LUND UNIVERSITY
School of Economics and Management

Inflation and Inflation Uncertainty in Sweden: a GARCH Modelling Approach

Anna Åkesson
June 3, 2016

Bachelor essay in Statistics (15 ECTS)
Lund University · Department of Statistics

Supervisor: Krzysztof Podgórski

Abstract

This essay investigates inflation and inflation uncertainty in Sweden from 1970:Q1 to 2014:Q4. GARCH models are used to generate a measure of inflation uncertainty estimated under the distributional assumption of Student's t-distribution and GED. The preferable model found for Swedish inflation was a EGARCH(1,1) estimated with Student's t-distribution. The coefficient of the asymmetry in the conditional variance of the Swedish inflation series was found to be negative, thereby implying that a positive shock to the Swedish inflation leads to less uncertainty about inflation. Relating these results to the two competing hypotheses about the causal relationship between inflation and its uncertainty, the results in this essay does not support the Friedman-Ball hypothesis. With regards to the Cukierman-Meltzer hypothesis, the results seem to support a version of this hypothesis put forward by Holland (1995), referred to as the "stabilisation hypothesis".

Key words: Inflation; Inflation uncertainty; GARCH models; leverage effects.

Contents

List of Tables

List of Figures

1 Introduction

Inflation brings about higher costs for both the economy and for society as a whole, one of the most important cost of inflation is that of uncertainty of future inflation rates, i.e inflation uncertainty. As an example one can look upon inflation uncertainty as a cause for uncertainty about contractual payments in the future. For instance, due to inflation uncertainty both employers and employees will be insecure about the real future value of wages. Thus uncertainty about future interest rates and other economic variables may have a negative effect on the economic activity.

From the above reasoning two competing hypotheses on the causal relationship between inflation and inflation uncertainty have emerged. The first hypothesis postulates that inflation causes inflation uncertainty, this was first suggested by Friedman (1977). Ball (1992) proposed a model in which higher inflation leads to increasing uncertainty over the monetary policy stance. The possibility of a negative effect of inflation on inflation uncertainty have been discussed by Pourgerami and Maskus (1987), who pointed out that in a world of accelerating inflation agents (e.g. firms and households) may invest more resources in inflation forecasting, and thereby reducing uncertainty (see also Ungar and Zilberfarb, 1993). This first hypothesis is usually denoted the Friedman-Ball hypothesis.

The second hypothesis on the causal relation between inflation and inflation uncertainty conversely postulates that inflation uncertainty causes inflation. This hypothesis has been examined by Cukierman and Meltzer (1986) who claim that when inflation uncertainty increases, it causes high rates of inflation. The authors proved that increased uncertainty about money supply and inflation raises the optimal inflation rate set by policymakers. That is, in the world of Cukierman and Meltzer (1986) the policymakers exhibits an opportunistic behaviour, meaning that they generate surprise inflation for economic agents, in order to obtain output gains. The same causality, but with a negative relationship between variables, was proposed by Holland (1995) suggesting that in the case of increased inflation uncertainty, the policymaker has a stabilising behaviour, meaning that they reduce money supply growth to reduce the negative welfare effects, this is sometimes denoted the "stabilisation hypothesis". It has been suggested that the stabilising behaviour of the monetary authorities is related to the level of central bank independence, in particular the higher the level of central bank independence, the lesser the rate of inflation (Grier and Perry 1998). The second hypothesis is usually referred to as the Cukierman-Meltzer hypothesis.

Based on the argument presented above the purpose of this essay is to investigate inflation and its uncertainty in Sweden using a GARCH-modelling approach. In particular, the primary purpose is to find an appropriate model for the inflation series in a GARCH modelling setting allowing for non-normal errors. The secondary purpose is to relate the results to the two competing hypotheses of the causal relationship between inflation and inflation uncertainty, the Friedman-Ball and the Cukierman-Meltzer hypotheses, and thus see in which camp the Swedish inflation series falls. The limitation made in the empirical study is to solely consider one country, which was in part made due to data availability and in part due to a desire to delve a bit deeper into the distributional assumption of the errors without overflowing the reader with results and tables and it was therefore preferable to focus the study to one country.

This essay contributes to the empirical literature on inflation and its uncertainty by allowing for non-normal error, which is in contrast to the existing literature that assumes normal errors for the GARCH model, and the empirical study in this essay did indeed find that the residuals does not follow a normal distribution but instead more appropriate distributional assumptions seemed to be the Student's t-distribution and the Generalised error distribution (GED). Under these two distributional assumptions three models emerge as contenders as the model of Swedish inflation, namely the EGARCH(1,1) and TGARCH(1,1) that is models that take asymmetry into account. After evaluating the assumptions underlying the model and on the basis of forecast evaluation statistics, the preferable model for Swedish inflation and its uncertainty that emerged was a EGARCH(1,1) estimated with Student's t-distribution.

The coefficient of the asymmetry in the conditional variance of the Swedish inflation series was found to be negative, thus implying that a positive shock to the Swedish inflation leads to less uncertainty about inflation. Relating these results to the two competing hypotheses, the results in this essay does not support the Friedman-Ball hypothesis. With regards to the Cukierman-Meltzer hypothesis, the results seem to support a version of this hypothesis put forward by Holland (1995).

The rest of this essay is structured as follows. In Section 2 presents a literature review of previous studies on inflation and inflation uncertainty. Section 3 presents the data and methodology used in this essay. In Section 4 the empirical study is presented and results are discussed. Finally, Section 5 concludes.

2 Literature Review

The first empirical papers that investigated inflation and inflation uncertainty assumed that differences in the standard deviation of inflation across countries were a valid estimate of the differences in inflation uncertainty across countries. When the literature turned to time-series tests, two uncertainty measures were typically used to estimate the uncertainty, namely the cross-sectional dispersion of individual forecasts from surveys or a moving standard deviation of the variable under consideration (Holland 1993, Golob 1993).

However, none of the techniques described above will be able to capture the uncertainty structure derived in the theoretical models of Ball (1992) and Cukierman and Meltzer (1986). This is because in these models the uncertainty is the variance of a stochastic process. A method that can incorporate this type of uncertainty is ARCH/GARCH-type of models which is why most empirical studies conducted on inflation and inflation uncertainty use these types of models.

Engle (1983) and Bollerslev (1986) were the first authors to employ ARCH/GARCH models to estimate the inflation rate for the UK and US, respectively. They then compare a graph of the conditional variance of inflation to the average inflation rate. Inflation uncertainty is highest in the late 1940s and early 1950s when inflation is not particularly high and inflation uncertainty is lower in the late 1970s and early 1980s when inflation is quite high. Given this, both authors independently conclude that high inflation levels are not correlated with unpredictable inflation. Such a conclusion is unsatisfactory in that it is not based on any kind of statistical test.

Baillie et al. (1996) tested the direction of causality between inflation and inflation uncertainty by including lagged inflation in the conditional variance equation and standard deviation in inflation equation. The authors demonstrated that there was an interrelation between inflation and inflation uncertainty in the US and some other countries with high inflation rates (Argentina, Brazil and Israel).

Caporale and McKiernan (1997) employed a GARCH model to monthly US inflation data. The authors established that there was a significant positive connection between inflation rate and conditional variance of inflation, that is a positive connection between inflation and its uncertainty. The authors further argued that the obtained results were also robust for alternative inflation models, and unlikely to depend on whether the data of US high inflation observed in the 1970s were included in or excluded from estimations.

Grier and Perry (1998) used the estimated conditional variance from a GARCH model and Granger causality test to determine the causality direction of inflation and inflation uncertainty. The authors found that inflation had a positive substantial impact on inflation uncertainty (hypothesis by Friedman–Ball) in all G7 countries; this impact, however, was not unambiguous. In three countries, including the US, increased inflation uncertainty lowered inflation, which contradicts the Cukierman–Meltzer theory. At the same time, with inflation uncertainty growing, inflation rose in Japan and France.

Studies that take leverage effects into account are for example Nazar et al. (2010) which use a time series of inflation uncertainty in Iran from 1959-2009 to estimate an EGARCH model. The authors confirm that positive shocks to inflation had a greater effect on inflation uncertainty as compared with negative effects. Granger causality test also verified that inflation Granger-caused inflation uncertainty.

Fountas et al. (2004) also employed EGARCH models to estimate inflation uncertainty and the authors find strong evidence for the Friedman-Ball hypothesis and mixed evidence for the second hypothesis (the Cukierman-Meltzer hypothesis) for a sample of six EU countries.

Using a TARCH model for the US inflation rate Caporale and Caporale (2002) find that negative inflationary shocks result in greater inflation uncertainty than positive shocks. That is, the authors find support for the Friedman-Ball hypothesis that inflation uncertainty leads to lower output growth.

Karanasos and Schurer (2008) modelled the conditional variance of inflation in Germany, the Netherlands and Sweden using APARCH models for the period 1962-2004 using monthly data. The authors find that for Germany and the Netherlands there is support for the Cukierman-Meltzer hypothesis. In Sweden, on the other hand, they find that uncertainty about the future inflation appears to have a negative impact on inflation.

3 Data and Methodology

3.1 Inflation

Before any further discussions are made a reasonable starting point is to begin with a brief discussion of the main variables of interest; inflation and inflation uncertainty. Inflation is usually defined as a continued increase in the general price level of goods and services in the economy over a period of time (month, quarter or year). As the price level rises each unit of currency buys fewer goods and services, therefore inflation reflects a reduction in the purchasing power per unit of money (Walgenbach et al. 1973). The measure for (price) inflation is the inflation rate which is defined as the percentage rate of change of a price index over time. Using the conventional notation for inflation rate from economic theory, π, we can write the inflation rate as:

$$\pi_t = \log(P_t) - \log(P_{t-1}) \tag{1}$$

where P is a price index. The most commonly used price index in the calculation of the inflation rate is the consumer price index (CPI) (Blanchard 2000). What CPI measures is changes over time in the general level of prices of goods and services that a reference population obtains, uses or pays for its consumption (OECD 2013). However there are other commonly used indices for the calculation of the inflation rate, such as producer price index (PPI), commodity price index, and core price index.

PPI measures the average changes in prices received by domestic producers for their output. The difference from the CPI is that price subsidisation, profits, and taxes may cause the amount the producer receives to differ from what is paid by the consumer (OECD 2005). Generally, there is also a delay between an increase in the PPI and any subsequent increase in the CPI. Another way to define PPI is as a measure of the pressure being put on producers by the costs of their raw materials, this pressure can end up being passed on to consumers (i.e. increased prices paid by the consumer), or it could be absorbed by profits, or possibly offset by increasing the productivity of the producers.

