LUND UNIVERSITY

MASTER'S THESIS

Forecasting the Volatility in Financial Assets using Conditional Variance Models

---

Authors:
Hugo Hultman
Jesper Swanson

Supervisor:
Dag Rydorff

LUND UNIVERSITY

DEPARTMENT OF ECONOMICS
SEMINAR DATE: MAY 31, 2017

Abstract

This thesis examines multiple ARCH-family models' volatility forecasting performance on the London Bullion Market Gold price, the OMXS30, and the USD/EUR exchange rate. Further, this thesis uses two different time periods to exploit differences and similarities in the forecast accuracy among the conditional variance models. The models we examine are the ARCH, the GARCH, the IGARCH, the EGARCH, and the GJR-GARCH model. Furthermore, we divide each period into an in-sample and an out-of-sample period. The models are estimated in the in-sample period and then used to forecast the volatility in the out-of-sample period. This thesis uses the squared returns as an unbiased approximation of the latent volatility. The forecasts are evaluated using two loss functions, the mean absolute error and the mean squared error. The results indicate a very inconsistent ranking among the models. None of the models seem superior to the other models, based on both loss functions, when forecasting the conditional volatility in the different assets and periods. However, the ARCH model seems to perform well relative the other models, when forecasting the volatility of the gold using only the mean absolute error as a tool to evaluate the forecasts.

Keywords: *Volatility Forecasting, Conditional Variance, ARCH, GARCH, IGARCH, EGARCH, GJR-GARCH*

Acknowledgement

We would like to express our gratitude to our supervisor, Dag Rydorff. Thank you for your guidance and help throughout this thesis.

Lund, May 24, 2017
Hugo Hultman & Jesper Swanson

Contents

1 Introduction

Modelling and forecasting the volatility is fundamental when pricing financial instruments, calculating measures of risk and hedging against portfolio risk. Many financial instruments require the future volatility. One example is when pricing an option where the only unknown variable is the underlying assets' volatility from now until the exercise date (Hull, 2011). The importance of the volatility leads to an extensive amount of financial and economic literature within the subject of modelling and forecasting the volatility in the financial market. Numerous models are developed for the purpose of capturing the behavior of the volatility. Financial data's volatility often exhibits periods of turmoil followed by relative tranquil periods (Enders, 2014). To capture these patterns, Engle (1982) proposes the autoregressive conditional heteroscedasticity (ARCH) model which has become a standard tool when modelling the conditional volatility. Through the years, extended versions of the ARCH model are introduced to capture different features of the volatility. However, the true generating process is not observable and may differ between different assets (Nelson and Foster, 1995). Consequently, there may not be a single best model for modelling the volatility.

There are several articles published with the purpose of finding suitable models that can describe the behavior of the volatility. The majority of these articles focus on stock indexes and exchange rates. Furthermore, some authors compare these assets and examine the similarities and the differences of the best fitting volatility models (Poon and Granger, 2003). To our knowledge, no previous research is comparing the models' ability to explain the volatility during different periods of time. Hence, this thesis aims to contribute to the literature by examining this area. We examine three different assets during two periods. Furthermore, we forecast the daily volatility in these assets and periods using five different volatility models. The forecasts are evaluated and hence we exploit the similarities and differences between these models for the different assets and periods. In this thesis, two main questions are considered. Firstly, if there is any model that has superior forecast accuracy among the assets. Secondly, if the models with the highest forecast accuracy differ between the two periods.

In this thesis, we examine the following assets: the London Bullion Market (LBMA) Gold price, the OMXS30 and the USD/EUR exchange rate. The first period we investigate covers the timespan January 1995 through December 1999, totaling 1305 observations. The Second period starts in January 2012 and ends in December 2016, also totaling 1305 observations. Furthermore we forecast the volatility using the ARCH, the GARCH, the IGARCH, the EGARCH, and the GJR-GARCH model. When estimating these models we consider two different distributions, the normal distribution and the Student's t-distribution. To be able to evaluate our forecasts we divide the two periods into an in-sample period and an out-of-sample period. The in-sample period is used to examine the data and estimate the volatility models. Furthermore we use the estimates from the in-sample periods to perform forecasts on the out-of-sample periods. We then compare the different models' forecast accuracy by using two loss functions, the mean absolute error (MAE) and the mean squared error (MSE).

Multiple major events affecting the economy such as the terror attacks on September 11 and the financial crisis 2008 have occurred in between our test periods. We believe that these shocks could affect people's way to react to future shocks. Hence, this may affect the behavior of the volatility for the examined assets. Another factor that may have affected the behavior of the volatility is the availability of information. Information has become more accessible which may have caused people to react to shocks more rapidly. Our hypothesis is therefore that the most accurate forecast model may differ between the two test periods.

The results indicate a very inconsistent ranking among the models. There is no single best volatility forecasting model in the different assets and periods based on the two loss functions. However, the results indicate that the GJR-GARCH is always outperformed by another model for all assets and periods. When only evaluating the forecasts using the mean absolute error, the ARCH model seems to be one of the better models. The ARCH model has the highest forecast accuracy for both periods when forecasting the conditional volatility for the gold. Hence, the choice of a loss function becomes vital and should depend upon the purpose of the forecasts.

The remainder of this thesis is organized as follows. Section 2 presents relevant previous research. Section 3 describes the empirical methodology and the theory used throughout this thesis. Section 4 presents the data and the tests we use to uncover patterns. This section also includes our approach when transforming the data into a desirable structure which can be used for estimating the volatility models. Section 5 presents our empirical results together with a coherent analysis. Finally, in Section 6 we summarize and conclude the results.

2 Previous Research

The importance of modelling the conditional variance in the financial sector has contributed to an extensive empirical research within the topic. Many articles focus on finding suitable models with the purpose of forecasting the conditional variance. However, different models may be superior to other models for different forecast purposes. Consequently, the evaluation process of the forecasted conditional volatility becomes crucial.

Numerous types of models, more or less sophisticated have been developed in the past decades for the purpose of modelling the conditional variance. Among these models, one can find Engle's (1982) ARCH model and Bollerslev's (1986) GARCH model. These models are frequently used, mostly due to their simplicity and capacity to capture the volatility clustering which often occurs in financial time series data (Enders, 2014). Hansen and Lunde (2001) compare 330 different ARCH-family models and their ability to capture the conditional variance in the daily DM/USD exchange rate and the daily IBM stock returns. The authors' purpose is to evaluate the GARCH(1,1) model's superior ability to predict the future volatility of the exchange rate and the stock return. The DM/USD data set consists of the daily spot exchange rate from October 1987 through September 1993, where the data is divided into an in-sample period and an out-of-sample period. The out-of-sample period is one year and it is used to compare the different models' forecast ability. Furthermore, the in-sample period for the IBM returns spans the period January 1990 through May 1999. The out-of-sample period for the IBM returns starts in June 1999 and ends in May 2000. Hansen and Lunde (2001) use intraday data to calculate their approximation for the latent volatility. Further, they use six different loss functions and find no evidence for the GARCH(1,1) model being inferior in terms of capturing the volatility of the exchange rate. However, in terms of capturing the volatility for the IBM stock returns, they find the GARCH(1,1) model to be inferior to the other models. Hansen and Lunde (2001) also conclude that the IGARCH model performs poorly, using the MAE and MSE as evaluation criterion, when forecasting the volatility of the IBM stock return. Apart from this, the authors' findings also indicate a more accurate forecast performance for the EGARCH model, in the IBM stock, when the error terms are assumed to follow the Student's t-distribution rather than the normal distribution.

Mckenzie and Mitchell (2002) investigate 17 currency pairs using 14 different ARCH-family models. They also find the GARCH(1,1) model to be suitable when modelling the volatility in exchange rates. However, they further argue that models accounting for leverage effects may perform better if there exist any asymmetry in the market's response to shocks. Other researchers restrict themselves to fewer models. Among them we find Franses and Van Dijk (1996). They examine the GARCH model together with two more exotic models, the Quadratic-GARCH and GJR-GARCH model. The authors' explore the three models' ability to forecast the weekly stock market volatility between 1990 and 1994. Their findings suggest that the GJR-GARCH model may not be an appropriate model to use when forecasting the volatility in exchange rates.

There exist a larger amount of published articles focusing on modelling the volatility in exchange rates and equity indexes compared to commodities. Trück and Lian (2012) investigate the volatility dynamic of the gold market. Their article explores different models that can be used to predict the future volatility of the gold return. To find an appropriate model for this, they implement three models from the ARCH-family. The forecasts are carried out over an out-of-sample period. Furthermore to evaluate the forecasts they consider three different loss functions, the mean squared error, the mean absolute error and the root mean squared error. The error statistics are calculated using the squared returns as a proxy for the latent volatility. However, their out-of-sample results are inconsistent. Hence they conclude that the model specification such as the lag order is crucial when specifying the models.

3 Method and Theory

3.1 In-sample and Out-of-sample

There is no general guideline on how a given data set should be divided into in-sample and out-of-sample subsets. Aronson (2011) argues that a common technique is to split your sample into two subsets. In this thesis we use this approach and divide our sample into one training set and one validation set, also called an in-sample and an out-of-sample period. The parameters for the conditional variance models are estimated in our in-sample period. Further, the estimated parameters are used to forecast the volatility in the out-of-sample period. The two periods we examine in this thesis consist of data over five years each. The length of Hansen and Lunde's (2001), and Andersen and Bollerslev's (1998) data is similar to our data length. They use a one-year forecast horizon and hence this is used in our thesis too. This leaves us a four-year in-sample period. Consequently, the in-sample period for our first period covers the timespan, January 1995 to December 1998 and the out-of-sample period starts in January 1999 and ends in December 1999. Furthermore, the second period's in-sample consists of data from January 2012 to December 2015 and the out-of-sample consists of data from January 2016 to December 2016.

3.2 Stationarity

A common assumption for a sequence of variables is that the sequence is independent. This might however not always be the case, especially when working with time series data. A stochastic process is said to be strictly stationary if its joint probability function is invariant over time. However, this condition is difficult to prove empirically. Hence, a weaker condition is often assumed. If a stochastic process's autocovariance, mean and variance do not depend on time it is said to be covariance stationary. Without stationarity, it would be hard to make inference about sample statistics, and the accuracy will vary at different time points. In this thesis, we use the Augmented Dickey-fuller test (ADF) to test for stationarity. The ADF tests the null hypothesis of a unit root in a time series sample. If the time series exhibits a unit root process it is said to be non-stationary (Tsay, 2001).

3.3 Log returns

Prices such as exchange rates, stock indexes, and commodities can typically be described by a unit root process, implying non-stationarity. This is also true for the prices examined in this thesis, see Appendix 8.1. A common approach is therefore to model the returns rather than the actual price itself. Meucci (2005) describes that returns are more invariant compared to prices. Additionally, Tsay (2001) argues that the statistical properties of the continuously compounded returns, also called log returns, are more tractable compared to the simple net returns. The log return is defined as

$$r_t = \ln p_t - \ln p_{t-1} \tag{3.1}$$

where p_t denotes the price of the asset at time t. Brownlees, Engle, and Kelly (2011) and Tse (1998) among others use the log returns when estimating the volatility for different financial assets. Hence, the log return is also used in this thesis.

