models from tests applied to WTI and Copper. The columns show number of hits, Kupiec and Christoffersen statistics for the 1%, 5%, 95% and 99% quantiles. Backtesting period is 30.12.1999 – 31.12.2015. Critical values for Kupiec and Christoffersen statistics at the 5% significance level is 3.84 and 5.99 respectively.

5.1 RiskMetrics and GARCH (1,1)

The results suggest that neither RiskMetrics nor GARCH(1,1) predicts the 1% VaR at an accurate level. Both models get rejected at all tests for the 1% VaR level for all five assets. The numbers of exceedances are systematically too high to pass the unconditional coverage tests (and the Christoffersen tests). The RiskMetrics model does especially predict too many violations for the 1% level.

Although the unconditional 5% VaR estimations are pretty accurate for both models, the overall becomes less reliable when we also consider the unconditional tests. Even though the estimates for the 5% VaR level is somewhat better for the GARCH(1,1) model than RiskMetrics, the two models fail to predict the left tails of the five assets distributions with any accuracy.

Coverage	1 %		5 %		95 %		99 %	
	Uncond.	Cond.	Uncond.	Cond.	Uncond.	Cond.	Uncond.	Cond.
RiskMetrics	0 %	0 %	60 %	20 %	100 %	100 %	80 %	60 %
GARCH(1,1)	0 %	0 %	100 %	80 %	60 %	80 %	100 %	100 %

Table 5: RiskMetrics and GARCH(1,1) aggregated success rate for unconditional and conditional coverage across the quantiles

On the contrary, the models are more reliable when predicting the right tails. We can see the GARCH(1,1) model passes both the conditional and unconditional tests at the 99% VaR level for all the five assets, while the RiskMetrics also delivers good results, the one exception being the total failure on the estimation of VaR at 99% for Copper. At the 95% VaR, both models also deliver mostly statistically accurate results, although the GARCH(1,1) model completely fails in its prediction of the Exxon stock at this level.

Considering the models are conditional normal, the results are not very surprising. Despite the models using dynamical volatility, the models do not capture the fat tails of the distributions that the descriptive statistics in table 1 showed are present in most of the distributions.

The difference of the left and right tail predictions may be due to the negative skewness of every return distribution (except for the ExxonMobile stock), which implies more extreme outcomes in the left tail. This implies the conditional normal models will underestimate the absolute value of the left tail VaR predictions. We get further evidence for this being the case from the rejected Kupiec tests for the lower quantiles due to far too many hits.

5.2 Historical Simulation (HS)

The HS methods underperforms at the 1% level. As was the case with the conditional normal models just assessed, the HS method also systematically overpredicts the number of hits for the 1% coverage. However, the overprediction is a great deal lower than that of the RiskMetrics and GARCH(1,1) models.

Coverage	1 %		5 %		95 %		99 %	
	Uncond.	Cond.	Uncond.	Cond.	Uncond.	Cond.	Uncond.	Cond.
HS	40 %	20 %	100 %	20 %	100 %	20 %	80 %	60 %

Table 6: HS aggregated success rate for unconditional and conditional coverage across the quantiles

The HS method performs much better at the 5% and 95% levels when it comes to unconditional coverage, which we see is not rejected for any Kupiec statistics across the five assets. Unconditional coverage on the other hand, is far from statistically satisfactory, as the HS method only achieve a 20% success rate at each of the levels.

This is in line with several other studies involving the HS method; it delivers good unconditional coverage, but not conditional coverage. This result can be attributed to the HS method's ability to capture the empirical return distribution, but without making it conditional on volatility. Hence time variant volatility features are not captured. Also, when the past not resembles the present in average, the unconditional predictions will be inaccurate as well.

5.3 Quantile Regression models

The three quantile regressions employed perform very well, with good results overall when evaluated at the 5% significance level for both unconditional and conditional coverage.

Coverage	1 %		5 %		95 %		99 %	
	Uncond.	Cond.	Uncond.	Cond.	Uncond.	Cond.	Uncond.	Cond.
EWMA QR	100 %	80 %	100 %	80 %	100 %	100 %	100 %	100 %
HAR QR	100 %	100 %	100 %	100 %	100 %	100 %	100 %	100 %
GARCH(1,1) QR	100 %	80 %	100 %	60 %	100 %	100 %	100 %	100 %

Table 7: QR models' aggregated success rate for unconditional and conditional coverage across the quantiles.

GARCH(1,1) QR delivers a success rate of 100% for SP500 and WTI, and 87.50% for Exxon, WTI, and Copper. The model delivers perfect unconditional coverage for all assets, but fails the conditional tests in three instances.

EWMA QR gets a passing grade of 100% for three of five assets and 87.5% for Copper and Exxon. As in the case with QR GARCH all the fails come from inability to provide correct conditional coverage.