Commodity price index is a fixed-weight index that measures the price of a selection of commodities. A commodity price index can be defined for all commodities or for particular subsets such as energy. Lastly, there is core price indices, in which the most volatile components of CPI are removed in order to easier detect the long run trend in the price levels (Gordon 1975). The most volatile parts of CPI are usually considered to be food and oil prices, since these two prices can change quickly due to changes in supply and demand conditions in food and oil markets.

Inflation uncertainty is uncertainty of future inflation rates. This uncertainty can lead to a reduction in economic output since it can adversely impact interest rates, consumer spending and saving as well as investment made by firms. Since inflation is seen as a stochastic process (when modelled over time) then inflation uncertainty is defined as the variance of this stochastic process. The variance is not expected to be constant over time, that is volatility clusters are expected to be present.

3.2 Stationarity

In order to be able to estimate any kind of time series model the stationarity of the series must be evaluated. If the stationarity hypothesis is not confirmed for the series, then the series must be transformed into a stationary series through one of the known traditional procedures, such as the first order difference of the series. Whether or not the inflation series is stationarity can be tested using the Augmented Dickey-Fuller (ADF) test, the Phillips-Perron test and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test.

The Dickey-Fuller test is performed for first order autoregressive process and is based on the following general equation (Enders 2010):

$$\pi_t = \phi\pi_{t-1} + \varepsilon_t \tag{2}$$

$$\Delta\pi_t = \alpha_0 + \beta t + \gamma\pi_{t-1} + \varepsilon_t \tag{3}$$

$$\text{where } \Delta\pi_t = \pi_t - \pi_{t-1} \text{ and } \gamma = \phi - 1$$

Equation (3) can be used to test whether or not a unit root is present, using the following hypotheses:

$$H_0 : \gamma = 0$$
$$H_1 : \gamma \neq 0$$

The procedure for an Augmented Dickey-Fuller (ADF) test is similar as for a Dickey-Fuller test but is performed for the following model

$$\Delta\pi_t = \alpha + \beta t + \gamma\pi_{t-1} + \delta_1\Delta\pi_{t-1} + \ldots + \delta_{p-1}\Delta\pi_{t-p+1} + \varepsilon_t \tag{4}$$

where α is a constant, β is a time trend coefficient and p is the number of lags in the autoregressive process. The presence of a unit root is tested using the same hypotheses as above and with the following test statistic

$$\tau = \frac{\hat{\gamma}}{SE(\hat{\gamma})} \tag{5}$$

If the test statistic is less than the (more negative) Dickey-Fuller critical value, then the null hypothesis of $\gamma = 0$ is rejected and no unit root is present.

In a similar vein as the ADF test, the Philips-Perron (PP) test addresses the issue that the process generating data for π_t might have a higher order of autocorrelation than is admitted in the test equation and thus making π_{t-1} endogenous and invalidating the Dickey–Fuller τ-test. Whilst the ADF test addresses this issue by introducing lags of $\Delta\pi_t$ as regressors in the test equation, the Phillips–Perron test makes a non-parametric correction to the τ-test statistic. The test is robust with respect to unspecified autocorrelation and heteroskedasticity in the error process of the test equation (Phillips and Perron 1988). However, as Davidson and MacKinnon (2004) report the PP test performs worse in finite samples than the augmented Dickey–Fuller test.

The third and final test for stationarity that is presented is the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test. This test analyses the properties of the residuals from the regression equation of π_t depending on the previous values. In the KPSS test the null hypothesis is stationarity, i.e. the opposite of the other two tests described above (Ruppert 2010).

3.3 Distributions for Financial Time Series

Time series of financial data are in general not consider to be normally distributed (Ruppert 2010), and since inflation is a financial time series it can therefore possibly not either be considered to be normally distributed. This clearly has consequences for estimation of model parameters, since if the wrong distributional assumption is made one will not obtain efficient estimators and less accurate predictions (Zivot 2009). Instead financial time series are usually better characterised by heavy-tailed distributions. This section will discuss a two commonly used heavy-tailed distributions used for financial time series, namely Student's t-distribution and the generalised error distribution (GED).

The Student's t distribution, denoted t_ν, is used in order to better describe the conditional heteroskedasticity (Bollerslev 1987). Student's t has the following density (Ruppert 2010):

$$f_\nu(y) = \left[\frac{\Gamma[(\nu+1)/2]}{(\pi\nu)^{1/2}\Gamma[\nu/2]}\right] \cdot \frac{1}{[1+(y^2/\nu)]^{(\nu+1)/2}} \tag{6}$$

where $\Gamma(\cdot)$ is the gamma function defined by:

$$\Gamma(u) = \int_0^\infty x^{u-1}e^{-x}dx,\ u > 0$$

If $\nu > 2$ then the variance of t_ν is finite and equals $\nu/(\nu-2)$. On the other hand if $0 < \nu \leq 1$ then the expected value of t_ν does not exist and its variance is not defined. If $1 < \nu \leq 2$ then t_ν has expected value equal to 0 and infinite variance. It should also be noted that t_ν is a symmetric distribution, i.e has skewness equal to 0, and has the desired heavy tails since it has kurtosis equal to $3 + \frac{\nu}{\nu-4}$, for $\nu > 4$, which is larger than three (kurtosis of a normal distribution).

Another heavy-tailed distribution commonly used for financial time series is the generalised error distribution (GED), which belongs to a family of symmetric exponential distributions (Nelson 1991). These type of distributions have the particular property, that depending on the value of a shape parameter ν, they can be described as both leptokurtic and platykurtic. Where a leptokurtic distribution is one where extreme values have a relatively high probability of occurring as compared to a normal distribution. This results in density with heavy tails and a spiky peak. The opposite is a platykurtic distribution, whose density has a flatter peak and wider middle part. The standardised generalised error distribution (GED) has the following density (Ruppert 2010):

$$f^{std}_{ged}(y|\nu) = \kappa(\nu)e^{-\frac{1}{2}|\frac{y}{\lambda_\nu}|^\nu}, \; -\infty < y < \infty \tag{7}$$

$$\text{where } \lambda_\nu = \left(\frac{2^{-2/\nu}\Gamma(\nu^{-1})}{\Gamma(3/\nu)}\right)^{1/2}$$

$$\text{and } \kappa(\nu) = \frac{\nu}{\lambda_\nu 2^{1+(1/\nu)}\Gamma(\nu^{-1})}$$

The constants $\kappa(\nu)$ and λ_ν are chosen so that the function integrates to 1 [1] and so the variance equals 1. The weight of the tails is determined by the shape parameter $\nu > 0$, where smaller values of ν will generate greater weight to the tails. Note that when $\nu = 2$ the GED is a normal distribution and when $\nu = 1$ GED is a double-exponential distribution. Thus the generalised error distribution can produce tail weights that are between those of a normal and double-exponential distribution if $1 < \nu < 2$. However, more extreme tails weights than that of a double-exponential is also possible if $\nu < 1$.

3.4 Modelling Conditional Heteroskedasticity

Conditional heteroskedasticity was first modelled by Engle (1983) and Bollerslev (1986) using just inflation data for the UK and US , respectively. This section will discuss the ARCH and GARCH model proposed by the authors above.

The Autoregressive Conditional Heteroskedasticity (ARCH) model was introduce by Engle (1982) and is of the following form

$$\pi_t = E_{t-1}\big[\pi_t\big] + \varepsilon_t \tag{8}$$

$$\varepsilon_t = z_t\sigma_t \tag{9}$$

$$\sigma^2_t = \alpha_0 + \alpha_1\varepsilon^2_{t-1} + \cdots + \varepsilon^2_{t-p} \tag{10}$$

where $E_{t-1}[\cdot]$ is the notation for the conditional expectation for information known up to time $t-1$. The sequence z_t is a sequence of independent

[1] Otherwise it would not be a density

identically distributed (i.i.d) variables with zero mean and unit variance. In order to have $\sigma_t^2 > 0$ some restrictions has to be imposed on the coefficients, namely $\alpha_0 > 0$ and $\alpha_i \geq 0$ for $i = 1, \ldots, p$. Note that equation (10) for σ_t^2 can be written as an AR(p)-process for ε_t^2 in the following way

$$\varepsilon_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \cdots + \alpha_p \varepsilon_{t-p}^2 + u_t \tag{11}$$

where $u_t = \varepsilon_t^2 - \sigma_t^2$. Furthermore note that u_t will be a so-called martingale difference sequence since it will have a finite expectation and its expectation with respect to the past will be zero, i.e $E\big[\varepsilon_t^2\big] < \infty$ and $E_{t-1}\big[u_t\big] = 0$ (Zivot 2009). In order for ε_t to be covariance stationary we must require $\alpha_1 + \cdots + \alpha_p < 1$, this sum of coefficients is also used to measure the persistence of ε_t^2 and σ_t^2, and the unconditional variance will be

$$\bar{\sigma}^2 = Var\big[\varepsilon_t\big] = E\big[\varepsilon_t^2\big] = \frac{\alpha_0}{1 - \alpha_1 - \cdots - \alpha_p}$$

One drawback of the ARCH model is that in order to fully capture the dynamic of the conditional heteroskedasticity one needs a model of high order. If the model is to be used for forecasting for a longer time period then this can lead to over-parametrisation (Enders 2010). That led to the generalisation of the ARCH model.

The generalised ARCH model, or GARCH, was introduced by Bollershev (1986), this model generalises the AR(p)-process in equation (10) with an ARMA(p,q)-process instead. That is σ_t^2 can be represented by:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2 \tag{12}$$

$$\text{where } \alpha_i > 0 \; \forall \; i \text{ and } \beta_j > 0 \; \forall \; j$$

The coefficients are assumed to be positive in order to guarantee that the conditional variance, σ_t^2, is always positive. For the most part a GARCH(1,1) with three parameters in the conditional variance equation is an adequate representation for time series of financial data. Telling evidence by Hansen and Lunde (2004) says that it is hard to find a model that performs better than a GARCH(1,1).