3.4 Autocorrelation function

Before estimating a conditional variance model to our data, we need to examine the properties of the data. In contrast to conventional sampling data, time series data are ordered. Hence, there might exist temporal patterns. Further, the residuals need to be described by a white noise process when estimating the conditional variance models. Any structure within the residuals therefore needs to be modelled separately. One way to do this is to include a mean equation in the conditional variance models which removes the structure in the residual, making it a white noise process (Tsay, 2001). The autocorrelation function (ACF) is commonly used as a tool to investigate these patterns. The ACF illustrates the correlation between data points by different time lags. The autocorrelation for the returns' residuals is calculated by Equation 3.2, where k denotes the number of lags (Enders, 2014).

$$\rho_k = \frac{cov(\varepsilon_t, \varepsilon_{t-k})}{\sqrt{var(\varepsilon_t) var(\varepsilon_{t-k})}} \tag{3.2}$$

3.5 Ljung-Box test

The autocorrelation function is used to get a visual overview of the structure in the returns' residuals. To further assure us if there exist any autocorrelation in the residuals we perform a Ljung-Box test. The Ljung-Box test tests for autocorrelation for multiple lags jointly. The null hypothesis is that autocorrelation up to lag k equal zero and the Ljung-Box Q-statistics is calculated by Equation 3.3, where n denotes the sample size, ρ denotes the autocorrelation at lag k, and h is the number of lags being tested Enders(2014).

$$Q_{LB} = n(n+2) \sum_{k}^{h} \frac{\rho_k^2}{n-k} \tag{3.3}$$

3.6 Approximation for the volatility

As mentioned in the introduction, the volatility is central in many asset pricing models. However, the volatility is not observed and hence it needs to be approximated. Consequently, one has to rely on proxies when forecasting the volatility. This makes the choice of a good proxy for the volatility crucial. In finance, the volatility is often approximated using the sample standard deviation or the sample variance. The sample variance is calculated by Equation 3.4 and the sample standard deviation is obtained by taking the square root of Equation 3.4.

$$\hat{\sigma}_t^2 = \frac{1}{n-1} \sum_{t=1}^{n} (r_t - \bar{r}_t)^2 \tag{3.4}$$

Using these two approximations is not entirely satisfactory when the sample is small. Instead, an alternative approach is to use the daily squared returns, calculated from the assets closing prices. The daily squared returns can be justified as a proxy since they are an unbiased estimator of the volatility. However, despite this fact, daily squared returns are an extremely noisy proxy. A noisy proxy for the latent volatility may lead to a small difference between the estimated forecasts (Poon and Granger, 2003).

The increased availability of high-frequency data has made it possible for new measures to evolve. Andersen and Bollerslev (1998) suggest an alternative proxy for the volatility based on intra-day returns. This measure has become frequently used within finance and has been shown to perform well on out–of-sample data (Poon and Granger, 2003). However such data is very time consuming to process. Hence, the squared returns are used as a proxy for volatility in this thesis.

3.7 Engle's ARCH test

The residuals for the returns can be uncorrelated and still exhibit conditional heteroscedasticity, meaning that the squared residuals are autocorrelated. If the squared residuals are autocorrelated, they are said to have an ARCH effect. Autocorrelated squared residuals implies that the ARCH-family models may be appropriate to use when modelling the conditional variance of the residuals. Furthermore, an Engle's ARCH test is used to examine if the residuals exhibit any ARCH effect with the null hypothesis of no ARCH effects in the residuals (Engle, 1982).

3.8 Conditional volatility models

In conventional econometric models, the variance of the error term is often assumed to be constant. However, this is not always an appropriate assumption. Financial time series, including the ones we examine, often exhibit periods of unusually high volatility followed by periods of relative tranquil volatility, also known as volatility clustering. Suppose that the returns can be described by the following process

$$r_t = \mu + \varepsilon_t. \tag{3.5}$$

Further, the returns' innovation process for the conditional variance models is shown in Equation 3.6, where v_t is a white-noise process such that $v_t \sim iid\,(0,1)$ and h_t is the conditional variance at time t (Enders, 2014).

$$\varepsilon_t = v_t \sqrt{h_t} \tag{3.6}$$

In this thesis, five different models from the ARCH-family are used to estimate and forecast the conditional variance, h_t.

3.8.1 ARCH Model

When modelling the volatility of a time series containing heteroscedasticity one may use the autoregressive conditional heteroscedasticity (ARCH) –family models as a tool to capture these effects. Robert F. Engle (1982) develops the ARCH model, which aims to model the conditional variance, see Equation 3.6, and the ARCH(1) model is defined in Equation 3.7.

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 \tag{3.7}$$

To ensure a nonnegative conditional variance, restrictions such as $\alpha_0 > 0$ and $\alpha_1 \geq 0$ are used. The conditional variance of an ARCH(1) model depends on its previous squared error. A problem with the ARCH(1) model is that to capture the effect of volatility clustering, a large number of lags are often required. Furthermore, a high number of lags may lead to the model not being parsimonious (Enders, 2014).

3.8.2 GARCH Model

Bollerslev (1986) extends the ARCH model by introducing a technique that allows the conditional variance to depend upon its own lags, known as the Generalized ARCH (GARCH) model. The number of ARCH lags is often reduced by enabling the conditional variance to depend on its previous values. The fewer lags make the model easier to identify as well as easier to estimate. The GARCH(1,1) model is defined in Equation 3.8.

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 h_{t-1} \tag{3.8}$$

One can easily observe that if β_1 is zero, then Equation 3.8 is equal to the ARCH(1) model in Equation 3.7. To guarantee a nonnegative conditional variance, Bollerslev (1986) imposes restrictions on the coefficients such as, $\alpha_0 > 0$, $\alpha_1 \geq 0$, $\beta_1 \geq 0$, and $\alpha_1 + \beta_1 < 1$ (Bollerslev, 1986).

3.8.3 IGARCH Model

It is often the the case that the sum of α_1 and β_1 is close to one when estimating a GARCH(1,1) model on financial time series data. A GARCH(1,1) model with the restrictions such that $\alpha_1 + \beta_1 = 1$ and $0 < \beta_1 < 1$ is called an Integrated GARCH model (IGARCH). These restrictions allow the volatility shocks to be permanent. In an IGARCH model, the process is forced to act as a unit root process (Enders, 2014). Nelson (1990) argues that the IGARCH model can yield a very parsimonious representation of the returns distribution. The IGARCH(1,1) model is defined in Equation 3.9 (Enders, 2014).

$$h_t = \alpha_0 + (1 - \beta_1)\, \varepsilon_{t-1}^2 + \beta_1 h_{t-1} \tag{3.9}$$

3.8.4 EGARCH model

A substantial disadvantage of the ARCH, the GARCH, and the IGARCH models is their lack of capturing asymmetries in the volatility with respect to the sign of past shocks. An interesting feature in financial volatility is that it usually behaves differently depending on whether the shock is positive or negative. Negative shocks tend to have a greater impact on the volatility compared to positive shocks. This is often called leverage effect (Enders, 2014). Nelson (1991) develops the Exponential GARCH (EGARCH) model, which allows for asymmetric effects and therefore solves one of the main drawbacks of the ARCH, the GARCH, and the IGARCH models. The EGARCH(1,1) model is defined as

$$\ln(h_t) = \alpha_0 + \alpha_1 \frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + \gamma_1 \left| \frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} \right| + \beta_1 \ln(h_{t-1}) \tag{3.10}$$

where the conditional variance is measured in a log-linear form. Furthermore, there is no need for positive restrictions on the coefficients since the implied value of h_t is always positive. As mentioned above, the main advantage of the EGARCH model compared to the ARCH, the IGARCH, and the GARCH model is the allowance for asymmetric effects. If $\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}$ is positive, then the effect of the shock onto the log conditional variance is $\alpha_1+\gamma_1$. In contrast, if $\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}$ is negative, the effect of the shock onto the log conditional variance is $\gamma_1 - \alpha_1$ (Nelson 1991).

3.8.5 GJR-GARCH Model

Glosten, Jagannathan, and Runkle (1993) discover an alternative way to model asymmetric effects in asset returns, the GJR-GARCH model. The model allows positive and negative returns to have different impact on the conditional variance, which makes it similar to the EGARCH model. The GJR-GARCH is defined in Equation 3.11.

$$h_t = \alpha_0 + (\alpha_1 + \gamma_1 I_{t-1})\,\varepsilon_{t-1}^2 + \beta_1 h_{t-1} \tag{3.11}$$

The parameters α_0, α_1, γ_1, and, β_1 need to be equal or greater than zero to guarantee a non-negative variance. Further, I_{t-1} is an indicator function and it is defined in Equation 3.12 (Glosten, Jagannathan and Runkle, 1993).

$$I_{t-1} := \begin{cases} 0\; if\; \varepsilon_{t-1} \geq 0 \\ 1\; if\; \varepsilon_{t-1} < 0 \end{cases} \tag{3.12}$$

3.9 The distribution of v_t

When estimating the conditional variance models using the maximum likelihood one must also make an assumption regarding the distribution of v_t, see Equation 3.6. The most common approach is to assume that v_t follows a normal distribution. However, when using financial data, this assumption may be violated. It is often the case that the distribution of financial data exhibits fatter tails than the normal distribution (Bradley and Taqqu, 2003). In such case, a distribution featuring higher kurtosis might be more appropriate. Student's t-distribution is symmetric and has a similar shape as the normal distribution. However, the Student's t-distribution is characterized by having heavier tails compared to the normal distribution (Zivot, 2009). Hence, in this thesis, we estimate the parameters in the conditional variance models using both distributions. The probability density function for the normal distribution is defined as

$$f(x) = \frac{1}{\sqrt{2\sigma^2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \tag{3.13}$$

whereas the probability density function for the Student's t-distribution is defined as

$$f(x) = \frac{\Gamma\left(\frac{df+1}{2}\right)}{\sqrt{df\pi}\,\Gamma\left(\frac{df}{2}\right)} \left(1+\frac{x^2}{df}\right)^{-\frac{df+1}{2}} \tag{3.14}$$

where df denotes the number of degrees of freedom and Γ denotes the gamma function: $\Gamma(x) = \int\limits_0^\infty y^{x-1}e^{-y}dy$ (Miller and Miller, 2004).

3.10 Q-Q plot and test for normality

A common approach to compare different distributions in statistics is to plot the distributions' quantiles in a quantile-quantile (q-q) plot. As mentioned in Section 3.9 we estimate the conditional variance models using two different assumptions on v_t, the normal distribution and the Student's t-distribution. Consequently, we use the q-q plot to examine if the returns' residuals for the different assets in the in-sample periods fit these distributions. The quantiles of the residuals are plotted on the y-axis, and the quantiles of the compared distribution are plotted on the x-axis. If the two distributions being compared are similar, the points in the q-q plot will be formed as a linear line. We also perform a Jarque-Bera test which is a goodness of fit test. The Jarque-Bera test tests if the residuals exhibit skewness and kurtosis matching a normal distribution. If the null hypothesis is rejected then one can conclude that the residuals do not follow the normal distribution (Keya and Rahmatullah, 2016).

3.11 Maximum Likelihood Estimation

A common method to estimate the parameters in the conditional variance models is the maximum likelihood method. The first step in the maximum likelihood estimation is to form the likelihood function. Let θ be a vector of k unknown parameters to be estimated and $f(x_t|I_{t-1})$ a conditional probability density function of x_t where I_{t-1} denotes the information available at time $t-1$. If $f(x_t|I_{t-1}) = f(x_t)$, then the joint density function, $f(x_1, x_2 \dots x_T|\theta)$, can be defined as $f(x_1, x_2 \dots x_T|\theta) = f(x_t|I_{t-1})\, f(x_{t-1}|I_{t-2}) \dots f(x_1)$. The likelihood function, $L(\theta|x_1, x_2, \dots,\ x_T)$, is the same function as the joint probability density but instead of seeing it as a function of the data given a set of parameters the likelihood function is viewed as a function of the parameters given a set of data. Hence, the likelihood function is defined in Equation 3.15.

$$L\left(\theta|x_1, x_2, \ldots,\ x_T\right) = \prod_{t=1}^{T} f(x_t|I_{t-1}) \tag{3.15}$$

$\hat{\theta}_{ML}$ is the argument that maximizes the likelihood function. For simplicity, one often uses the loglikelihood function rather than the likelihood function itself. Due to the fact that the loglikelihood function just being a monotonic transformation of the likelihood function, the argument that maximizes the likelihood function is also the one maximizing the loglikelihood function. The exact form of the loglikelihood function depends on the distribution of v_t, assuming v_t to follow the normal distribution generates the following loglikelihood function when estimating the conditional variance models

$$\ln L = -\frac{T}{2}\ln\left(2\pi\right) - \frac{1}{2}\sum_{t=1}^{T}\ln h_t - \frac{1}{2}\sum_{t=1}^{T}\frac{\varepsilon_t^2}{h_t} \tag{3.16}$$

where h_t is being replaced depending on which conditional variance model being estimated. When assuming v_t to follow the Student's t-distribution rather than the normal distribution the loglikelihood function is defined in Equation 3.17, where $df > 2$.

$$\ln L = T\ln\left[\frac{\Gamma\left(\frac{df+1}{2}\right)}{\sqrt{\pi\left(df-2\right)}\Gamma\left(\frac{df}{2}\right)}\right] - \frac{1}{2}\sum_{t=1}^{T}\ln h_t - \frac{df+1}{2}\sum_{t=1}^{T}\ln\left[1+\frac{\varepsilon_t^2}{h_t\left(df-2\right)}\right] \tag{3.17}$$

If the maximum likelihood function is misspecified in terms of distribution but can be argued to be asymptotically consistent it is sometimes called the Quasi-maximum likelihood estimator. It can then be shown that the estimates are asymptotically normal distributed (Verbeek, 2004).