The HAR-QR model on the other hand, delivers perfect results for unconditional as well as conditional coverage and hence gets a perfect score of 100% for all of the five assets. A closer look at the test statistics reveal that while there are no systematic differences in Kupiec statistics between the QR models, the HAR-QR tends to produce the lowest Christoffersen. This is further proof of the HAR-QR models superior ability to provide correct conditional coverage. By having three different horizons of volatility as independent variables, the model can isolate effects from each horizon to better capture the variation of conditional return distributions. This is the model's strength as outlined by Haugom et. al(2014).

The two former quantile regression models do not have that same kind of dimensionality as they are regressed on one independent variable as opposed to three. Equally for all three models though, is that by using one sample method for the quantile regression, we model the relationship between the returns and the volatility estimate(s) to be stable over time. For the QR GARCH model do we also, by using one set of estimates for the entire period, model the GARCH volatility to be developing by the same magnitude over the sample. Both these assumptions may be too restrictive or simplistic. Different sample periods will in general yield different estimates depending on changing market conditions.

Contrary to the QR GARCH (1,1) model where the input as mentioned is derived from one estimation based on the entire sample, this restriction does not apply to either HAR-QREG or the QR EWMA. This is because the inputs in the HAR-QREG model are as optimal as they can be at every instant by the model's construction. The same applies to QR EMWA, though this model also only use one set of parameters by construction.

From the estimated GARCH(1,1) parameters in A.4, we can see the estimated parameters are different from the parameters in the EWMA model, although the GARCH(1,1) output resemble the EWMA parameters, particularly in regards to Freeport and WTI. Both models put relatively much weight on past variance. Judging from the very reliable performance of both the QR EWMA and QR GARCH(1,1) does this not necessarily have any negative effect

on the reliability of the models volatility forecast. However, too much weight on previous variance can become an issue during sudden changes in market regimes, where the EWMA and GARCH(1,1) models may adapt to the new market conditions somewhat slower than the HAR-QR.

Despite the potential rigidity in constructing the models based on one set of quantile regression parameters estimated on the basis of the entire data sample, we see that all three QREG models perform very well over the sample both on their own merit and when compared to the alternative models. However, do we not see any substantial difference in performance between the QR EWMA and QR GARCH(1,1). When also taking into account that the latter in fact had a slightly worse success rate and is slightly more time consuming to implement, the results do not provide any basis for suggesting the QR GARCH(1,1) should be preferred over the QR EWMA model.

6 Conclusion

In this study different VaR models have been examined and evaluated with respect to forecasting the one-day-ahead market risk measured by Value-at-Risk. We have been particularly interested in how the quantile regression models perform both when compared to the benchmark models and when compared individually. We have tested the models abilities to predict the day-ahead VaR for long positions (1% and 5% levels) and short positions (95% and 99%) for a small selection of equities and commodities. The models have been implemented over a sample of 4025 trading days in an in-of-sample study for which the models have been evaluated. Our main finding is that the quantile regression models perform better than the benchmark models RiskMetrics, GARCH(1,1) and Historical Simulation.

The benchmarks models deliver substantially inferior results compared with the quantile regressions over this data sample of financial assets. All of the models examined do however perform better when predicting the upper tail of the return distributions versus the lower tail. This could be due to the fact most of the series exhibit negative skewness and leptokurtosis. This suggests there are more observations between the mean of the distribution and the maximum value compared to the number of returns in the left end of the distribution. Also the conditional normal models do slightly better at predicting the right tail of the equities compared to the commodities. This could be due to the assets difference with respect to leverage effects; the stocks volatility is in general more stable in periods with positive returns than what is the case for commodities.

Among the quantile regression models is HAR-QR the model that performs best of all with a perfect score for all unconditional and conditional coverage tests. EWMA QR and GARCH(1,1) QR have a success rate of 95% and 92.5% respectively. HAR-QR only requires observed returns to model day-ahead VaR and is therefore perhaps the easiest method to implement. It is an interesting feature of the study that the most easily implemented model performs the best.

However, as emphasized by Alexander (2008) is VaR forecasting sensitive to sample period and choice of financial assets. This means we should be reluctant to make any claims of one model's universal superiority over another based solely on the basis of this study. However, we will conclude that the findings in this study contribute to further proof of the ability of quantile regression models to deliver robust VaR forecasts in general. We also draw the conclusion that based on this data sample, the proposed GARCH(1,1) QR doesn't seem to offer any significant improvement over the EWMA QR model.

While the data samples in this study contain five assets which are qualitatively different in terms of volatility ranging from the low volatility S&P500 index to the highly volatile Freeport stock, the assets distributions in terms of skewness and kurtosis are not that different. Further studies where the performance of linear quantile regression models are tested should possibly consider applying assets with more variations in distributional properties as well as volatility in order to further stress the models abilities to produce reliable VaR estimates.