A GARCH model can be expressed as an ARMA-process for the squared residuals, in a similar vein as the ARCH model was written an AR-process. For simplicity, consider a GARCH(1,1) of the following form

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \tag{13}$$

Next utilise that $E_{t-1}\left[\varepsilon_t^2\right] = \sigma_t^2$ and then (13) can be rewritten as

$$\varepsilon_t^2 = \alpha_0 + (\alpha_1 + \beta_1)\varepsilon_{t-1}^2 + u_t - \beta_1 u_{t-1} \tag{14}$$

which is an ARMA(1,1) where $u_1 = \varepsilon_1^2$. This means that the properties of a GARCH(1,1) can be derived using the ARMA(1,1) representation. The process for the conditional variance will be covariance stationary if $\alpha_1 + \beta_1 < 1$. The persistence of σ_t^2 is also measured by the sum of the coefficients. Furthermore the covariance stationary GARCH(1,1) has unconditional variance equal to

$$\bar{\sigma}^2 = \frac{\alpha_0}{1 - \alpha_1 - \beta_1}$$

If instead a general GARCH(p,q) process, like the one in equation (12), is consider then the squared residuals will behave like an ARMA(max(p,q),q) process. For the general GARCH covariance stationarity requires that $\sum_{i=1}^{p} \alpha_i + \sum_{j=1}^{q} \beta_j < 1$ and have the following unconditional variance

$$\bar{\sigma}^2 = \frac{\alpha_0}{1 - \left(\sum_{i=1}^{p} \alpha_i + \sum_{j=1}^{q} \beta_j \right)}$$

The order of the GARCH model is usually determined by using a information criteria, such as the Akaike information criterion (AIC), the Bayesian information criterion (BIC) and the Hannan-Quinn information criterion. In particular, for GARCH models with $p, q \leq 2$ the order is usually determined by the means of a information criteria (Zivot 2009). The possibility to use information criterion comes from the fact that GARCH models can be expressed as an ARMA-process as was mentioned above.

3.5 Asymmetric Leverage Effects

In financial time series leverage effects are often present, i.e. the fact that shocks (or news) have asymmetric effects on the volatility. In particular, negative shocks tend to have a greater impact on volatility than positive shocks. The GARCH(p,q) model described above cannot capture these type of leverage effects (Ruppert 2010), and this section will discuss three extensions of the GARCH model that can.

The first model that can be used to incorporate leverage effects is the Exponential GARCH (EGARCH) model proposed by Nelson (1991). Following Nazar et al. (2010) and Fountas et al. (2004) the conditional variance in an EGARCH model has the following expression:

$$h_t = \log\left(\sigma_t^2\right) = \alpha_0 + \sum_{i=1}^{p} \alpha_i \frac{|\varepsilon_{t-i}| + \gamma_i \varepsilon_{t-i}}{\sigma_{t-i}} + \sum_{j=1}^{q} \beta_j h_{t-j} \tag{15}$$

If ε_{t-i} is positive then the total effect of ε_{t-i} will be $(1+\gamma_i)|\varepsilon_{t-i}|$, on the other hand if ε_{t-i} is negative then the total effect of ε_{t-i} will be $(1-\gamma_i)|\varepsilon_{t-i}|$. Note that in this specification negative shocks can have a greater impact on volatility and one would then expect γ_i to be negative. The conditional variance, σ_t^2, will be positive regardless of the value on the coefficients in (15), this because the logarithm of σ_t^2 is modelled instead of σ_t^2 itself. This is an advantage of the EGARCH model over the GARCH model. The EGARCH model is covariance stationary if $\sum_{j=1}^{q} \beta_j < 1$.

The second model to incorporate leverage effect is the threshold GARCH (TGARCH) proposed by Zakoian (1994) and Glosten et al. (1993). In a similar vein as in Caporale and Caporale (2002) the conditional variance has the following expression:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^{p} \gamma_i S_{t-i} \varepsilon_{t-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^2 \tag{16}$$

where

$$S_i = \begin{cases} 1 & \text{if } \varepsilon_{t-i} < 0 \\ 0 & \text{if } \varepsilon_{t-i} \geq 0 \end{cases}$$

The effect of ε_{t-i}^2 on the conditional variance will depend on whether or not ε_{t-i} is above or below the threshold. If ε_{t-i} is positive then the effect of ε_{t-i} will be $\alpha_i \varepsilon_{t-i}^2$, and if ε_{t-i} is negative then the effect of ε_{t-i} will be $(\alpha_i + \gamma_i)\varepsilon_{t-i}^2$. These shocks have an asymmetric effect on inflation uncertainty if the γ_is are statistically different from zero. Consider a TGARCH(1,1) model, this model is covariance stationary if

$$\gamma_t^2 < 1 - \alpha_1^2 - \beta_1^2 - 2\alpha_1\beta_1\nu_1 \tag{17}$$

where ν_1 is $\sqrt{\frac{\pi}{2}}$ if ε_t follows a normal distribution and $\sqrt{\frac{(\nu-2)}{\pi}}\frac{\Gamma((\nu-1)/2)}{\Gamma(\nu/2)}$ if ε_t follows a Student's t-distribution with ν degrees of freedom (He and Teräsvirta 1999).

The third model that incorporates leverage effects is the Asymmetric Power ARCH (APARCH) model proposed by Ding et al. (1993). Following Karanasos and Schurer (2008) the conditional standard deviation has the following expression:

$$\sigma_t^\delta = \alpha_0 + \sum_{i=1}^{p} \alpha_i \big(|A_{t-i}| - \gamma_i A_{t-i}\big)^\delta + \sum_{j=1}^{q} \beta_j \sigma_{t-j}^\delta \tag{18}$$

where

$$A_t = \sigma_t \sigma_t$$

where $\delta > 0$ and $-1 < \gamma_i < 1, i = 1, \ldots, p$. Note that when $\delta = 2$ and $\gamma_1 = \ldots = \gamma_p = 0$ the standard GARCH(p,q) model is obtained.

3.6 Testing for ARCH/GARCH Effects

If the data has volatility clusters, or that is ARCH/GARCH effects are present, it will manifest itself as autocorrelation in the squared time series. A test for the significance of the autocorrelation is the modified Ljung-Box test, what is usually referred to as the McLeod-Li test which has the following test statistic (Zivot 2009)

$$MQ(p) = T(T+2)\sum_{j=1}^{p} \frac{\hat{\rho}_j^2}{T-j} \tag{19}$$

where $\hat{\rho}_j$ is the jth lag of the sample autocorrelation. Under the null the time series is a white noise process, so the MQ(p)-statistic asymptotically follows a $\chi^2(p)$-distribution. If a significant value of MQ(p)-statistic is found, then the test provides evidence of time varying conditional variance.

Engle (1982) showed that if an ARCH model is employed with an AR-process for the squared residuals, then a Lagrange Multiplier (LM) test can be used to test for ARCH effects. The LM test for ARCH effects is based on equation (10). The null hypothesis is that there are no ARCH effect, i.e. $\alpha_1 = \cdots = \alpha_p = 0$, test statistic is the following

$$LM = TR^2 \tag{20}$$

where T is the sample size and R^2 is computed from the regression of (10) based on estimated residuals. Under the null the test asymptotically follows a $\chi^2(p)$-distribution. Lee and King (1993) showed that this LM test also can be used as a test for GARCH effects.

3.7 Parameter Estimation

The most common method used for estimation of the parameters in a GARCH model is maximum likelihood (ML) estimation. Before proceeding with ML estimation of the parameters the distributional assumption must be considered. As an example, consider a financial time series that cannot be considered to be characterised by a normal distribution, then if the parameters are estimated under the assumption of normal distribution they may be less effective. If the distributional assumption is changed to the Student's t-distribution, then the following conditional log-likelihood function is obtained

$$\begin{aligned}\log(L) = \sum_{t=1}^{T} \Bigg[\log\left(\Gamma\left(\frac{\nu+1}{2}\right)\right) - \frac{1}{2}\log(\pi\nu) \\ - \log\left(\Gamma\left(\frac{\nu+1}{2}\right)\right) - \frac{\nu+1}{2}\log\left(\frac{\varepsilon_t^2}{\nu}\right)\Bigg]\end{aligned} \tag{21}$$

For the case when the distributional assumption is GED, the following conditional log-likelihood function is obtained

$$\log(L) = \sum_{t=1}^{T} \log\big(\kappa(\nu)\big) - \frac{1}{2}\Big|\frac{\varepsilon_t}{\lambda_\nu}\Big| \tag{22}$$

3.8 Evaluation of the Estimated GARCH Model

If the estimated GARCH model is correctly specified then the behaviour of the standardised residuals should be the same as in classical linear regression. That is, the standardised residuals should exhibit neither serial correlation, conditional heteroskedasticity nor any non-linear dependence. Moreover, the distribution of the standardised residuals should reflect the distributional assumption made in the estimation of the model (Zivot 2009).

The serial correlation, or ARCH/GARCH effects, of the standardised residuals can be evaluated by the same test that are used to test for ARCH/GARCH effects (the McLeod test and the LM test), and if the model is correctly specified then no such effects should be found. The distributional assumption for the standardised residuals can be evaluated by using QQ-plots.

3.9 Forecasting Using the Estimated GARCH Model

The estimated GARCH model can be used for forecasting both the future value of the time series itself and the conditional volatility. In order to make forecast of π_T from a GARCH model, with the assumption that $E_T[\pi_{T+1}] = c$. Then for this model, the h-step-ahead forecast of π_T will be c, which clearly does not depend on the parameters in the GARCH model, this forecast will have the following forecast error (Enders 2010)

$$\varepsilon_{T+h} = \pi_{T+h} - E_T[\pi_{T+h}] \tag{23}$$

The forecast error in (23) will have the following conditional variance

$$Var_T[\varepsilon_{T+h}] = E_T[\sigma^2_{T+h}] \tag{24}$$

Note that the conditional variance in (24) does not depend on the parameters from the GARCH model. So the h-step-ahead volatility forecast is needed in order to provide confidence bands for the h-step-ahead forecast.

Computing $E_T[\sigma^2_{T+h}]$ can be done using recursion. Consider a GARCH(1,1) model with $\varepsilon_t = z_t\sigma_t$ and z_t are i.i.d. with zero mean and unit variance from a symmetric distribution. Then for $h = 1$

$$\begin{aligned} E_T[\sigma^2_{T+1}] &= \alpha_0 + \alpha_1 E_T[\varepsilon^2_T] + \beta_1 E_T[\sigma^2_T] \\ &= \alpha_0 + \alpha_1\varepsilon^2_T + \beta_1\sigma^2_T \end{aligned} \tag{25}$$

Note that in (25) it is assumed that ε^2_T and σ^2_T are known. However, in practice none of the true parameters are known and instead the estimated parameters from the model is used (Zivot 2009). In the case when $h = 2$

$$\begin{aligned} E_T[\sigma^2_{T+2}] &= \alpha_0 + \alpha_1 E_T[\varepsilon^2_{T+1}] + \beta_1 E_T[\sigma^2_{T+1}] \\ &= \alpha_0 + (\alpha_1 + \beta_1) E_T[\sigma^2_{T+1}] \end{aligned} \tag{26}$$

since

$$E_T[\varepsilon^2_{T+1}] = E_T[\sigma^2_{T+1} z^2_{T+1}] = E_T[\sigma^2_{T+1}]$$

Thus, h-step-ahead volatility forecast for $h \geq 2$ will be

$$\begin{aligned} E_T[\sigma^2_{T+h}] &= \alpha_0 + (\alpha_1 + \beta_1) E_T[\sigma^2_{T+h-1}] \\ &= \alpha_0 \sum_{i=0}^{h-1} (\alpha_1 + \beta_1)^i + (\alpha_1 + \beta_1)^{h-1} (\alpha_1\varepsilon^2_T + \beta_1\sigma^2_T) \end{aligned} \tag{27}$$

The forecast obtained from (27) is for the conditional variance, and the forecast for the conditional variance is usually defined as the square root of the conditional variance forecast. The volatility forecast will approach the unconditional variance, $\bar{\sigma}^2$, as $h \to \infty$ if the process is covariance stationary (Enders 2010). How fast the forecast approaches $\bar{\sigma}^2$ is determined by $(\alpha_1 + \beta_1)$. A similar argument can be made for forecasting models that take asymmetry in to account.