3.12 Forecast Procedure

As described in Section 3.1, the parameters for the conditional variance models are estimated in the in-sample periods. Further, the estimated models are used to forecast the daily conditional volatility in the out-of-sample period. Hence, the parameters are only estimated once for each model and period. These estimates are then used for all forecasts in the out-of-sample period. However, the GARCH, the IGARCH, the EGARCH, and the GJR-GARCH model require an initial value for h_t. Zivot (2009) argues that a common approach is to use the unconditional variance for the in-sample period. Hence, this method is used in this thesis as well.

3.13 Forecast Evaluation

A common approach when evaluating forecasts is to use a loss function. The loss function compares the estimated forecast with the true value where the objective is to minimize the forecast error. There exist a broad range of loss functions, different loss functions are good for different purposes. Hence, the choice of loss function when evaluating forecasts should depend upon the purpose of the forecasts. Patton (2011) suggests the mean squared error (MSE) as a great loss function when evaluating the forecasts of the volatility. Patton (2011) describes the MSE to have a consistent rank among models disregarding the choice of proxy for the volatility. Hence, the MSE is used as a tool to evaluate the forecasts in this thesis. The formula for the MSE is shown in Equation 3.18 where $\hat{h}_t$ denotes the forecasted volatility and h_t denotes the approximated volatility.

$$MSE = \frac{1}{n}\sum_{t=1}^{n}\left(\hat{h}_t - h_t\right)^2 \tag{3.18}$$

Vilhelmsson (2006) characterizes the MSE as being sensitive to outliers. Furthermore, Vilhelmsson (2006) argues that the mean absolute error (MAE) is more robust to outliers. The in-sample data in our test periods includes some outliers which gives us incentives to believe that the forecast period includes some as well. Thus, we also include the MAE when evaluating our forecasts. Equation 3.19 illustrates the formula for the MAE.

$$MAE = \frac{1}{n}\sum_{t=1}^{n}\left|\hat{h}_t - h_t\right| \tag{3.19}$$

4 Data

4.1 Gold

The London Bullion Market (LBMA) Gold price is obtained through Thomson Reuters Datastream. Further, the gold price is measured in USD per troy ounce. Figure 4.1 illustrates the price trend from January 1995 to December 2016, including the in-sample and out-of-sample periods. The red color denotes Period 1 and the blue color denotes Period 2. The black color represents the data in between our test periods and it is not used in this thesis.

Figure 4.1: *The daily London Bullion Market Gold price in USD per troy ounce.*

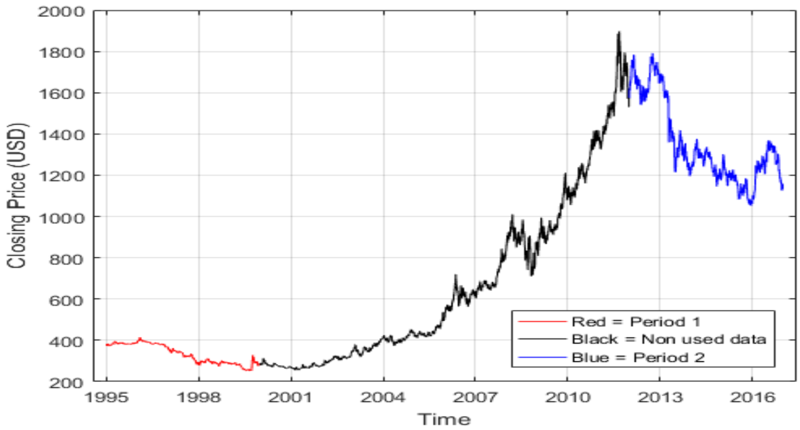

As mentioned in Section 3.3, we are interested in modelling the volatility of the returns rather than the volatility of the actual closing price. Figure 4.2 illustrates the daily returns for the in-sample period where the upper plot represents the in-sample period for Period 1 and the lower plot represents the in-sample period for Period 2. By the looks of the graphs, one may observe that the returns exhibit turbulent periods followed by relative tranquil periods also known as a clustering effect. It appears as if the clustering effect is more noticeable in Period 2. The two return processes also seem to be mean reverting with a mean close to zero.

Figure 4.2: *The two graphs illustrate the daily returns of the gold in percent. The upper graph represents Period 1 and the lower graph represents Period 2.*

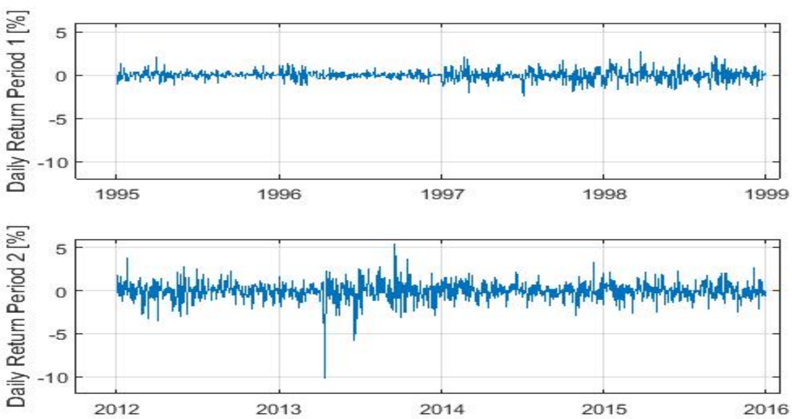

The next step is to investigate the process of the residuals from the returns. As mentioned in section 3.4 we use the ACF to exploit the structure of the residuals. Figure 4.3 shows the autocorrelation function (ACF) for the residuals in the two in-sample periods. The upper plot represents the ACF for Period 1 and the lower plot represents the ACF for Period 2. The ACF suggests that the residuals may be described by a white noise process, for both periods. This is supported by a Ljung-Box test for the residuals in which the null hypothesis of serial independence cannot be rejected, see Appendix 8.2.

Figure 4.3: *The two plots show the sample autocorrelations for the in-sample residuals from the gold returns. The upper plot corresponds to Period 1 whereas the lower plot corresponds to Period 2. The blue lines denote the confidence interval at 95 percent.*

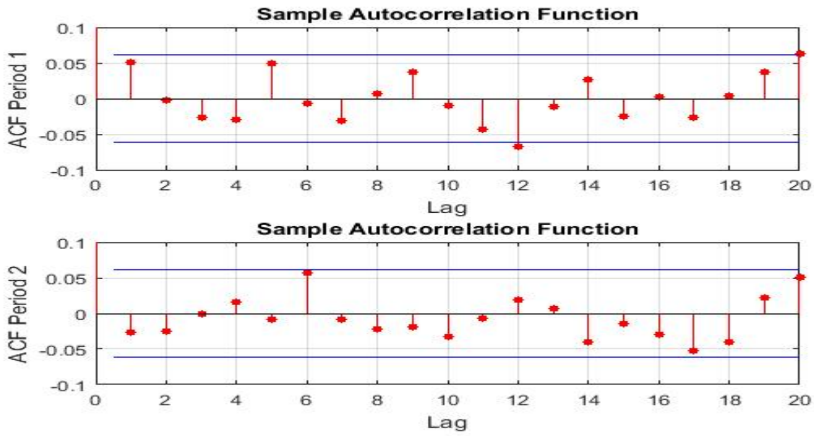

As mentioned in Section 3.6, the approximation for the volatility in this thesis is the squared residuals. One can therefore use the ACF for the squared residuals as a tool to examine whether the volatility in the in-sample period exhibits any clustering effect. Figure 4.4 shows the sample ACF for squared residuals in the two in-sample periods. The upper plot represents the ACF for the squared residuals in Period 1 and it clearly shows that the squared residuals are autocorrelated, implying that the volatility exhibits clustering effect. The lower plot illustrates the ACF for Period 2 which also demonstrates the existence for autocorrelation, although only for the first lag. This is also confirmed by Engle's ARCH test which is presented in Appendix 8.3.

Figure 4.4: *The two plots show the sample autocorrelations for the in-sample squared residuals from the gold returns. The upper plot corresponds to Period 1 whereas the lower plot corresponds to Period 2. The blue lines denote the confidence interval at 95 percent.*

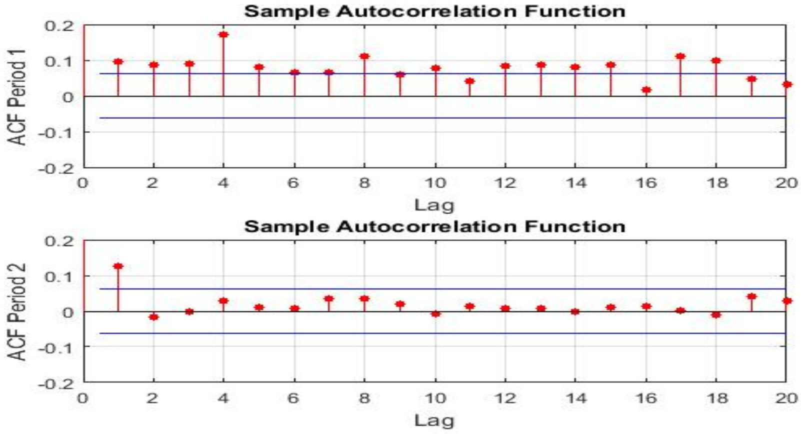

As described in Section 3.9 the parameters are estimated using two different distributions, the normal distribution and the Student's t-distribution. Figure 4.5 shows the residuals in the in-sample periods plotted against the normal distribution in a q-q plot. For both periods, the residuals seem to exhibit heavier tails than the normal distribution which goes along with the theory of the distribution for financial data described in Section 3.9. Hence, the quantiles of the residuals do not seem to match the quantiles of the normal distribution. This is also confirmed by a Jarque-Bera test where the null hypothesis that the distribution of the residuals follows the normal distribution is rejected, see Appendix 8.4.

Figure 4.5: *The left plot illustrates a q-q plot of the residuals from the gold returns in Period 1 against the normal distribution. The right plot shows the residuals from the gold returns in Period 2 against the normal distribution.*

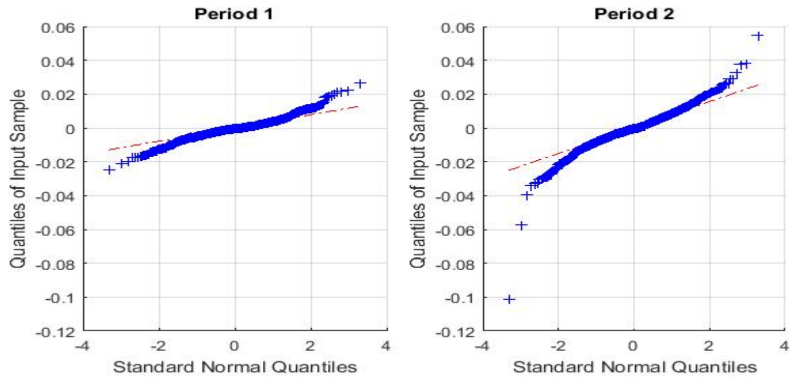

Figure 4.6 shows the quantiles of the residuals in the in-sample periods plotted against the quantiles of the Student's t-distribution in a q-q plot. One may observe that the quantiles from the residuals, for both periods, fit the quantiles from the Student's t-distribution in a more suitable way than the normal distribution since they lay closer to the linear red line in the two plots.