Another natural extension of this study could be to implement quantile regression with volatility estimated from models that are more specifically optimized for each asset. Even if the GARCH(1,1) delivers reasonable volatility forecasts, input from more other GARCH models could further enhance the performance of the model. For example, asymmetric GARCH models have in general proven to reliably forecast volatility at longer horizons, as demonstrated in Marcucci (2005). Or implementation of two-regime models such as Markov regime-switching GARCH model (see for example Hull (2012)) could be considered. It is safe to say that there are several possibilities for model selection in the pursuit of finding a model that may match the performance of the HAR-QR model.

References

- Alexander, C. (2009). *Market Risk Analysis: Volumes I, II, III and IV*, Vol. 4. Wiley.
- Akgiray, V. (1989). Conditional heteroskedasticity in time series of stock returns: Evidence and forecasts. *The journal of business* **62**(1), 55–80.
- Barone-Adesi, G. Bourgoin, F and Giannopoulos, K. (1998). Don't look back. *Risk* **11**(8), 100-104.
- Barone-Adesi, G. Giannopoulos, K and Vosper, L. (1999). VaR without correlations for non-linear portfolios. *Journal of Futures Markets* **19**(5), 583-602.
- Bollerslev, T., Chou, R., Kroner, K., and Jayaraman, N. (1992). ARCH modeling in finance: a review of the theory and empirical evidence. *Journal of Econometrics* **52**(1-2), 5–59.
- Boudouk, J., Richardson, M., and Whitelaw, R. (1998). The best of both worlds. *Risk* **11**(5), 64-67.
- Christoffersen, P. (1998). Evaluating interval forecasts. *International Economic Review* **39**(4), 841-862.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* **7**(2), 174-196.
- Dickey, D., and Fuller, W. (1979). Distribution of the estimation for autoregressive time series with a unit root. *Journal of the American Statistical Association* **74**(366), 427-431.
- Engle, R. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of UK inflation. *Econometrica* **50**(4), 987-1008.
- Hansen, P., and Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1)?. *Journal of applied econometrics* **20**(7), 873-889.
- Haugom, E., Ray, R., Ullrich, J. Veka, S., and Westgaard, S.(2014). A simple quantile regression model to forecast day-ahead Value-at-Risk.
- Hull, J. (2012). *Options, futures and other derivatives*. 8th edition. Pearson.
- Jorion, P. (2001) *Value at Risk: The new benchmark for managing financial risk*. New York: McGraw-Hill.
- Koenker, R., and Bassett, G.(1978). Regression quantiles. *Econometrica* **46**(1), 33-50.
- Kupiec, P. (1995). Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives* **3**(2), 73-84.

- Markowitz, H. (1952). Portfolio selection. *Journal of Finance* **7**(1), 77-91.
- Marcucci, Juri. (2005). Forecasting stock market volatility with regime-switching GARCH models. *Studies in Nonlinear Dynamic & Econometrics* **9**(4), 1-55.
- McMillan, D. G., and Kambouroudis, D. (2009). Are RiskMetrics forecasts good enough? Evidence from 31 stock markets. *International Review of Financial Analysis* **18**(3), 117-124.
- RiskMetrics Group (1996) RiskMetrics technical document. J. P. Morgan/ Reuters, fourth edition. <https://www.msci.com/documents> (accessed January 25, 2016)
- Sharma M. (2012). The historical simulation method for Value-at-Risk: A research based evaluation of the industry favorite. *Indian Institute of Management*. PGDM.
- Steen, M., Gjølberg, Ole., and Westgaard, S. (2015). Commodity value-at-risk modeling: comparing RiskMetrics, historic simulation and quantile regression. *Journal of Risk Model Validation* **9**(2), 49-79.
- Taylor, J. (1999). A quantile regression approach to estimating the distribution of multiperiod returns. *Journal of Derivatives* **7**(1), 64-78.

Appendix

A.1




Figure 1: Distributions with negative and positive skew

A.2




Figure 2: Illustrations of meokurtic (A), leptokurtic (B) and Platykurtic (C) distributions

A.3




Figure 3: Illustration of VaR with critical value of 5% , $\alpha = 5\%$, so that the confidence level is 95% (a) is the significance interval and (b) is the confidence interval

A.4

	ω	α_1	β_1
SP500	1.74e-06	.0917648	.895721
Freeport	7.35e-06	.0566837	.9372174
ExxonMobile	3.37e-06	.076373	.9087458
WTI	3.15e-06	.0574707	.9381223
Copper	1.66e-06	.0425723	.9520863

Tabell 1: The resulting coefficient output from GARCH(1,1) estimation for the five assets in the data sample



Norges miljø- og biovitenskapelig universitet
Noregs miljø- og biovitenskapelige universitet
Norwegian University of Life Sciences

Postboks 5003
NO-1432 Ås
Norway