The volatility forecasts from competing models can be evaluated by forecast evaluation statistics in order to find the most appropriate model, the two defined here are the mean square error (MSE) and mean absolute error (MAE). Assume that for a GARCH model i the h-step ahead forecast of σ^2_{T+h} at time T is $E_{i,T}[\sigma^2_{T+h}]$ with corresponding forecast error defined as $e_{i,T+h|T} = E_{i,T}[\sigma^2_{T+h}] - \sigma^2_{T+h}$. So for model i the forecast statistics are based on N out-of-sample forecasts from $T = T+1, \ldots, T+N$ are

$$MSE_i = \frac{1}{N} \sum_{j=T+1}^{T+N} e^2_{i,j+h|j} \tag{28}$$

$$MAE_i = \frac{1}{N} \sum_{j=T+1}^{T+N} |e_{i,j+h|j}| \tag{29}$$

The model that produces the smallest value of the two measures above is considered to be the preferable model (Zivot 2009).

4 Empirical Study - Swedish Inflation

The empirical study in this essay consists of finding an appropriate GARCH model for Swedish inflation. In Sweden, inflation is measured by Consumer Price Index (CPI), and the inflation rate between period t and period $t-1$, using the definition in (1), is then:

$$\pi_t = \log\big(CPI_t\big) - \log\big(CPI_{t-1}\big) \tag{30}$$

This essay will use inflation measured a quarterly frequency and a sample period of 1970:Q1-2014:Q4 and the data is taken from Statistics Sweden. Some summary statistics of the series is found in Table 1 below. In Table 1 the Jarque-Bera statistic and corresponding p-value are reported, and recalling that the null hypothesis the Jarque-Bera test is that the inflation series is normally distributed. On the basis of the p-value reported in Table 1, the null is rejected and other distribution will be considered in section 4.1. Based on Table 1 it is also worth to note that the kurtosis of the inflation series is much lower than three (kurtosis of a normal distribution). Distributions with kurtosis less than three are said to be platykurtic, meaning that the distribution produces fewer and less extreme outliers than the normal distribution.

Table 1: Descriptive statistics

Mean	Std.Dev	Skewness	Kurtosis	JB-statistic	p-value
0.0115	0.0120	0.8620	0.7306	26.2930	1.952e-06

Figure 1: Time series plot over inflation

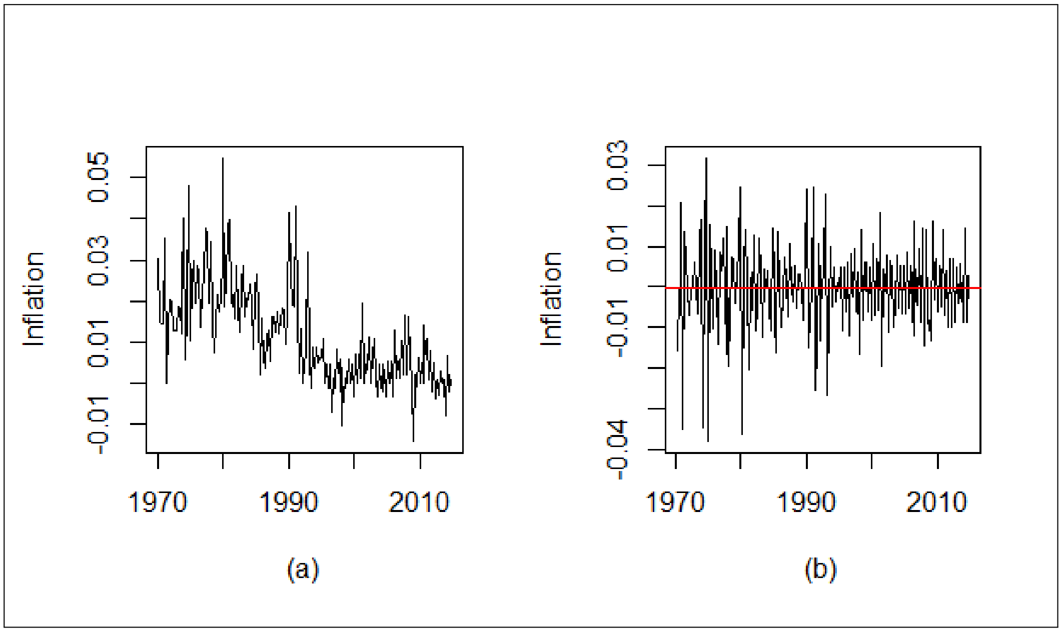

The series of inflation is depicted in Figure 1 above. In panel (a), the raw inflation series is depicted and a downward trend appear to be present, which is to be expected as the Riksbank (the Swedish central bank) adopted a inflation targeting regime in 1991 in order to bring inflation down. The trend is, however, not of interest for an analysis of inflation uncertainty (i.e. volatility analysis) and therefore the differenced series is used. The differenced series is depicted in panel (b) of Figure 1 and the series appears to be centered around its long-run mean (indicated by the red line). What is also present in panel (b) are volatility clusters, which indicates that the variance is not constant over time. The presence of volatility clusters means that ARMA(p,q) models cannot be employed, and thus indicating that a GARCH(p,q) model may be a better fit for the inflation series since such a model would take volatility clusters into account.

The presence of ARCH/GARCH effects in the inflation series is tested using the two formal tests described in section 3.6, namely the LM test proposed by Engle (1982) and the McLeod-Li test. In the LM test the null hypothesis is that there are no ARCH effects and for the inflation series the observed test statistic is $LM = 31.647$ with a corresponding p-value equal to 0.0016, so the null is rejected. In the second test, the McLeod-Li test the null hypothesis is again that there are no ARCH effects and for all 20 lag the null is rejected, the results are illustrated graphically in Appendix A.1 Figure 7. Thus based on the results from the these two tests we can conclude that GARCH effects are present.

Thus, a GARCH(p,q) model would appear to be appropriate for the inflation series. However, before any further analysis can be conducted the stationary of the series must be investigated. The three stationarity tests described in section 3.2 are conducted and the results are found in Table 2. For both the ADF test and the PP test the null hypothesis of a unit root can be rejected and one can conclude that the series is stationary. Note that in the KPSS test the null hypothesis is that the series is stationarity and at the conventional significance levels (0.05 or 0.01), the null cannot be rejected and once again the conclusion is that the series is stationary.

Table 2: Stationarity results

ADF	p-value	PP	p-value	KPSS	p-value
-6.5022	0.01	-221.61	0.01	0.02004	0.1

4.1 The Distribution of the Swedish Inflation Series

Previous sections mentioned that financial time series, such as our considered inflation series are in general not consider to be normally distributed (Ruppert 2010). Which has consequences for estimation of model parameters, since if the wrong distributional assumption is made one will not obtain efficient estimators and less accurate predictions (Zivot 2009). Thus, one important part in the GARCH modelling process is to find the distribution of the residuals, because it is this distribution that is used in the maximum likelihood estimation of the parameters.

In case the standardised residuals can be assumed to follow a normal distribution then the log-likelihood function that estimates the parameters in the GARCH-model based on the assumption of normality. Thus initially it is investigated if the Swedish inflation series follows a normal distribution. In Table 1 above, the results for the Jarque-Bera normality test are found and the null hypothesis of a normal distribution is rejected. The distributional assumption can also be graphically analysed, below in Figure 2 a histogram and QQ-plot for the inflation series is depicted. In panel (a) the red curve is the density curve from a normal distribution and as can be seen the histogram from the inflation series does not fit the curve very well. In panel (b) of Figure 2 the QQ-plot is depicted and as can be seen the inflation series seem to fit the line in the center of the distribution, while the tails deviates quite a bit from the line. Which is also mirrored in panel (a), thus on the whole the graphical analysis further strengthen the argument that the inflation series is not normally distributed.

Figure 2: Histogram and QQ-plot - Normal distribution

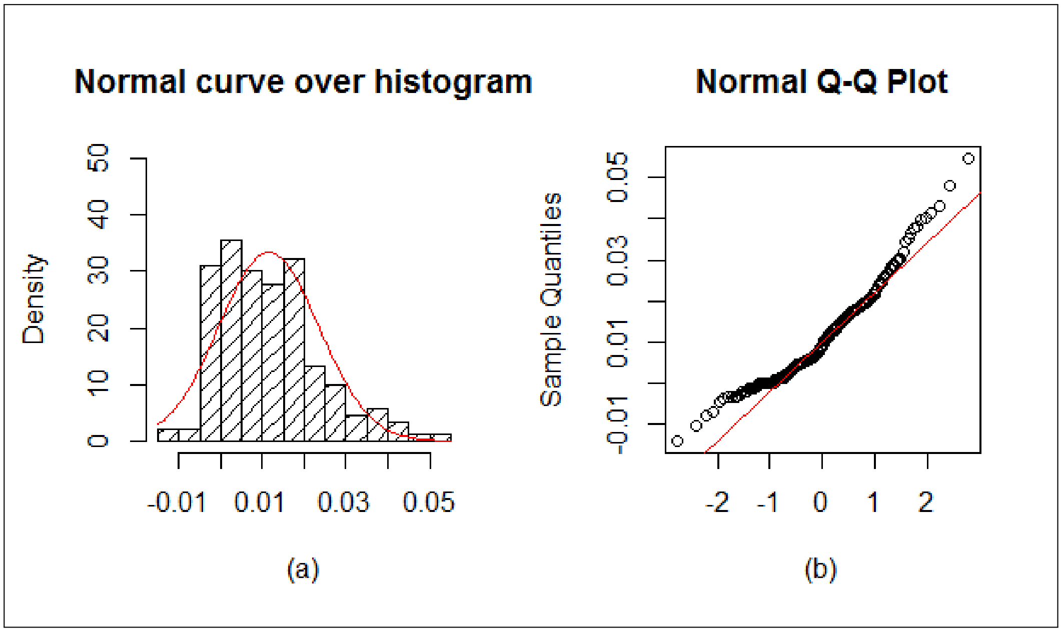

Since the normality assumption has been rejected, the next step is to consider the two distributions presented in section 3.3, i.e the Student's t-distribution and GED for the inflation series. The degrees of freedom (ν) for both of these distributions can be estimated using maximum likelihood which is done in order to find a theoretical counterpart that can be used to compare the empirical distribution with. The results from the ML-estimation is found in Table 3 and graphically analysed in Figure 3 below.