Figure 4.6: *The left plot shows a q-q plot of the residuals from the gold returns in Period 1 against the Student's t-distribution with three degrees of freedom. The right plot shows a q-q plot of the residuals from the gold returns in Period 2 against the Student's t-distribution with also three degrees of freedom.*

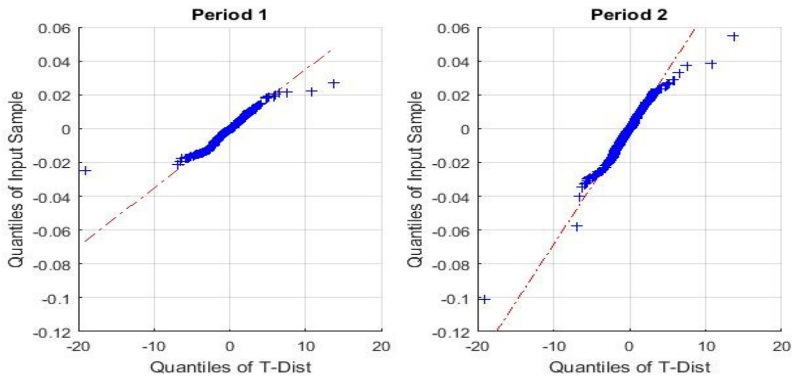

4.2 OMXS30

The data for the OMXS30 is obtained through Thomson Reuters Datastream. The OMXS30 consists of the 30 most traded stocks on the Stockholm Stock Exchange. Figure 4.7 shows the daily closing price, in USD, for the OMXS30 from January 1995 through December 2016, including the in-sample and out-of-sample periods. The red color denotes Period 1 and the blue color Period 2. The black color represents the data in between our test periods and it is not used in this thesis.

Figure 4.7: *The daily closing price for the OMSX30 in USD.*

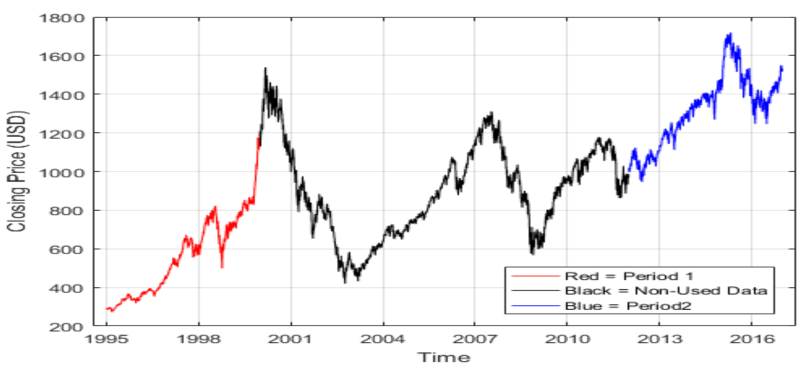

The same procedure as in Section 4.1 is used to investigate the data for the OMXS30. The daily returns for the two in-sample periods are plotted in Figure 4.8. This gives an overview of the behavior of the two return processes. The upper plot represents Period 1 and the lower plot represents Period 2. One can observe mean reversion, where the means seem to be close to zero. One can also notice clustering effects, where Period 1 seems to exhibit more turbulence at the end of the period.

Figure 4.8: *The two graphs illustrates the daily returns of the OMXS30 in percent where the upper graph represents Period 1 and the lower graph represents Period 2.*

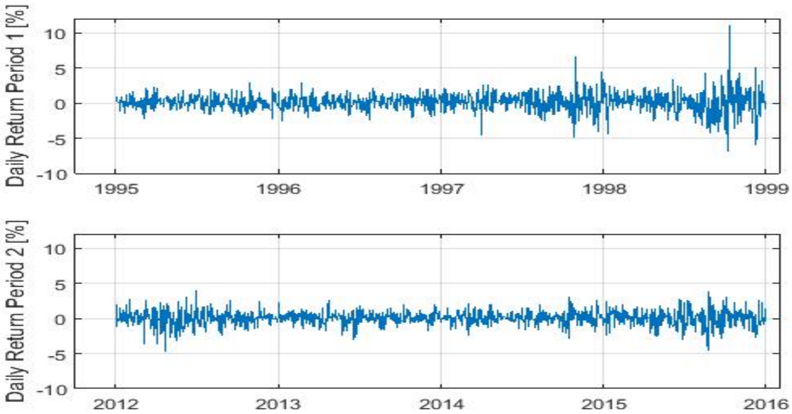

Furthermore, we plot the ACF for both periods to investigate whether the residuals from the returns can be described by a white noise process. Figure 4.9 shows the ACF for the residuals from the returns in the two in-sample periods. The upper plot represents Period 1 and the lower plot Period 2. The ACF for Period 1 suggests that the residuals may exhibit autocorrelation for lags greater than four. This is however rejected by the Ljung-Box test where we cannot reject the null hypothesis of serial independence, see Appendix 8.2. Consequently, we assume that the residuals for Period 1 is described by a white noise process. The ACF for Period 2 suggests that the residuals may exhibit autocorrelation for lag one, which possible could be described by an autoregressive process of order one. The Ljung-Box test confirms that the residuals for Period 2 exhibit autocorrelation, see Appendix 8.2.

Figure 4.9: *The two plots show the sample autocorrelation for the in-sample residuals from the return. The upper plot corresponds to Period 1 whereas the lower plot corresponds to Period 2. The blue lines denote the confidence interval at 95 percent.*

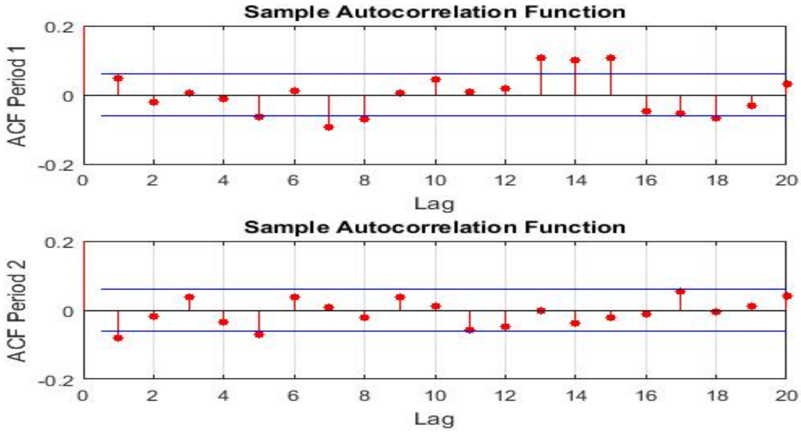

The ACF and the Ljung-Box test indicate that the returns for Period 2 may not be described by only an intercept and an error term as in Equation 3.5. Hence, the structure in the residuals needs to be modelled separately before estimating the conditional variance models. We succeed in eliminating the structure in the residuals by describing the returns with an autoregressive process of order one, see Equation 4.1. Thus the residuals can now be described by a white noise process. See Appendix 8.2, 8.3 and 8.5 for the results from a Ljung-Box test, ACF of the residuals and the estimated coefficients.

$$r_t = \alpha_0 + \alpha_1 r_{t-1} + \varepsilon_t \tag{4.1}$$

The ACF for the squared residuals is plotted in Figure 4.10 to examine whether the volatility exhibits any clustering effect. Both the upper plot, Period 1 and the lower plot, Period 2, show that the squared returns, our approximate for the volatility, are autocorrelated. Hence there exist clustering effects in the volatility for the two in-sample periods. This is also confirmed by Engle's ARCH test which is presented in Appendix 8.3.

Figure 4.10: *The two plots show the sample autocorrelation for the in-sample squared residuals from the return. The upper plot corresponds to Period 1 whereas the lower plot corresponds to Period 2. The blue lines denote the confidence interval at 95 percent*

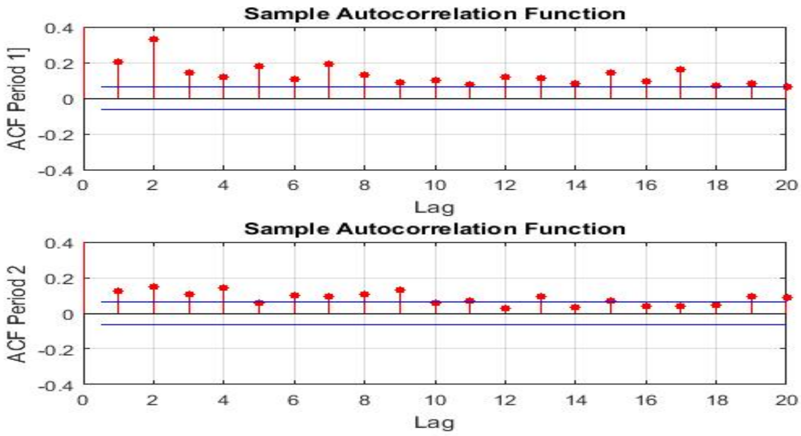

To examine the distribution of the residuals we use the q-q plot. Figure 4.11 shows the residuals from the OMSX30's returns plotted against the normal distribution. The distributions for each period have heavier tails compared to the normal distribution, and it seems like the residuals for Period 1 may have heavier tails than the residuals for Period 2. The null hypothesis in the Jarque-Bera test is rejected for both periods, meaning that the residuals do not follow the normal distribution, see Appendix 8.4.

Figure 4.11: *The left plot shows a q-q plot of the residuals from the OMXS30's returns in Period 1 against the Student's t-distribution with three degrees of freedom. The right plot shows a q-q plot of the residuals from the OMXS30's returns in Period 2 against the Student's t-distribution with also three degrees of freedom.*

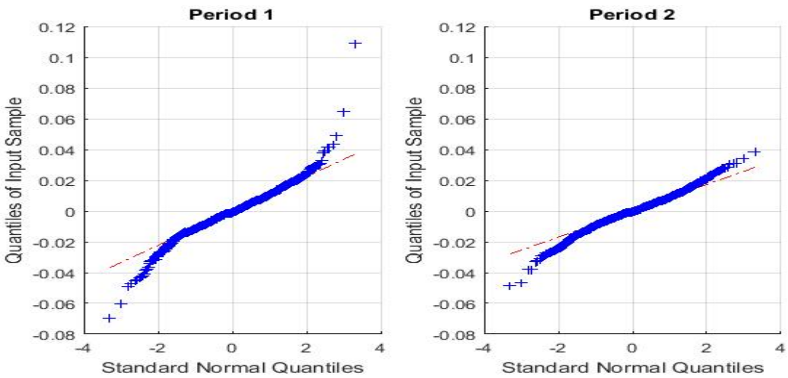

Figure 4.12 illustrates the residuals in the in-sample periods plotted against the Student's t-distribution in a q-q plot. The residuals, for both periods, seem to have lighter tails than the Student's t-distribution. However, the distributions of the residuals seem to be closer to the Student's t-distribution compared to the normal distribution.

Figure 4.12: *The left plot shows a q-q plot of the residuals from the OMXS30's returns in Period 1 against the Student's t-distribution with three degrees of freedom. The right plot shows a q-q plot of the residuals from the OMXS30's returns in Period 2 against the Student's t-distribution with also three degrees of freedom.*

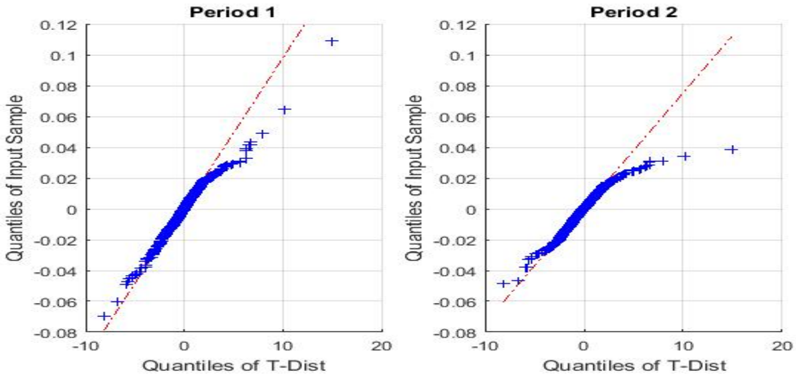

4.3 USD/EUR

The data is obtained through Thomson Reuters Datastream and consists of the daily closing price of the exchange rate USD/EUR. Figure 4.13 shows the daily closing price for the USD/EUR exchange rate from January 1995 through December 2016, including both in-sample and out-of-sample periods. The red color represents Period 1 and the blue color represents Period 2. The black color represents the data in between our test periods and it is not used in this thesis.