Table 3: Maximum likelihood estimated based on distributional assumption

Student's t			GED		
$\hat{\mu}_{ML}$	$\hat{\sigma}_{ML}$	ν_{ML}	$\hat{\mu}_{ML}$	$\hat{\sigma}_{ML}$	ν_{ML}
0.01077	0.01074	48.00149	0.01120	0.01192	2.79927

Figure 3: QQ-plots for the two distributional assumption for π_t

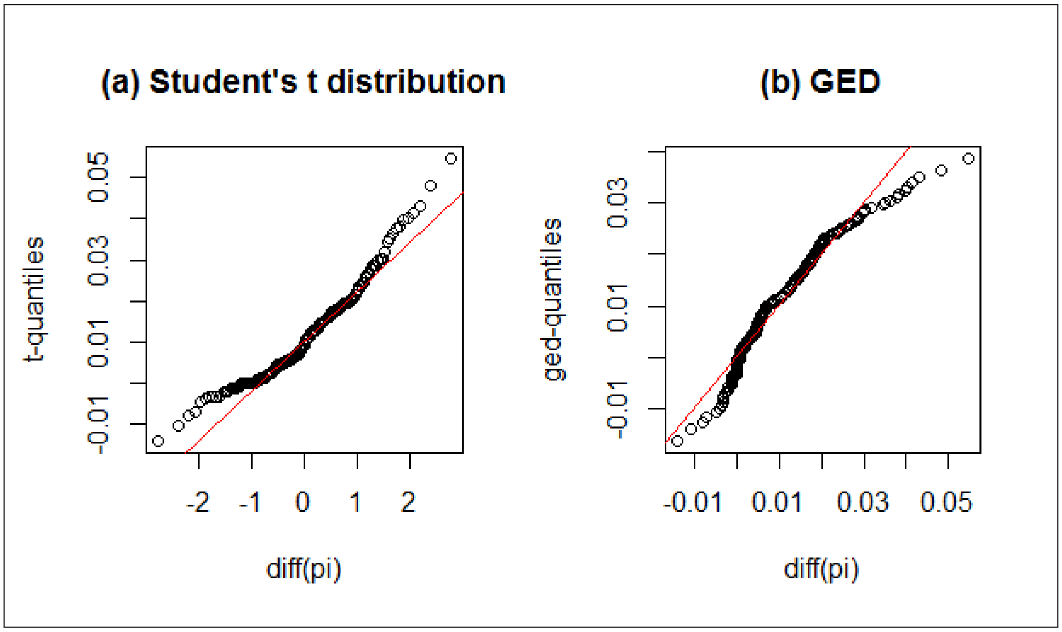

As can be seen in Figure 3 the fit of a Student's t-distribution does not significantly improve as compared to the normal fit, while the fit of a GED seems to ever so slightly improve the fit for the inflation series. It is, however important to stress that the unconditional distribution for the inflation series (i.e. the one obtained from ML estimation) is not necessarily the distribution that the residual belong to, since the estimates will be affected by the heteroskedasticity. But, on the other hand the unconditional distribution of the inflation series can indicate which distribution the conditional series and the residuals belong to (Ruppert 2010), the distribution of the errors is further discussed in the next section.

4.2 Estimating the GARCH Models

The first step of the GARCH-modelling procedure is to determine the adequate autoregressive model for the inflation series, i.e. to determine the mean equation. In earlier studies the specification for the mean equation has been found to be an AR-model of quite high order (three or more significant lags), the AR model found to fit our Swedish inflation series is an AR(4) model. The estimated model is presented below with standard errors in parenthesis below:

$$\pi_t = \underset{(0.0003)}{-0.0001} \underset{(0.0747)}{-0.7994}\,\pi_{t-1} \underset{(0.0928)}{-0.6035}\,\pi_{t-2} \underset{(0.0929)}{-0.3090}\,\pi_{t-3} \underset{(0.0754)}{+0.0669}\,\pi_{t-4} \tag{31}$$

Once the specification for the mean equation is found one can move on to the specification for the variance equation, or in other words the GARCH model is estimated. As there exist no consensus based on neither previous studies nor economic theory over which GARCH model is to be used for modelling inflation all models described in section 3.4 and 3.5 are estimated. However before this can be executed the distribution of the errors will be investigated further.

The results of the previous section was somewhat inconclusive on which distributional assumption to make for the errors. That is since the errors are effected by the heteroskedasticity the reasoning in the previous section does not definitely indicate that the error terms are non-normal. Therefore a GARCH model is estimated with the normal distribution and the distribution of the residuals from this model are investigated in Figure 4 and Table 4 below. As can be seen in Figure 4 the residuals deviate from the line in the tails of their distribution and thereby does not seem to fit the normal distribution very well. In Table 4 the normality of the residuals are tested with a Jarque-Bera test and the null hypothesis of the residuals belonging to the normal distribution is rejected. The kurtosis of the residuals is also presented in Table 4 and as can be seen the kurtosis is greater than three, and distributions with kurtosis greater than three are said to be leptokurtic. The same analysis was made for the residuals from GARCH models described in section 3.5 and for none of the models can the residuals be considered to be normally distributed, see Table 8 and Figure 8 in Appendix A.2.

Figure 4: QQ-plot of the residual from a GARCH(1,1)

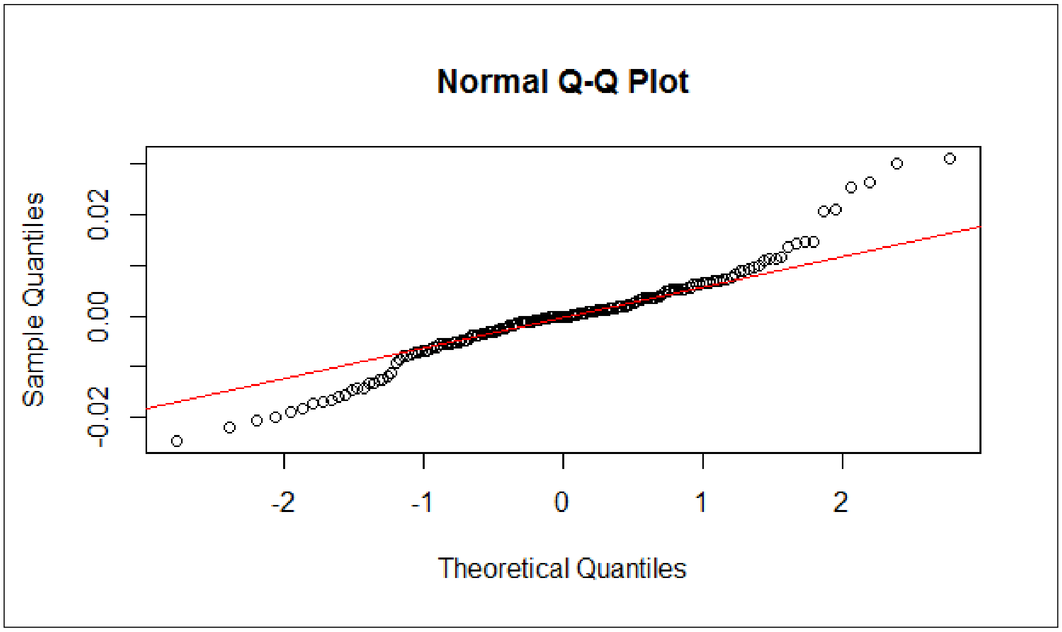

Table 4: Kurtosis and Jarque-Bera test for the residuals

Kurtosis	JB-statistic	p-value
4.0252	33.931	4.285e-08

Based on these results the GARCH models presented in section 3.4 and 3.5 will be estimated with distributions that have heaver tails, namely Student's t-distribution and GED. Therefore, and in contrast to the studies presented in the literature review[2], all four model specifications are estimated under both distributional assumptions. In Table 5, on the next page, all estimated coefficients are found for the four different types of GARCH models.

[2] Except Karanasos and Schurer (2008), who estimate a APARCH model using both the normal distribution and the Student's t-distribution.

Table 5: Estimated parameters from GARCH-models

Distribution	Student's t				GED			
Coefficients	GARCH(1,1)	EGARCH(1,1)	TGARCH(1,1)	APARCH(1,1)	GARCH(1,1)	EGARCH(1,1)	TGARCH(1,1)	APARCH(1,1)
Mean Equation								
μ	-0.0001	0.0003***	0.0001	0.0000	0.0000**	0.0003***	0.0000***	0.0000***
	(0.0002)	(0.0000)	(0.0002)	(0.0002)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
φ_1	-0.8239***	-0.7571***	-0.7552***	-0.7960***	-0.7531***	-0.7164***	-0.7452***	-0.7449***
	(0.0792)	(0.0024)	(0.0723)	(0.0681)	(0.0039)	(0.0272)	(0.0086)	(0.0071)
φ_2	-0.5185***	-0.3861***	-0.4563***	-0.4769***	-0.4494***	-0.4172***	-0.4440***	-0.4433***
	(0.0951)	(0.0079)	(0.0785)	(0.0837)	(0.0034)	(0.0181)	(0.0138)	(0.0156)
φ_3	-0.3645***	-0.2647***	-0.3312***	-0.3431***	-0.3033***	-0.2954***	-0.3040***	-0.3031***
	(0.0678)	(0.0021)	(0.0514)	(0.0823)	(0.0030)	(0.0032)	(0.0129)	(0.0157)
φ_4	0.1102*	0.1047***	0.1008*	0.1050	0.1661***	0.1270***	0.1615***	0.1624***
	(0.0651)	(0.0094)	(0.0559)	(0.0665)	(0.0016)	(0.0323)	(0.0064)	(0.0055)
Variance Equation								
α_0	0.0000	-0.1965***	0.0000***	0.0001	0.0000	-0.2260***	0.0000	0.0001
	(0.0000)	(0.0010)	(0.0000)	(0.0000)	(0.0000)	(0.0004)	(0.0000)	(0.0003)
α_1	0.1431	0.4392***	0.2592***	0.1228*	0.1284	0.4395***	0.2364	0.1136
	(0.2180)	(0.0032)	(0.0525)	(0.0651)	(0.1445)	(0.0110)	(0.2987)	(0.0707)
β_1	0.8416***	0.9784***	0.8947***	0.8894***	0.8427***	0.9757***	0.8723***	0.8815***
	(0.2096)	(0.0248)	(0.0199)	(0.0462)	(0.2341)	(0.0572)	(0.1820)	(0.0522)
γ_1		-0.1961***	-0.3430***	-1.0000***		-0.2104***	-0.2576	-1.0000***
		(0.0013)	(0.0599)	(0.0014)		(0.0002)	(0.3762)	(0.0014)
δ_0				1.1233				1.2237
				(0.6059)				(0.6750)
$\nu^{\star}$	3.4981**	4.6500***	3.3061***	3.4964***	0.9322***	1.1373***	0.9267***	0.9396***
	1.6248	0.0807	0.5486	0.9416	0.1634	0.0308	0.2544	0.1386
Information criteria								
AIC	-6.7925	-6.9168	-6.8113	-6.8011	-6.8203	-6.9256	-6.8300	-6.8227
BIC	-6.6322	-6.7387	-6.6333	-6.6052	-6.6601	-6.7475	-6.6519	-6.6269
Hannan-Quinn	-6.7275	-6.8446	-6.7391	-6.7216	-6.7554	-6.8534	-6.7578	-6.7433

Robust standard errors in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.
* are the maximum likelihood estimate of the degrees of freedom for the conditional distribution.

When reading Table 5 and looking for a preferable model for the Swedish inflation, a crude start is to look at the information criteria in the bottom panel of the table as an indicator of which models are preferable. The models that have the lowest value (i.e. most negative) of the information criteria are the models indicated to be preferable. Based on the criteria four preferable models emerge, namely the EGARCH and TGARCH under both distributional assumptions. For a definition of the information criteria and brief discussion see Appendix A.3. Narrowing the model selection further from these four, the next step is to look at significant parameters.

Out of the four relevant models, three pass through based on having significant parameters. That is the candidates to be a reasonable model for Swedish inflation are EGARCH(1,1) and TGARCH(1,1) under the assumption of Student's t-distribution and the EGARCH(1,1) model under the assumption of GED. Models of higher order was considered but none was found to have significant parameters (see Table 9 in Appendix A.4). Before any further model evaluation is performed, the parameters of the models are discussed further.

Recall that in a GARCH modelling setting the parameter α_1 reflects the influence of random deviations in the previous period on the conditional variance, whereas β_1 in turn measures the part of the realized variance in the previous period that is carried over into the current period. Thus the size of the parameters α_1 and β_1 determine the short-run dynamics of the resulting volatility time series. In particular large GARCH error coefficients, α_1, mean that volatility reacts intensely to movements in the economy. While large GARCH lag coefficients, β_1, indicate that shocks to conditional variance take a long time die out, or in other words the volatility is persistent. So referring back to Table 5 we see that all our three preferable models have coefficients in similar size and sign. It is also worth to note that across all three models we find quite high values for all the βs[3], meaning that the volatility in our inflation series is very persistent. Since volatility is our measure for inflation uncertainty, this means that if there is a shock to the Swedish inflation process it will take a long time for this uncertainty to die out. When comparing

[3] Recall from section 3.5 that the condition for covariance stationary for a EGARCH(1,1) model is $\beta_1 < 1$ and for a TARCH(1,1) model is $\gamma_1 < 1 - \alpha_1^2 - \beta_1^2 - 2\alpha\beta\nu_1$, conditions that are satisfied for the preferable models.

the results from this study to that of other studies that have also looked at small open economies, like Sweden, they have also found highly significant and persistent lag coefficients, see for example Fountas et al. (2004).

Next let's turn our attention to the leverage effects, that is the γs in Table 5, where we see that all estimated coefficients are statistically significant and negative thus indicating that there is asymmetry in the conditional variance. This implies that negative and positive shocks have different effects on inflation and inflation uncertainty. In particular, since γ_1 is negative this implies that a positive shock to the Swedish inflation leads to less uncertainty about inflation. It has been proposed that a negative asymmetry coefficient can be explained by the strong commitment of anti-inflationary policies by the central bank (Rogoff 1985), which was also empirically verified by Alesina and Summers (1993) who found that more independent central banks are linked to both lower inflation and inflation uncertainty. During the 1990s the central bank in Sweden, the Riksbank, gained more independence and its independence was written into law in 1999 (the Riksbank 2011). At this point it would be interesting refer back to the two competing hypotheses of the causal relationship between inflation and its uncertainty that were presented in the introduction. The first one denoted the Friedman-Ball hypothesis postulates that inflation causes inflation uncertainty. Whereas the second hypothesis is denoted the Cukierman-Meltzer hypothesis which postulates that inflation uncertainty causes inflation. Before any further discussion is made some caveats should be stated, first in the other studies mention throughout this study Granger-causality tests are employed to detangle the direction of causality,[4] however such tests are not employed in this essay. Thus the results found in this study will only indicate the direction of causality between inflation and its uncertainty, keeping in mind that there could be reverse causality. Now let's consider the Swedish inflation series. The results found in Table 5 does not support the Friedman-Ball hypothesis, this becomes clear by using the argument of Ball (1992). That is, an increase in the Swedish inflation would not lead to an increase in inflation uncertainty as the Riksbank has a strong commitment to anti-inflationary policies and would thus be willing to bear the cost of a decrease in the rate of inflation. With regards

[4] Note that the Granger-causality test tests for the presence of Granger causality which is not necessarily true causality. That is the Granger test finds only what is referred to as predictive causality (Diebold 2001).

to the Cukierman-Meltzer hypothesis, the results seem to support a version of this hypothesis put forward by Holland (1995). This version, sometimes called the "stabilisation hypothesis", assumes the same causality but with a negative relationship, suggesting that in an economy where inflation uncertainty leads to inflation one would expect the central bank to try to stabilise inflation and thus have a negative effect of inflation uncertainty on inflation.

In order to validate the analysis made above the next step is to evaluate our preferable models using the methods described in section 3.8. Thus, if one has estimated a correct model then the standardised residuals should behave in the same way as in classical linear regression. So the standardised residuals should exhibit neither serial correlation, conditional heteroskedasticity nor any non-linear dependence. Moreover, the distribution of the standardised residuals should reflect the distributional assumption made in the estimation of the model. The results from the model evaluation are found in Table 6 and Figure 5. The first panel of Table 6 tests whether or not there are serial correlation in the standardised residuals, and recall that in the Ljung-Box test the null hypothesis is that there is no serial correlation. As can be seen there are significant lags (the 11th and 19th lag) however since the significant lags are quite far away they are deemed to not carry that much weight and thus allow our models to pass this requirement. The second panel of Table 6 tests whether or not the squares of the standardised residuals exhibits serial correlation or conditional heteroskedasticity. Again the Ljung-Box test is employed, and here the null cannot be rejected for all lags. Thus we can conclude that there is no conditional heteroskedasticity in the standardised residuals, indicating that the models fit the data. The third and final panel of Table 6 shows the results from the ARCH LM test, recall that the null hypothesis in this test is that there are no ARCH effects. As can be seen the null cannot be rejected for all but one lag, the 7th lag for the EGARCH(1,1) estimated with GED. Thus based on these three tests our preferable models seem to fit the data quite well. The next step is to investigate the distribution of the standardised residuals.

Table 6: Model evaluation

p-values			
Distribution	Student's t		GED
Lag	EGARCH(1,1)	TGARCH(1,1)	EGARCH(1,1)
Weighted Ljung-Box Test on Standardised Residuals			
1	0.9280	0.7439	0.6447
11	0.0000	0.0015	0.0000
19	0.0166	0.0703	0.0028
Weighted Ljung-Box Test on Standardised Squared Residuals			
1	0.4645	0.6310	0.3616
5	0.5398	0.6365	0.2378
9	0.6095	0.6824	0.2557
Weighted ARCH LM Tests			
3	0.6546	0.7476	0.7622
7	0.2589	0.3852	0.0644
7	0.3968	0.5139	0.1150

Null hypothesis in the ARCH LM test is "No ARCH effects" and the null hypothesis in the Ljung-Box test is "No serial correlation".

Degrees of freedom for the tests on the standardized residuals and their squares are, respectively 4 and 2.

Lag for * is calculated as Lag[2*(p+q)+(p+q)-1] and lag for ** is calculated as Lag[4*(p+q)+(p+q)-1] where p=1 and q=1

Figure 5: QQ-plots of the standardised residuals

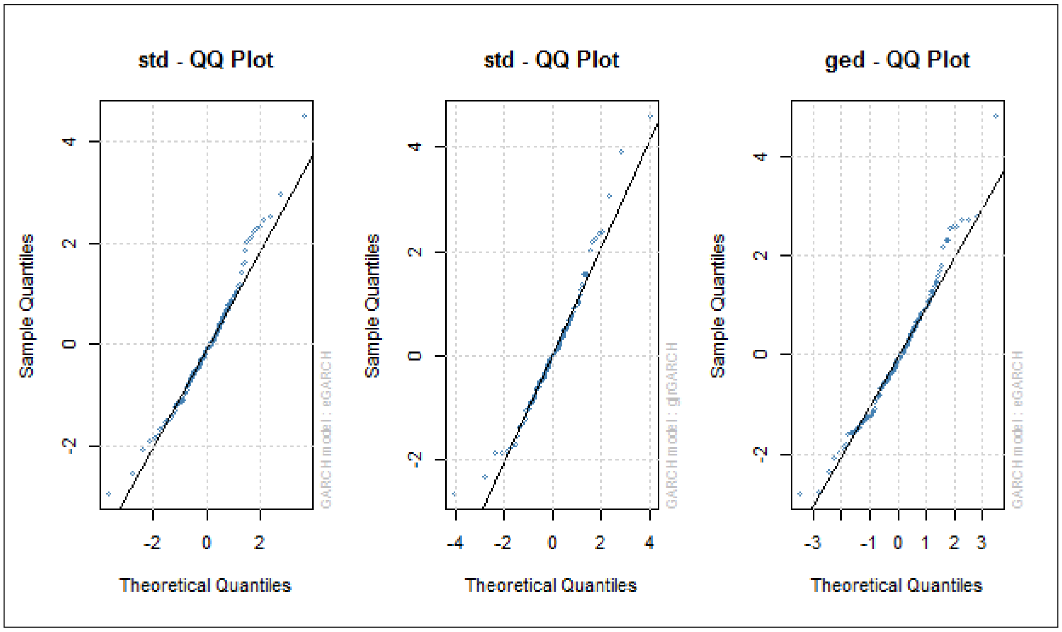

The distribution of the standardised residuals is investigated by the means of QQ-plots which are depicted in Figure 5. Note that in Figure 5 the panel to left depicts the QQ-plots of residuals from the EGARCH(1,1) estimated with the Student's t distribution, under the same distributional assumption the middle panel depicts the QQ-plots of residuals from the TARCH(1,1) model, and final the panel to the right depicts the QQ-plots of residuals from the EGARCH(1,1) model estimated with GED. Based on the figure above we can conclude that the residuals seem to fit the distributional assumption made for each of the three models. Thus on the whole the conclusion that can be drawn based on Table 6 and Figure 5 is that none the preferable models for the Swedish inflation series yield any evidence of misspecification. In the next section we move on to forecasting the volatility of the inflation series in another way to evaluate the models.