Figure 4.13: *The daily closing price for the USD/EUR exchange rate.*

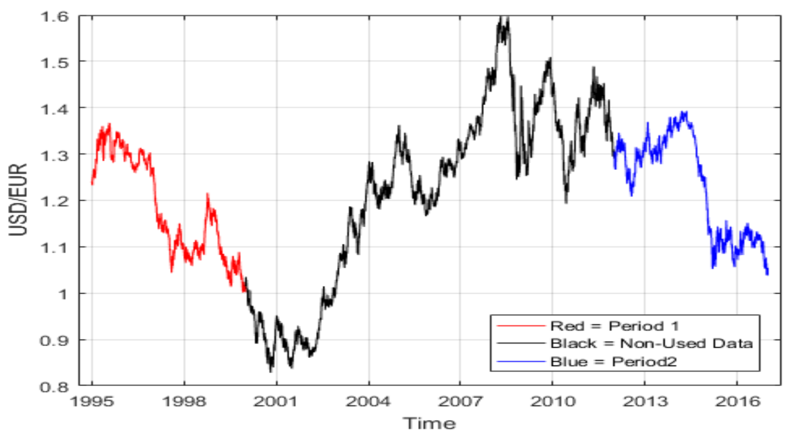

The daily returns for the USD/EUR exchange rate for the two period's in-sample periods are presented in Figure 4.14. One can observe periods of turbulence followed by relative calm periods. The mean seems to be close to zero for both periods. The returns for the two periods also seem to be mean reverting.

Figure 4.14: *The two graphs illustrates the daily returns of the USD/EUR exchange rate in percent where the upper graph represents Period 1 and the lower graph Period 2.*

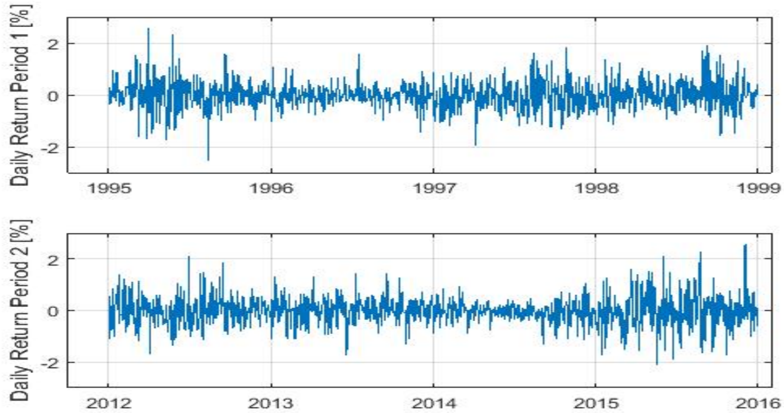

We then follow the same procedure as for the two previously indexes and plot the ACF for both periods to examine if there is any structure left in the residuals or if they can be described by white noise processes. Figure 4.15 illustrates the ACF for the residuals in the two in-sample periods. The two plots indicate that the residuals, for both periods, may be described by white noise processes. This is also confirmed by a Ljung-Box test for the residuals in which the null hypothesis of serial independence cannot be rejected, see Appendix 8.2.

Figure 4.15: *The two plots show the sample autocorrelation for the in-sample residuals from the return. The upper plot corresponds to Period 1 whereas the lower plot corresponds to Period 2. The blue lines denote the confidence interval at 95 percent.*

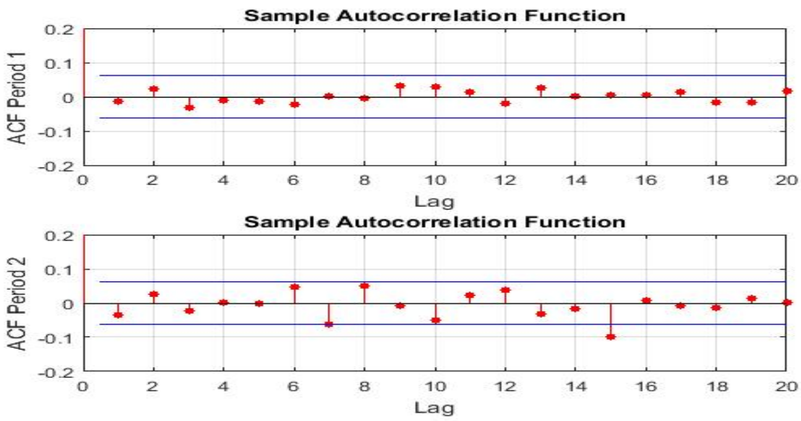

Figure 4.16 shows the ACF for the squared residuals. The ACF suggests that the volatility for both periods exhibit autocorrelation, hence there exist clustering in the volatility. The clustering effect is also confirmed by Engle's ARCH test which is presented in Appendix 8.3.

Figure 4.16: *The two plots show the sample autocorrelation for the in-sample squared residuals from the return. The upper plot corresponds to Period 1 whereas the lower plot corresponds to Period 2. The blue lines denote the confidence interval at 95 percent.*

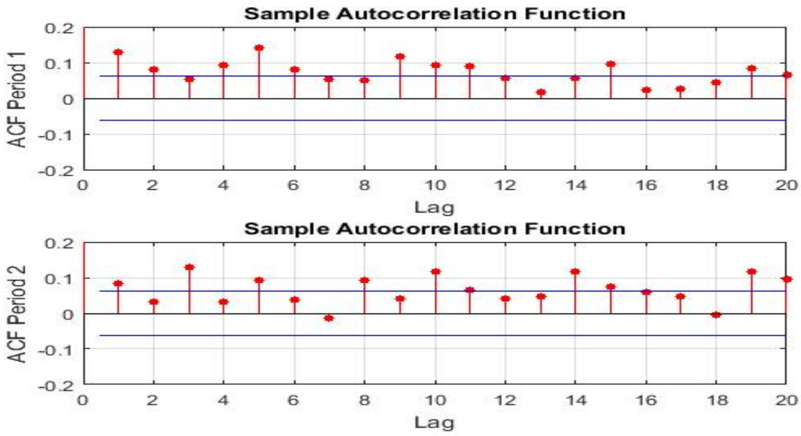

The distribution of the residuals is examined in the same way as for the other two indexes. Starting with a q-q plot where the residuals are plotted against the normal distribution which is illustrated in Figure 4.17. The distributions of the residuals, for both periods, seem to have fatter tails than the normal distribution. Furthermore, the null hypothesis that the distributions of the residuals follow the normal distribution is rejected for both periods, see Appendix 8.4.

Figure 4.17: *The left plot shows a q-q plot of the residuals from the USD/EUR's returns in Period 1 against the Student's t-distribution with three degrees of freedom. The right plot shows a q-q plot of the residuals from the USD/EUR's returns in Period 2 against the Student's t-distribution with also three degrees of freedom.*

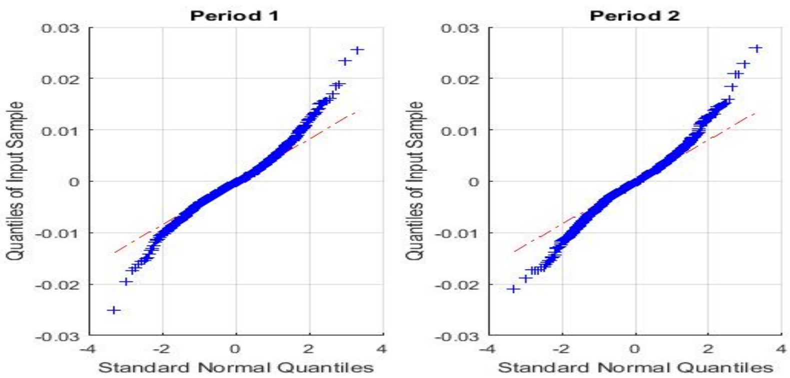

Figure 4.18 shows the residuals in the in-sample periods plotted against the Student's t-distribution in a q-q plot. The distributions of the residuals seem to be similar to the distribution of the returns' residuals for the OMXS30 and the gold. They seem to fit the Student's t-distribution better than the normal distribution even though it appears as if they exhibit lighter tails compared to the Student's t-distribution.

Figure 4.18: *The left plot shows a q-q plot of the residuals from the USD/EUR's returns in Period 1 against the Student's t-distribution with three degrees of freedom. The right plot shows a q-q plot of the residuals from the USD/EUR's returns in Period 2 against the Student's t-distribution with also three degrees of freedom.*

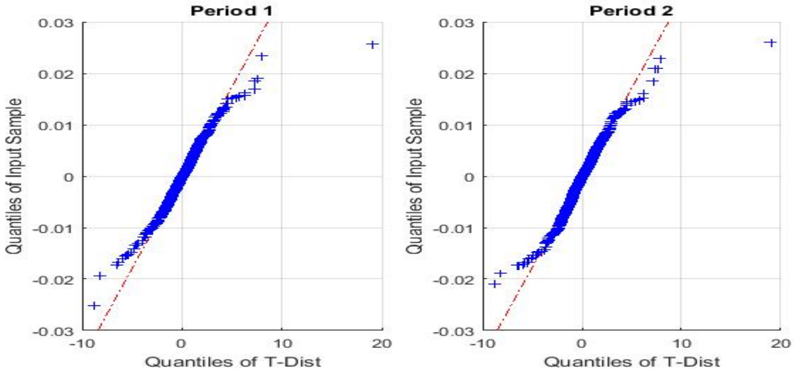

5 Empirical Results and Analysis

5.1 Parameter Estimation

The models are evaluated based on their out-of-sample forecast accuracy rather than their in-sample fit. Hence, insignificant parameters and violated restrictions for the models are not as important as if we would examine the models' in-sample fit. However, such models may indicate a poor out-of-sample performance which may be seen when evaluating the forecasts.

5.1.1 Gold

Table 5.1 and Table 5.2 summarize the estimated parameters in the two in-sample periods. In Period 1, the estimated α_1 in the ARCH model is quite high using the Student's t-distribution compared to the same parameters in the ARCH model using the normal distribution. A high α_1 in the ARCH model implies that a shock in the current period will have a large effect in the next period when forecasting the conditional volatility. One can also observe that the GARCH model, in Period 1 using the Student's t-distribution, violates its restrictions since the sum of α_1 and β_1 is greater than one. Further, the sum of α_1 and β_1 for the rest of the GARCH models are close to unity. This indicates that the conditional volatility is very persistent and therefore the effect of a shock will show very little tendency to dissipate. Furthermore, the leverage effects in the EGARCH models are higher in Period 2 compared to Period 1. However, in both periods and distributions, the estimated leverage effects for the GJR-GARCH and the EGARCH model indicate similar effects; a negative shock will increase the volatility for the next period more than a positive shock. The difference between the estimates when assuming the two different distributions are in general quite small for the two periods. One may also observe that the ARCH effect in the models is very small compared to the GARCH effect. This implies that the effect of a shock will have little impact on the next period's volatility.

Gold (Period 1)	
ARCH(1) normal dist.	$h_t = 2.8300\text{E-}05 + 0.1092\varepsilon_{t-1}^2$
ARCH(1) t-dist.	$h_t = 3.7000\text{E-}05 + 0.6509\varepsilon_{t-1}^2$
GARCH(1,1) normal dist.	$h_t = 1.9600\text{E-}07 + 0.0495\varepsilon_{t-1}^2 + 0.9457h_{t-1}$
GARCH(1,1) t-dist.	$h_t = 3.4600\text{E-}07 + 0.1179\varepsilon_{t-1}^2 + 0.8943h_{t-1}$
IGARCH(1,1) normal dist.	$h_t = (1 - 0.9618)\varepsilon_{t-1}^2 + 0.9618h_{t-1}$
IGARCH(1,1) t-dist.	$h_t = (1 - 1.0026)\varepsilon_{t-1}^2 + 1.0026h_{t-1}$
EGARCH(1,1) normal dist.	$ln(h_t) = -0.2647 - 0.0046\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.1272\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.9834ln(h_{t-1})$
EGARCH(1,1) t-dist.	$ln(h_t) = -0.4070 - 0.0102\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.2768\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.9784ln(h_{t-1})$
GJRGARCH(1,1) normal dist.	$h_t = 2.0000\text{E-}07 + (0.0438 + 0.0114I_{t-1})\varepsilon_{t-1}^2 + 0.9456h_{t-1}$
GJRGARCH(1,1) t-dist.	$h_t = 3.6100\text{E-}07 + (0.0871 + 0.0282I_{t-1})\varepsilon_{t-1}^2 + 0.8988h_{t-1}$

Table 5.1: *Period 1's estimated parameters in the conditional variance models for the gold. The corresponding p-values are presented in Appendix 8.7.*

Gold (Period 2)	
ARCH(1) normal dist.	$h_t = 9.2500\text{E-}05 + 0.0996\varepsilon_{t-1}^2$
ARCH(1) t-dist.	$h_t = 0.0001 + 0.0272\varepsilon_{t-1}^2$
GARCH(1,1) normal dist.	$h_t = 3.6400\text{E-}06 + 0.0507\varepsilon_{t-1}^2 + 0.9164h_{t-1}$
GARCH(1,1) t-dist.	$h_t = 1.1100\text{E-}06 + 0.0175\varepsilon_{t-1}^2 + 0.9724h_{t-1}$
IGARCH(1,1) normal dist.	$h_t = (1 - 1.0007)\varepsilon_{t-1}^2 + 1.0007h_{t-1}$
IGARCH(1,1) t-dist.	$h_t = (1 - 0.9826)\varepsilon_{t-1}^2 + 0.9826h_{t-1}$
EGARCH(1,1) normal dist.	$ln(h_t) = -0.7101 - 0.0863\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.1209\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.9324ln(h_{t-1})$
EGARCH(1,1) t-dist.	$ln(h_t) = -7.2096 - 0.1702\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} - 0.1194\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.1987ln(h_{t-1})$
GJRGARCH(1,1) normal dist.	$h_t = 7.7282\text{E-}06 + (0.0020 + 0.0920I_{t-1})\varepsilon_{t-1}^2 + 0.8763h_{t-1}$
GJRGARCH(1,1) t-dist.	$h_t = 2.2800\text{E-}06 + (0.0272 + 0.0013I_{t-1})\varepsilon_{t-1}^2 + 0.9570h_{t-1}$

Table 5.2: *Period 2's estimated parameters in the conditional variance models for the gold. The corresponding p-values are presented in Appendix 8.7.*

5.1.2 OMXS30

The results for the estimated parameters in the conditional variance models are shown in Table 5.3 and Table 5.4. In period 1, one can observe a reduced effect in α_1 when using a heavier tailed distribution. Thus, for the second period the opposite behavior is observed. All the models, for both periods, satisfy their restrictions. However, at least one of the estimated parameters in all GJR-GARCH models is insignificant, see Appendix 8.7. Regarding the GARCH models, their parameters are all very close to one. This is often the case when estimating GARCH (1, 1) models over a long span of financial data, see Section 3.8.3. As for both periods the GARCH effect, β_1, is a bit higher for the normal distribution compared to the student's t-distribution. Further, the estimated parameters in the EGARCH and GJR-GARCH models indicate that a negative shock will have a greater impact on the conditional variance than a positive shock. Similarly to the estimated parameters in Section 5.1.1, the differences between the estimated parameters between the two distributions are very small.

OMXS30 (Period 1)

ARCH(1) normal dist.	$h_t = 0.0001 + 0.3358\varepsilon_{t-1}^2$
ARCH(1) t-dist.	$h_t = 0.0001 + 0.3232\varepsilon_{t-1}^2$
GARCH(1,1) normal dist.	$h_t = 4.3600\text{E-}06 + 0.1222\varepsilon_{t-1}^2 + 0.8533h_{t-1}$
GARCH(1,1) t-dist.	$h_t = 4.7100\text{E-}06 + 0.1199\varepsilon_{t-1}^2 + 0.8522h_{t-1}$
IGARCH(1,1) normal dist.	$h_t = (1 - 0.9155)\varepsilon_{t-1}^2 + 0.9155h_{t-1}$
IGARCH(1,1) t-dist.	$h_t = (1 - 0.9172)\varepsilon_{t-1}^2 + 0.9172h_{t-1}$
EGARCH(1,1) normal dist.	$ln(h_t) = -0.4296 - 0.0850\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.1907\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.9685ln(h_{t-1})$
EGARCH(1,1) t-dist.	$ln(h_t) = -0.4585 - 0.0872\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.1929\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.9655ln(h_{t-1})$
GJRGARCH(1,1) normal dist.	$h_t = 5.6700\text{E-}06 + (0.0403 + 0.1616I_{t-1})\varepsilon_{t-1}^2 + 0.8457h_{t-1}$
GJRGARCH(1,1) t-dist.	$h_t = 6.0700\text{E-}06 + (0.0386 + 0.1641I_{t-1})\varepsilon_{t-1}^2 + 0.8427h_{t-1}$

Table 5.3: *Period 1's estimated parameters in the conditional variance models for OMXS30. The corresponding p-values are presented in Appendix 8.7.*

OMXS30 (Period 2)

ARCH(1) normal dist.	$h_t = 9.4200\text{E-}05 + 0.1080\varepsilon_{t-1}^2$
ARCH(1) t-dist.	$h_t = 9.8500\text{E-}05 + 0.1201\varepsilon_{t-1}^2$
GARCH(1,1) normal dist.	$h_t = 1.5300\text{E-}06 + 0.0505\varepsilon_{t-1}^2 + 0.9350h_{t-1}$
GARCH(1,1) t-dist.	$h_t = 1.9100\text{E-}06 + 0.0667\varepsilon_{t-1}^2 + 0.9179h_{t-1}$
IGARCH(1,1) normal dist.	$h_t = (1 - 0.9608)\varepsilon_{t-1}^2 + 0.9608h_{t-1}$
IGARCH(1,1) t-dist.	$h_t = (1 - 0.9535)\varepsilon_{t-1}^2 + 0.9535h_{t-1}$
EGARCH(1,1) normal dist.	$ln(h_t) = -0.5715 - 0.1529\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.0979\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.9468ln(h_{t-1})$
EGARCH(1,1) t-dist.	$ln(h_t) = -0.5707 - 0.1689\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.1222\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.9489ln(h_{t-1})$
GJRGARCH(1,1) normal dist.	$h_t = 4.5100\text{E-}06 + (0.000 + 0.1672I_{t-1})\varepsilon_{t-1}^2 + 0.8709h_{t-1}$
GJRGARCH(1,1) t-dist.	$h_t = 3.7300\text{E-}06 + (0.000 + 0.1848I_{t-1})\varepsilon_{t-1}^2 + 0.8734h_{t-1}$

Table 5.4: *Period 2's estimated parameters in the conditional variance models for OMXS30. The corresponding p-values are presented in Appendix 8.7.*

5.1.3 USD/EUR

Table 5.5 and Table 5.6 summarize the two in-sample periods estimated parameters in the conditional variance models for the USD/EUR. The tables show that the clustering effects in the ARCH models for Period 1 are slightly higher than the clustering effects in the ARCH models for Period 2. This implies that the effect of a shock in 1995 to 1998 is higher compared to the effect of a shock in 2012 to 2015. One can also observe that the sum of α_1 and β_1, in all GARCH models for the USD/EUR, are close to one. In general, the estimates from the different distribution are similar. An interesting result is that the estimated leverage effects vary between the periods and distributions. For example, the EGARCH models in Period 1 suggest that positive shocks affect the volatility more than negative shocks. In contrast, the GJR-GARCH models in Period 2 indicate the opposite, negative shocks have a greater effect on the volatility compared to positive shocks. Further, one can observe that the leverage effects in the GJR-GARCH models for Period 1 and the EGARCH models for Period 2 show mixed results depending on which distribution one assumes.

USD/EUR (Period 1)

ARCH(1) normal dist.	$h_t = 2.4300\text{E-}05 + 0.1155\varepsilon_{t-1}^2$
ARCH(1) t-dist.	$h_t = 2.4900\text{E-}05 + 0.1484\varepsilon_{t-1}^2$
GARCH(1,1) normal dist.	$h_t = 6.6900\text{E-}07 + 0.0535\varepsilon_{t-1}^2 + 0.9225h_{t-1}$
GARCH(1,1) t-dist.	$h_t = 3.8100\text{E-}07 + 0.0634\varepsilon_{t-1}^2 + 0.9272h_{t-1}$
IGARCH(1,1) normal dist.	$h_t = (1 - 0.9630)\varepsilon_{t-1}^2 + 0.9630h_{t-1}$
IGARCH(1,1) t-dist.	$h_t = (1 - 0.9491)\varepsilon_{t-1}^2 + 0.9491h_{t-1}$
EGARCH(1,1) normal dist.	$ln(h_t) = -0.3194 + 0.0092\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.1074\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.9773ln(h_{t-1})$
EGARCH(1,1) t-dist.	$ln(h_t) = -0.2692 + 0.0072\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.1333\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 0.9838ln(h_{t-1})$
GJRGARCH(1,1) normal dist.	$h_t = 4.0200\text{E-}07 + (0.0434 + 0.0025I_{t-1})\varepsilon_{t-1}^2 + 0.9412h_{t-1}$
GJRGARCH(1,1) t-dist.	$h_t = 2.9100\text{E-}07 + (0.0619 - 0.0042I_{t-1})\varepsilon_{t-1}^2 + 0.9334h_{t-1}$

Table 5.5: *Period 1's estimated parameters in the conditional variance models for USD/EUR. The corresponding p-values are presented in Appendix 8.7.*

USD/EUR (Period 2)

ARCH(1) normal dist.	$h_t = 2.8000\text{E-}05 + 0.0748\varepsilon_{t-1}^2$
ARCH(1) t-dist.	$h_t = 2.9500\text{E-}05 + 0.1145\varepsilon_{t-1}^2$
GARCH(1,1) normal dist.	$h_t = 7.9500\text{E-}08 + 0.0280\varepsilon_{t-1}^2 + 0.9698h_{t-1}$
GARCH(1,1) t-dist.	$h_t = 5.1000\text{E-}08 + 0.0343\varepsilon_{t-1}^2 + 0.9655h_{t-1}$
IGARCH(1,1) normal dist.	$h_t = (1 - 0.9738)\varepsilon_{t-1}^2 + 0.9738h_{t-1}$
IGARCH(1,1) t-dist.	$h_t = (1 - 0.9684)\varepsilon_{t-1}^2 + 0.9684h_{t-1}$
EGARCH(1,1) normal dist.	$ln(h_t) = 0.006 - 0.0518\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} - 0.0078\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert + 1.0000ln(h_{t-1})$
EGARCH(1,1) t-dist.	$ln(h_t) = -15.5055 + 0.1455\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + 0.1929\lvert\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}}\rvert - 0.4857ln(h_{t-1})$
GJRGARCH(1,1) normal dist.	$h_t = 2.0000\text{E-}07 + (0.0183 + 0.0214I_{t-1})\varepsilon_{t-1}^2 + 0.9643h_{t-1}$
GJRGARCH(1,1) t-dist.	$h_t = 2.0000\text{E-}07 + (0.0264 + 0.0216I_{t-1})\varepsilon_{t-1}^2 + 0.9575h_{t-1}$

Table 5.6: *Period 2's estimated parameters in the conditional variance models for USD/EUR. The corresponding p-values are presented in Appendix 8.7.*

5.2 Forecast evaluation

The error statistics from the forecasts in Period 1 are presented in Table 5.7. The underlined values represents the most accurate model for each index and loss function. When using the MAE as critera, the ARCH model when assuming the errors to follow the normal distribution, is the most accurate model for both gold and USD/EUR. Further, the MAE indicates that the EGARCH model is the most accurate model to predict the volatility in the OMXS30. Hence, these results could imply that the OMXS30 exhibits more asymmetry in its market's response to shocks compared to the gold and the USD/EUR exchange rate. Furthermore, the MSE shows that the more advanced models tend to have a more accurate forecast precision when forecasting the volatility in the three assets.