4.3 Forecasting Volatility

Another approach to evaluation, and thus to find the preferred model, is to look at the forecasted volatility of our competing three models and then look at the forecast evaluation statistics presented in section 3.9. Figure 6 shows the h-step ahead volatility forecasts ($h = 1, \ldots, 250$) from our three models, the horizontal line in the figure is the unconditional standard deviation (also presented in the first column in Table 7). As is expected from models that take asymmetry into account, we find very steep curves depicted in Figure 6 especially in the beginning of the forecast horizon but as the horizon moves further and further away the volatility forecasts approach the unconditional standard deviation.

Figure 6: Volatility Forecasts

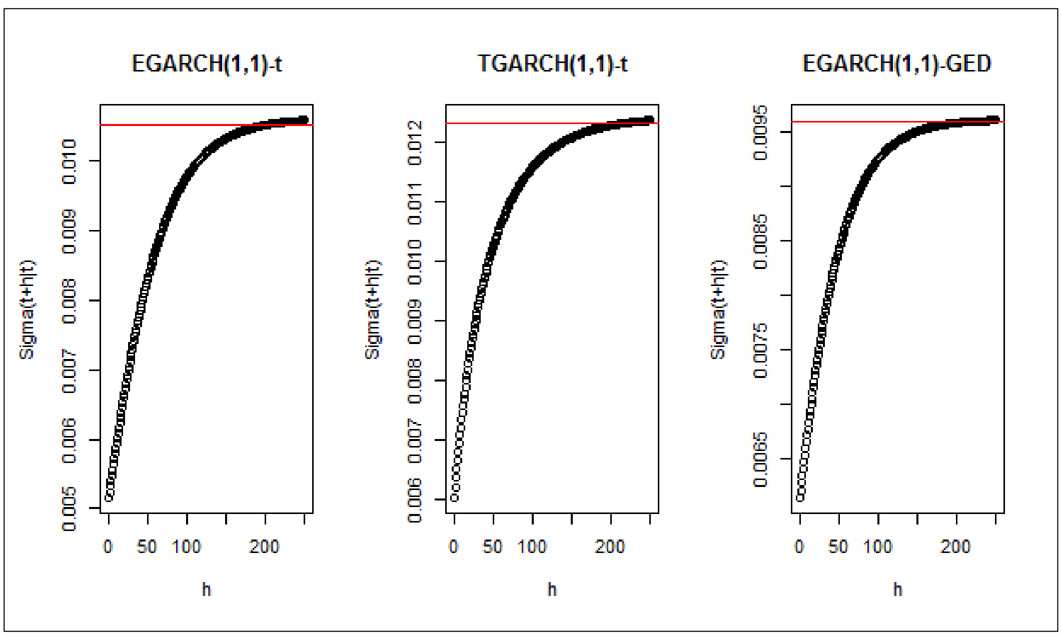

Table 7: Unconditional standard deviation and forecast evaluation statistics

Model	$E[\bar{\sigma}]$	MSE	MAE
EGARCH(1,1)-t	0.0105	3.515e-05	0.0046
TGARCH(1,1)-t	0.0123	4.798e-05	0.0056
EGARCH(1,1)-GED	0.0096	4.512e-05	0.0054

As previously mention the forecasted volatility from our three preferable models can be evaluated with forecasts evaluation statistics defined in section 3.9, these are presented in the last two column of Table 7. Recall that the model that produces the smallest value of the two measures is considered to be the preferable model. Therefore the model that the forecast evaluation statistics advocates is the EGARCH(1,1) estimated with Student's t distribution.

Summarising the results found from all the model evaluation techniques used in both this section and the previous one, in the previous section we found no evidence of misspecification for either of the models, so not much aide in selection of a preferable model. However, based on the forecast evaluation statistics the most appropriate model for the Swedish inflation series is the EGARCH(1,1) estimated with Student's t distribution.

5 Summary and Concluding Remarks

Inflation and its uncertainty for Sweden during the period 1970:Q1-2014:Q4 have been investigated in this essay. The primary goal of this essay was to find an appropriate model for the inflation series in a GARCH modelling setting allowing for non-normal errors. The secondary goal was to relate the results of the model estimation to the two competing hypotheses of the causal relationship between inflation and inflation uncertainty, the Friedman-Ball and the Cukierman-Meltzer hypotheses, and thus see in which camp the Swedish inflation series falls.

In contrast to other studies that have looked at this relationship this essay investigated the distribution of the errors, and did indeed find that the residuals did not follow a normal distribution but instead more appropriate distributional assumptions seemed to be the Student's t-distribution and the Generalised error distribution (GED). Under these two distributional assumptions three models emerge as contenders for the model for the Swedish inflation, namely the EGARCH(1,1) and TGARCH(1,1) that is models that take asymmetry into account. After evaluating the assumptions underlying the model and on the basis of forecast evaluation statistics, the preferable model for Swedish inflation and its uncertainty that emerged was an EGARCH(1,1) estimated with a Student's t-distribution.

For the estimated model highly significant and persistent lag coefficient was found, meaning that the volatility in the Swedish inflation series takes a long time to die out. Furthermore a significant leverage effect was found, or in other words the results found that there is asymmetry in the conditional variance, thus implying that negative and positive shocks have different effects on inflation and inflation uncertainty. In particular, since γ_1 is negative this implies that a positive shock to the Swedish inflation leads to less uncertainty about inflation. Relating these results to the two competing hypotheses, the results in this essay does not support the Friedman-Ball hypothesis, i.e. an increase in the Swedish inflation would not lead to an increase in inflation uncertainty as the Riksbank has a strong commitment to anti-inflationary policies and would thus be willing to bear the cost of a decrease in the rate of inflation. With regards to the Cukierman-Meltzer hypothesis, the results seem to support a version of this hypothesis put forward by Holland (1995).

This version, is sometimes called the ”stabilisation hypothesis” and assumes the same causality but with a negative relationship, suggesting that in an economy where inflation uncertainty leads to inflation one would expect the central bank to try to stabilise inflation and thus have a negative effect of inflation uncertainty on inflation.

An interesting question for further research would be to further delve into the distributional assumptions made for the errors, perhaps considering other heavy-tailed distribution used for financial time series. Another interesting approach is to consider other long memory GARCH-models.

6 References

Alesina, A. & Summers, L. H., 1993. Central Bank Independence and Macroeconomic Performance: Some Comparative Evidence. *Journal of Money, Credit & Banking* , 25(2), pp. 151-162.

Baillie, R. T., Chung, C.-F. & Tieslau, M. A., 1996. Analysing Inflation by the Fractionally Integrated ARIFIMA-GARCH Model. *Journal of Applied Econometrics*, 11(1), pp. 23-40.

Ball, L., 1992. Why Does High Inflation Raise Inflation Uncertainty?. *Journal of Monetary Economics.*, 29(3), pp. 371-388.

Blanchard, O., 2000. *Macroeconomics.* 2 ed. Englewood Cliffs, N.J: Prentice Hall.

Bollerslev, T., 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 31(3), pp. 307-327.

Bollerslev, T., 1987. A Heteroskedastic Time Series Model for Speculative Prices Rates of Return. *The Review of Economics and Statistics*, 69(3), pp. 542-547.

Burnham, K. P. & Anderson, D. R., 2002. *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach.* 2nd ed. Berlin Heidelberg: Springer-Verlag.

Caporale, B. & Caporale, T., 2002. Asymmetric Effects of Inflation Shocks on Inflation Uncertainty. *Atlantic Economic Journal.*, 30(4), pp. 385-389.

Caporale, T. & McKiernan, B., 1997. High and variable inflation: Further evidence on the Friedman hypothesis. *Economics Letters* , 54(1), pp. 65-68.

Cukierman, A. & Meltzer, A. H., 1986. A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information. *Econometrica*, 54(5), pp. 1099-1128.

Davidson, R. & MacKinnon, J. G., 2009. *Econometric Theory and Methods.* International ed. New York: Oxford University Press.

Diebold, F. X., 2001. *Elements of Forecasting.* 2nd ed. Cincinnati: South Western.

Enders, W., 2010. *Applied Econometric Time Series.* 3rd ed. Hoboken, NJ: Wiley.

Engle, R. F., 1982. Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of the United Kingdom Inflation. *Econometrica*, 50(4), pp. 987-1007.

Engle, R. F., 1983. Estimates of the variance of U.S. inflation based on the ARCH model. *Journal of Money, Credit and Banking*, 15(3), pp. 286-301.

Fountas, S., Ioannidis, A. & Karanasos, M., 2004. Inflation, Inflation Uncertainty and Common European Monetary Policy. *The Manchester School*, 72(2), pp. 221-242.

Friedman, M., 1977. Nobel Lecture: Inflation and Unemployment. *Journal of Political Economy*, 85(3), pp. 451-472.

Glosten, L., Jagannathan, R.& Runkle, D. E., 1993. On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance*, 48(5), pp. 1779-1801.

Golob, J., 1993. *Inflation, inflation uncertainty and relative price variability: A survey*, Working paper no. 93-15: Federal Reserve Bank of Kansas City.

Gordon, R. J., 1975. Alternative Responses of Policy to External Supply Shocks. *Brookings Papers on Economic Activity*, 1975(1), pp. 183–206.

Grier, K. B.& Perry, M. J., 1998. On inflation and inflation uncertainty in the G7 countries. *Journal of International Money and Finance*, 17(4), pp. 671-689.

Hansen, P. & Lunde, A., 2004. A Forecast Comparison of Volatility Models: Does Anything Beat a GARCH(1,1)?. *Journal of Applied Econometrics*, 20(7), pp. 873–889.

He, C. & Teräsvirta, T., 1999. Properties of moments of a family of GARCH processes. *Journal of Econometrics*, 92(1), pp. 173–192.

Holland, A. S., 1995. Inflation and Uncertainty: Tests for Temporal Ordering. *Journal of Money, Credit and Banking* , 27(3), pp. 827-837.