	Period 1			
	Mean Absolute Error		Mean Squared Error	
	normal dist.	t-dist.	normal dist.	t-dist.
Gold				
ARCH(1)	<u>9.515E-05</u>	1.259E-04	1.563E-07	1.714E-07
GARCH(1,1)	1.142E-04	1.288E-04	1.562E-07	1.581E-07
IGARCH(1,1)	1.147E-04	9.868E-05	1.577E-07	1.681E-07
EGARCH(1,1)	1.023E-04	1.153E-04	1.567E-07	<u>1.518E-07</u>
GJRGARCH	1.131E-04	1.198E-04	1.565E-07	1.552E-07
OMXS30				
ARCH(1)	1.612E-04	1.607E-04	8.131E-08	8.084E-08
GARCH(1,1)	1.601E-04	1.593E-04	7.839E-08	7.820E-08
IGARCH(1,1)	1.582E-04	1.580E-04	7.806E-08	<u>7.801E-08</u>
EGARCH(1,1)	1.544E-04	<u>1.543E-04</u>	7.839E-08	7.840E-08
GJRGARCH(1,1)	1.551E-04	1.547E-04	7.907E-08	7.905E-08
USD/EUR				
ARCH(1)	<u>3.204E-05</u>	3.297E-05	3.398E-09	3.432E-09
GARCH(1,1)	3.247E-05	3.344E-05	<u>3.334E-09</u>	3.367E-09
IGARCH(1,1)	3.285E-05	3.301E-05	3.387E-09	3.416E-09
EGARCH(1,1)	3.287E-05	3.376E-05	3.379E-09	3.406E-09
GJRGARCH(1,1)	3.276E-05	3.369E-05	3.387E-09	3.432E-09

Table 5.7: *Period 1's error statistics for gold, OMXS30 and USD/EUR.*

Table 5.8 shows the MAE and MSE for the compared volatility models in Period 2. The two loss functions shows different results for the three assets. The MAE indicates that the ARCH model is the most accurate forecasting model for gold and OMXS30. The EGARCH model is the most accurate for USD/EUR using the same criteria. Further, the MSE selects the GARCH model for gold, the EGARCH model for OMXS30 and the ARCH model for USD/EUR.

	Period 2			
	Mean Absolute Error		Mean Squared Error	
	normal dist.	t-dist.	normal dist.	t-dist
Gold				
ARCH(1)	9.697E-05	9.736E-05	5.870E-08	5.932E-08
GARCH(1,1)	1.139E-04	1.144E-04	4.986E-08	4.989E-08
IGARCH(1,1)	1.162E-04	1.142E-04	4.997E-08	5.009E-08
EGARCH(1,1)	1.153E-04	1.199E-04	5.159E-08	5.001E-08
GJRGARCH	1.130E-04	1.210E-04	5.130E-08	5.023E-08
OMXS30				
ARCH(1)	1.736E-04	1.760E-04	3.015E-07	3.013E-07
GARCH(1,1)	1.869E-04	1.885E-04	2.918E-07	2.917E-07
IGARCH(1,1)	1.987E-04	1.979E-04	2.951E-07	2.949E-07
EGARCH(1,1)	1.630E-04	1.688E-04	2.886E-07	2.905E-07
GJRGARCH(1,1)	1.862E-04	1.925E-04	3.011E-07	3.060E-07
USD/EUR				
ARCH(1)	3.415E-05	3.531E-05	3.954E-09	3.973E-09
GARCH(1,1)	3.510E-05	3.553E-05	3.996E-09	4.009E-09
IGARCH(1,1)	3.505E-05	3.501E-05	4.001E-09	4.006E-09
EGARCH(1,1)	2.960E-05	3.569E-05	4.078E-09	4.045E-09
GJRGARCH(1,1)	3.455E-05	3.514E-05	3.997E-09	4.008E-09

Table 5.8: *Period 2's error statistics for gold, OMXS30 and USD/EUR.*

5.2.1 Analysing the forecast evaluation

The results indicate that none of the models seem superior the other models when forecasting the conditional volatility for the different assets and periods. The ranking between the models forecast accuracy differs depending on which loss function is used. Hence, the results are inconclusive and difficult to interpret. Franses and Van Dijk (1996) show that the GJR-GARCH model has a poor forecast ability and therefore not recommended for forecasting the conditional volatility. They examine different stock indexes, including a Swedish stock index. This is in line with our findings, the GJR-GARCH model is always outperformed by another model, not only for the OMXS30 but for all indexes and periods. One reason for the GJR-GARCH model always being outperformed could be its poor in-sample fit. At least one of the estimated parameters in the GJR-GARCH models, for all assets and periods, is insignificant.

Furthermore, Mckenzie and Mitchell (2002) examine the conditional volatility of 17 different exchange rates. They find the GARCH(1,1) model to be the most suitable model when modelling the conditional volatility of exchange rate data. However, they also mention that models accounting for leverage effects may perform better if there exist any asymmetry in the market's response to shocks. Our results, for Period 1, indicate that the GARCH model perform well compared to other models for the exchange rate. When evaluating the forecasts through MSE, the GARCH model is the most accurate model. However, when evaluating the forecasts using MSE for the second period, the results for the exchange rate indicate that the EGARCH model has the highest forecast accuracy. By the looks of the estimated parameters in the two test periods, the exchange rate in Period 1 seem to exhibit more asymmetry in its market's response to shocks compared to Period 2. This could explain our results and hence support Mckenzie and Mitchell's (2002) findings.

Hansens and Lunde's (2001) results, when forecasting the volatility of the IBM stock return, indicate that the EGARCH model performs better when assuming the residuals to follow the Student's t-distribution rather than the normal distribution. In comparison to Hansen and Lunde's (2001) findings, our results vary. For the OMXS30 in Period 1, the MAE suggests that the EGARCH model performs better when assuming normal distribution compared to the Student's t-distribution. However, the opposite is shown when evaluating the forecasts using the MSE. As for Period 2, the EGARCH model has a higher forecast accuracy when assuming the residuals to follow the normal distribution compared the Student's t-distribution. Furthermore, Hansen and Lunde's (2001) findings show that the IGARCH model has a poor forecast accuracy when forecasting the volatility of the IBM stock return. This is in line with our results for the OMXS30 in Period 2; the IGARCH model is always outperformed by another model. However, in Period 1, the IGARCH model is the best performing model for the OMXS30 using MSE as criteria.

Trück and Liang (2012) examines the volatility of the gold. Their results are however quite inconclusive. They find none of their models to be superior the others. Our forecast results, when evaluated using the MAE, indicate that the ARCH model have the highest forecast accuracy when forecasting the volatility of the gold returns. However, the forecast errors are very similar for all models and should therefore be treated carefully.

Poon and Granger (2003) argues that the squared residuals are a noisy proxy for the volatility. Further, they describe that this noisy proxy may result in small differences between the forecast estimations. Hence, our choice of proxy for the latent volatility may explain our inconsistent results. The fact that the number of outliers may differ between the two test periods could be another reason for our inconsistent results. The loss functions used in this thesis are different in the sense of their robustness to outliers. As mentioned in Section 3.13, the MAE is more robust to outliers whereas the MSE is more penalizing. This implies that models who succeed in capturing these outliers will be ranked higher using the MSE compared to the MAE. Hence, the inconsistent result between the test periods may be explained by possible differences in the number of outliers in the two test periods. Possible outliers could also be a reason for why the ARCH(1) model performs well. The ARCH(1) model accounts for less persistence compared to the more advanced models, and hence it is less affected by outliers.

6 Summary and Concluding Remarks

We compare five different conditional variance models using daily returns from OMXS30, USD/EUR and gold for two different time periods assuming two different distributions. The models are compared in terms of their out-of-sample volatility forecasting ability. As a proxy for the volatility, we use squared returns based on inter-day returns. Furthermore, to uncover patterns and extract information from the different data sets we test for autocorrelation and stationarity using an ADF-test and a Ljung-Box test. Further, we test for ARCH-effects using an Engel's ARCH test. The Ljung-Box test indicates that the residuals for OMXS30 in period 2 are autocorrelated. This structure is removed by including an AR(1) process in the conditional variance models. The forecast accuracy of the five models is evaluated using two loss functions, the mean squared error and the mean absolute error.

In terms of out-of-sample forecasting, the loss functions select different models for the same period and between periods. This makes our results very inconclusive and it is therefore difficult to conclude if any model is superior to the other models. Our inconsistent results could support our hypothesis that the most accurate forecast model may differ between the periods. However, several factors could have affected our inconsistent result. One explanation could be our choice of proxy for the volatility. This noisy proxy may result in small differences between the forecasts, which make them indistinguishable. Another factor might be the difference in the numbers of possible outliers in the two out-of-sample periods. Some model may capture these outliers better than other models and therefore affect our result.

This thesis contributes to the literature by comparing the forecast accuracy between five different conditional variance models during two different periods. This thesis is limited to a proxy for the volatility based on inter-day data. Using intra-day data to approximate the latent volatility could lead to a less inconsistent result. Hence, the models' forecast ability could be more distinguishable and easier to interpret. An interesting topic for further research could be the comparison between the volatility of the OMXS30 and the SIXVX. The SIXVX is a volatility index for the OMXS30 and it is derived from prices of the OMXS30 index options. One could therefore forecast the volatility of the OMXS30 using conditional variance models and then compare the results with the SIXVX.

7 Bibliography

Andersen, Torben. G. & Bollerslev, Tim. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39 (4), pp. 885-905.

Aronson, David. (2011). *Evidence-Based Technical Analysis: Applying the Scientific Method and Statistical Inference to Trading Signals.* New Jersey: John Wiley & Sons Inc. E-book.

Bollerslev, Tim. (1986). Generalized Autoregressive Conditional Heteroscedasticity. *Journal of Econometrics*, 31, pp.307:327.

Bradley, Brendan. O., & Taqqu, Murad.S.(2003). Chapter 2: Financial risk and heavy tails. In: Rachev, Svetlozar .T (editor.). *Handbook Of Heavy Tailed Distributions In Finance.* Amsterdam: North Holland Publishing Co. pp. 35-103.

Brownless , Christian., Engle, Robert & Kelly, Bryan. (2011). A practical guide to volatility forecasting through calm and storm. *The Journal of Risk*, 14(2), pp.3-22.

Enders, Walter. (2014). *Applied Econometric Time Series.* (4th ed.). New Jersey: John Wiley & Sons Inc.

Engle, Robert .F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50(4), pp.987-1008.

Franses, Philip & Van Dijk, Dick. (1996). Forecasting Stock Market Volatility Using (Non-Linear) Garch Models. *Journal of Forecasting*, 15(3), pp.229-235.

Glosten, L. R., Jagannathan, R. & 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, pp.1779–1801.

Hansen, Peter & Lunde, Asger. (2001). A Comparison of Volatility Models: Does Anything Beat a GARCH (1, 1)?. *Working Paper*, Brown University.

Hull, John. C. (2011). *Options, Futures and Other Derivatives.* (8th ed.). New Jersey: Pearson Education.

Keya, Rani Das & A.H.M: Rahmatulla Imon. (2016).A brief Review of Tests for Normality. *American Journal of Theoretical and Applied Statistics*, 5(1), pp.5-12.

McKenzie, Michael & Mitchell, Heather. (2002). Generalized Asymmetric Power ARCH modelling of Exchange Rate Volatility. *Applied Financial Economics*, 12(8), pp.555-564.

Meucci, Attilio. (2005). *Risk and Asset Allocation.* Heidelberg: Springer-Verlag.

Miller, Irwin. & Miller, Marylees. (2004). *John E. Freund's Mathematical Statistics: With Applications.* (7th ed.). New Jersey: Pearson Education.

Nelson, Daniel. B. (1990). Stationarity and Persistence in the GARCH (1, 1) Model. *Econometric Theory*, 6(3), pp.318-334.

Nelson, Daniel. B. (1991). Conditional Heteroscedasticity in Asset Returns: A New Approach. *Econometrica*, 59(2). pp. 347-370.

Nelson, Daniel. B. & Foster, Dean. (1995). Filtering and Forecasting with Misspecified ARCH Models II: Making the Right Forecast with the Wrong Model. *Journal of Econometrics*, 67(2), pp.303:335.

Patton, A.J. (2011).Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160. pp.246-256.

Poon, Ser-Huang. & Grange, Clive.W. J. (2003). Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature*, 41(2), pp.478-539.
Trück, Stefan & Liang, Kevin.(2012). Modelling and forecasting volatility in the gold market. *International Journal of Banking and Finance*,9(1),pp.48-80.
Tsay, Ruey.S. (2001). *Analysis of Financial Time Series.* (1st ed.). New Jersey: John Wiley & Sons Inc.
Tse, Y.K. (1998). The Conditional Heteroscedasticity of the Yen-Dollar Exchange Rate. *Journal of Applied Econometrics*, 13, pp.49-55.
Verbeek, Marno. (2004). *A Guide to Modern Econometrics.* (2nd ed.). New Jersey: John Wiley & Sons Inc.
Vilhelmsson, Anders. (2006). Garch Forecasting Performance under Different Distribution Assumptions. *Journal of Forecasting*, 25, pp.561-578.
Zivot, Eric. (2009). Practical Issues in the Analysis of Univariate GARCH Models. In: Andersen, T. G., R. A. Davis, J. Kreib, and T. Mikosch (ed.). *Handbook of Financial Time Series*. Heidelberg: Springer-Verlag, pp.113-155.

8 Appendix

8.1 Augmented Dickey-Fuller Test

Gold

	Period 1	Period 2
P-value	0.107	0.207
F-statistic	-1.586	-1.217
Critical Value	-1.942	-1.942

Table 8.1: *P-values, F-statistics and Critical Values from the ADF-test for the gold price in Period 1 and Period 2.*

OMXS30

	Period 1	Period 2
P-value	0.956	0.882
F-statistic	1.351	0.785
Critical Value	-1.942	-1.942

Table 8.2: *P-values, F-statistics and Critical Values from the ADF-test for the OMXS30 price in Period 1 and Period 2.*

USD/EUR

	Period 1	Period 2
P-value	0.503	0.284
F-statistic	-0.406	1.005
Critical Value	-1.942	-1.942

Table 8.3: *P-values, F-statistics and Critical Values from the ADF-test for the USD/EUR Exchange rate in Period 1 and Period 2.*

8.2 Ljung-Box Test

Gold

	Period 1	Period 2
P-value	0.233	0.887
Q-statistic	6.836	1.713
Critical Value	11.071	11.071

Table 8.4: *P-values, Q-Statistics and Critical values from the Ljung-Box test on the residuals from the gold returns in Period 1 and Period 2.*

OMXS30		
	Period 1	Period 2
P-value	0.193	0.010
Q-statistic	7.394	15.065
Critical Value	11.071	11.071

Table 8.5: *P-values, Q-Statistic and Critical values from the Ljung-Box test on the residuals from the OMXS30 returns in Period 1 and Period 2.*

USD/EUR		
	Period 1	Period 2
P-value	0.842	0.784
Q-statistic	2.049	2.452
Critical Value	11.071	11.071

Table 8.6: *P-values, Q-Statistic and Critical values from the Ljung-Box test on the residuals from the USD/EUR returns in Period 1 and Period 2.*

OMXS30	
	Period 2
P-value	0.123
Q-statistic	8.682
Critical Value	11.071

Table 8.7: *P-values, Q-Statistic and Critical values from the Ljung-Box test on the residuals from the OMXS30 returns in Period 2. Returns are described by an AR(1) process.*

8.3 Engle's ARCH Test

Guld		
	Period 1	Period 2
P-value	3.561E-04	1.483E-04
F-statistic	15.881	17.633
Critical Value	5.992	5.992

Table 8.8: *P-values, F-Statistics and Critical values from the Engle's ARCH test on the squared residuals from the gold returns in Period 1 and Period 2.*

OMXS30		
	Period 1	Period 2
P-value	0.000	1.955E-08
F-statistic	132.484	35.501
Critical Value	5.992	5.992

Table 8.9: *P-values, F-Statistics and Critical values from the Engle's ARCH test on the squared residuals from the OMXS30 returns in Period 1 and Period 2.*

USD/EUR		
	Period 1	Period 2
P-value	2.160E-05	0.020
F-statistic	21.178	7.781
Critical Value	5.992	5.992

Table 8.10: *P-values, F-Statistics and Critical values from the Engle's ARCH test on the squared residuals from the EUR/USD returns in Period 1 and Period 2.*

8.4 Jarque-Bera Test

Gold		
	Period 1	Period 2
P-value	1.000E-03	1.000E-03
JB-statistic	276,7	5.650E+03
Critical Value	5.931	5.931

Table 8.11: *P-values, F-Statistics and Critical values from the Jarque-Bera test on the residuals from the gold returns in Period 1 and Period 2.*

OMXS30		
	Period 1	Period 2
P-value	1.000E-03	1.000E-03
JB-statistic	1.813E+03	134.0
Critical Value	5.931	5.931

Table 8.12: *P-values, F-Statistics and Critical values from the Jarque-Bera test on the residuals from the OMXS30 returns in Period 1 and Period 2.*

USD/EUR		
	Period 1	Period 2
P-value	1.000E-03	1.000E-03
JB-statistic	212.4	146.4
Critical Value	5.931	5.931

Table 8.13: *P-values, F-Statistics and Critical values from the Jarque-Bera test on the residuals from the EUR/USD returns in Period 1 and Period 2.*

8.5 Autocorrelation Function

Figure 8.1: *The two plots show the sample autocorrelation for the in-sample residuals from the return of lags 0-20. The upper plot corresponds to Period 1 whereas the lower plot corresponds to Period 2. The blue line denotes the confidence interval at five percent.*

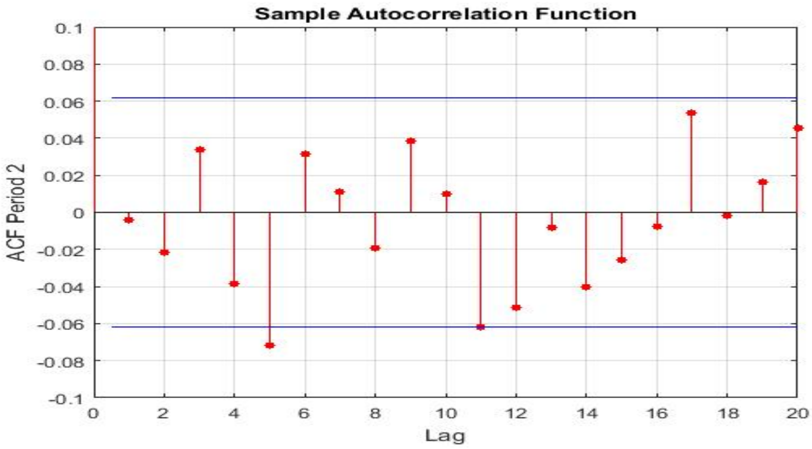

8.6 Mean Equation

OMXS30 (Period 2)				
Variable	Coefficient	Std. Error	t-Statistic	P-value
Constant	0.000384	0.000319	1.202733	0.2294
r_{t-1}	-0.079017	0.030879	-2.558910	0.0106

Table 8.14: *Estimated parameters for the AR(1) process for the OMXS30 returns.*

8.7 P-values

P-values for ARCH(1) model

	Period 1		Period 2	
	α_0	α_1	α_0	α_1
Gold (n-dist)	0.0000	0.0045	0.0000	0.0000
Gold (t-dist)	0.0059	0.0365	0.0000	0.3531
OMXS30 (n-dist)	0.0000	0.0000	0.0000	0.0001
OMXS30 (t-dist)	0.0000	0.0000	0.0000	0.0210
USD/EUR (n-dist)	0.0000	0.0000	0.0000	0.0075
USD/EUR (t-dist)	0.0000	0.0068	0.0000	0.0321

Table 8.15: *P-values for the estimated parameters from an ARCH(1) model for gold, OMXS30 and USD/EUR in Period 1 and Period 2.*

P-values for GARCH(1,1) model

	Period 1			Period 2		
	α_0	α_1	β_1	α_0	α_1	β_1
Gold (n-dist)	0.0007	0.0000	0.0000	0.0000	0.0000	0.0000
Gold (t-dist)	0.0635	0.0001	0.0000	0.1418	0.0257	0.0000
OMXS30 (n-dist)	0.0076	0.0000	0.0000	0.0141	0.0000	0.0000
OMXS30 (t-dist)	0.0159	0.0000	0.0000	0.0658	0.0002	0.0000
USD/EUR (n-dist)	0.0005	0.0000	0.0000	0.0103	0.0000	0.0000
USD/EUR (t-dist)	0.0743	0.0002	0.0000	0.4287	0.0011	0.0000

Table 8.16: *P-values for the estimated parameters from a GARCH(1,1) model for gold, OMXS30 and USD/EUR in Period 1 and Period 2.*

P-values for IGARCH(1,1) model

	Period 1		Period 2	
	α_0	β_1	α_0	β_1
Gold (n-dist)	0.0000	0.0000	0.0000	0.0000
Gold (t-dist)	0.0000	0.0000	0.0000	0.0000
OMXS30 (n-dist)	0.0000	0.0000	0.0000	0.0000
OMXS30 (t-dist)	0.0000	0.0000	0.0000	0.0000
USD/EUR (n-dist)	0.0000	0.0000	0.0000	0.0000
USD/EUR (t-dist)	0.0000	0.0000	0.0000	0.0000

Table 8.17: *P-values for the estimated parameters from an IGARCH(1,1)model for gold, OMXS30 and USD/EUR in Period 1 and Period 2.*

P-values for EGARCH(1,1) model

	Period 1				Period 2			
	α_0	α_1	γ_1	β_1	α_0	α_1	γ_1	β_1
Gold (n-dist)	0.0000	0.5825	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
Gold (t-dist)	0.0003	0.7109	0.0000	0.0000	0.0181	0.0026	0.0980	0.5540
OMXS30 (n-dist)	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000
OMXS30 (t-dist)	0.0001	0.0001	0.0000	0.0000	0.0000	0.0000	0.0005	0.0000
USD/EUR (n-dist)	0.0001	0.4088	0.0000	0.0000	0.1614	0.0000	0.1833	0.0000
USD/EUR (t-dist)	0.0096	0.7005	0.0000	0.0000	0.0000	0.0095	0.0115	0.0043

Table 8.18: *P-values for the estimated parameters from an EGARCH(1,1) for gold, OMXS30 and USD/EUR in Period 1 and Period 2.*

P-values for GJR-GARCH(1,1) model

	Period 1				Period 2			
	α_0	α_1	γ_1	β_1	α_0	α_1	γ_1	β_1
Gold (n-dist)	0.0000	0.0002	0.3378	0.0000	0.1940	0.7538	0.0173	0.0000
Gold (t-dist)	0.1709	0.0042	0.3260	0.0000	0.0439	0.5741	0.1624	0.0000
OMXS30 (n-dist)	0.0011	0.0863	0.0000	0.0000	0.0022	0.6726	0.0000	0.0000
OMXS30 (t-dist)	0.0016	0.1392	0.0001	0.0000	0.0074	0.5519	0.0000	0.0000
USD/EUR (n-dist)	0.0514	0.0003	0.6749	0.0000	1.0000	0.6371	0.0000	0.0000
USD/EUR (t-dist)	0.1198	0.0000	0.8119	0.0000	1.0000	0.4529	0.0005	0.0000

Table 8.19: *P-values for the estimated parameters from a GJR-GARCH(1,1) for gold, OMXS30 and USD/EUR in Period 1 and Period 2.*