Holland, S. A., 1993. Comment on inflation regimes and the sources of inflation uncertainty. *Journal of Money, Credit and Banking* , 25(3), pp. 514-520.

Karanasos, M. & Schurer, S., 2008. Is the Relationship between Inflation and Its Uncertainty Linear?. *German Economic Review*, 9(3), pp. 265–286.

Lee, J. H. H. & King, M. L., 1993. A Locally Most Powerful Based Score Test for ARCH and GARCH Regression Disturbances. *Journal of Business and Economic Statistics*, 11(1), pp. 17-27.

Nazar, D., Farshid, P.& Mojtaba, K. Z., 2010. Asymmetry Effects of Inflation on Inflation Uncertainty in Iran: Using from EGARCH Model 1959-2009. *American Journal of Applied Sciences*, 7(4), pp. 535-539.

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

OECD, 2005. *Glossary of Statistical terms.* [Online]
Available at: http://stats.oecd.org.ludwig.lub.lu.se/glossary/detail.asp?ID=2146
[Accessed 17 January 2016].

OECD, 2013. *Glossary of Statistical Terms.* [Online]
Available at: http://stats.oecd.org.ludwig.lub.lu.se/glossary/detail.asp?ID=427
[Accessed 17 January 2016].

Phillips, P. C. B. & Perron, P., 1988. Testing for a Unit Root in Time Series Regression. *Biometrika*, 75(2), pp. 335-346.

Pourgerami, A. & Maskus, K. E., 1985. The Effects of Inflation on the Predictability of Price Changes in Latin America: Some Estimates and Policy Implications. *World Development*, 15(2), pp. 287-290.

Riksbank, 2011. *The Riksbank: Sveriges Riksbank, 10 years as an independent central bank*. [Online]
Available at: http://www.riksbank.se/en/The-Riksbank/History/Sveriges-Riksbank-10-years-as-an-independent-central-bank/ [Accessed 16 Febuary 2016].

Rogoff, K., 1985. The Optimal Degree of Commitment to an Intermediate Target. *The Quarterly Journal of Economic*, 100(4), pp. 1169-1189.

Ruppert, D., 2010. *Statistics and Data Analysis for Financial Engineering.* 1 ed. New York: Springer Verlag.

Ungar, M. & Zilberfarb, B.-Z., 1993. Inflation and its unpredictability – theory and empirical evidence. *Journal of Money, Credit and Banking* , 25(4), pp. 709-720.

Walgenbach, P. H., Dittrich, N. E. & Hanson, E. I., 1973. *Financial Accounting.* New York: Harcourt Brace Javonovich, Inc..

Zakoian, J.-M., 1994. Threshold heteroskedastic models. *Journal of Economic Dynamics and Control*, 18(5), pp. 931–955.

Zivot, E., 2009. Practical Issues in the Analysis of Univariate GARCH Models. In: T. G. Andersen, R. A. Davis, J. Krieß & T. Mikosch, eds. *Handbook of financial time series.* Berlin Heidelberg: Springer Verlag, pp. 114-155.

7 Appendix

A.1 The McLeond-Li test

Figure 7: McLeod-Li test

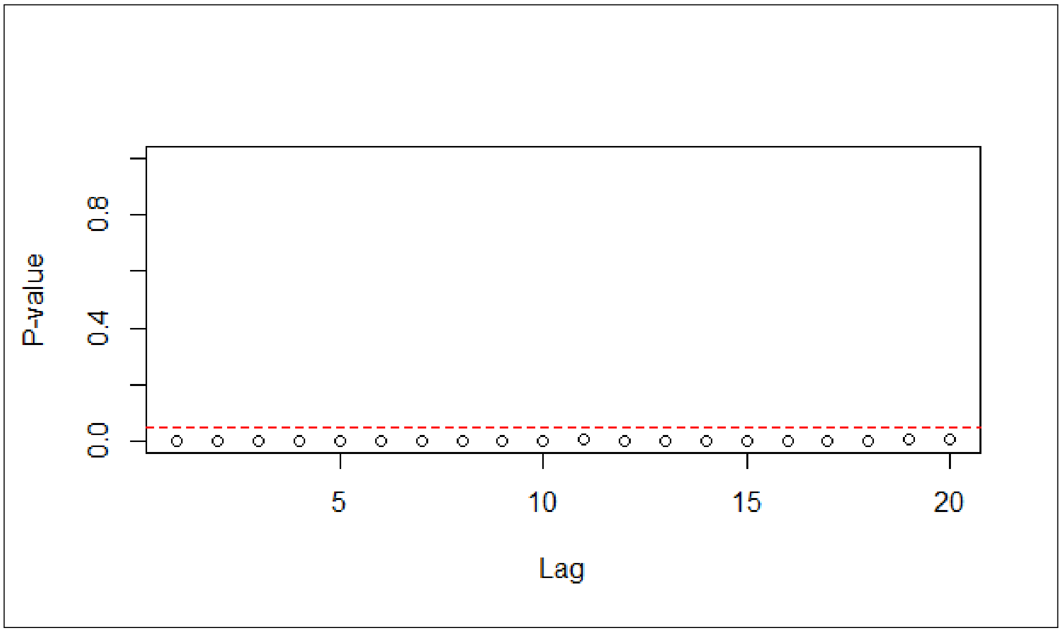

In the McLeod-Li test the null hypothesis is that there are no ARCH effects and the result is illustrated graphically in Figure 7 above. The red barred line in Figure 7 indicate the significance level and as can be seen for all 20 lags the p-values are less than the significance and the null is rejected.

A.2 Residuals from Models Estimated with the Normal Distribution

Figure 8: QQ-plots for the residuals from models estimated with the normal distribution

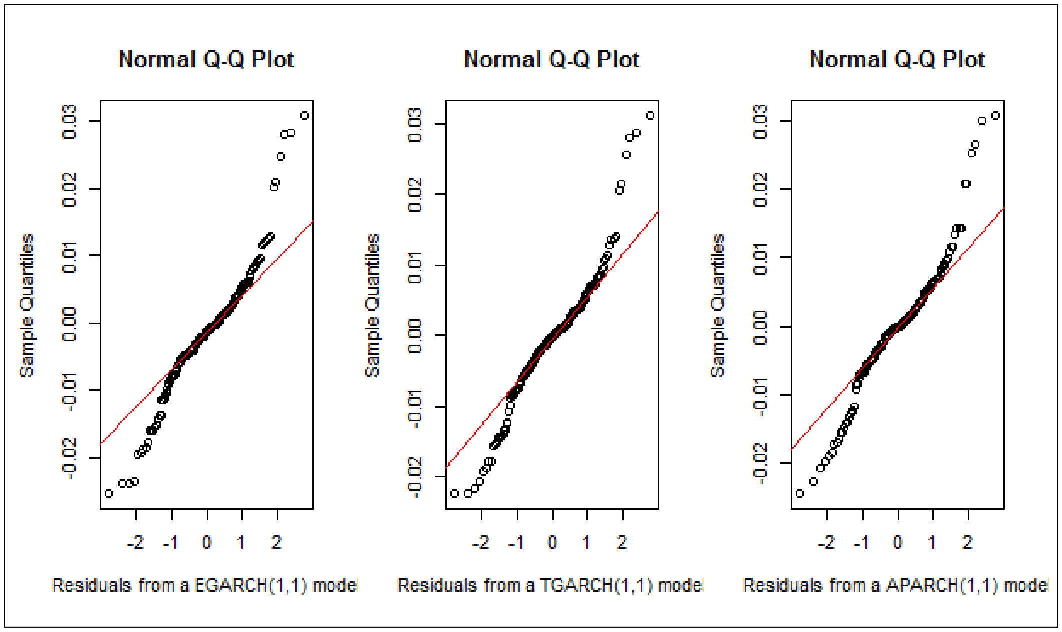

Table 8: Kurtosis and Jarque-Bera test for the residuals

Model	Kurtosis	JB	p-value
EGARCH(1,1)	5.2566	42.565	5.718e-10
TGARCH(1,1)	4.0767	36.812	1.015e-08
APARCH(1,1)	3.9898	32.528	8.643e-08

A.3 Information Criteria

A.3.1 The Akaike Information Criterion (AIC)

Define L as the maximum value of the likelihood function for the estimated model, furthermore let k be the number of estimated parameters in the model. Then the AIC value of the model is the following:

$$AIC = 2k - 2\log(L) \tag{32}$$

Thus, given a set of possible model specifications, the preferred model will be the one with the minimum AIC value. That is the AIC rewards goodness of fit which is assessed by the likelihood function, but it also includes a penalty that is an increasing function of the number of estimated parameters. The penalty thereby discourages overfitting (Burnham and Anderson 2002).

A.3.2 The Bayesian Information Criterion (BIC)

Define L as the maximum value of the likelihood function for the estimated model, furthermore let k be the number of estimated parameters in the model, and n is the number of observations. Then the BIC value of the model is the following:

$$BIC = -2\log(L) + 2k \cdot \log(n) \tag{33}$$

Note that both BIC and AIC resolve the problem of overfitting by introducing a penalty term for the number of parameters in the model, and the penalty term is larger in BIC than in AIC (Burnham and Anderson 2002).

A.3.3 The Hannan–Quinn Information Criterion

Define ℓ as the log-likelihood function for the estimated model, furthermore let k be the number of estimated parameters in the model, and n is the number of observations. Then the Hannan–Quinn value of the model is the following (Burnham and Anderson 2002):

$$Hannan\text{–}Quinn = -2\ell + 2k \log\ \log\ n \tag{34}$$

A.4 GARCH Models of Higher Orders

Table 9: GARCH models of higher order

Distribution:	Student's t-distribution			
Model	Significant parameters *	AIC	BIC	Hannan-Quinn
GARCH(1,2)	No	-6.7820	-6.6039	-6.7098
GARCH(2,1)	No	-6.7813	-6.6033	-6.7091
GARCH(2,2)	No	-6.7708	-6.5749	-6.6914
EGARCH(1,2)	No	-6.8888	-6.6751	-6.8021
EGARCH(2,1)	No	-6.8888	-6.6751	-6.8021
EGARCH(2,2)	No	-6.8790	-6.6475	-6.7851
TGARCH(1,2)	No	-6.8348	-6.6389	-6.7553
TGARCH(2,1)	**			
TGARCH(2,2)	No	-6.7832	-6.5517	-6.6893
APARCH(1,2)	No	-6.7899	-6.5762	-6.7032
APARCH(2,1)	No	-6.7781	-6.5466	-6.6842
APARCH(2,2)	No	-6.7646	-6.5153	-6.6635
Distribution:	GED			
Model	Significant parameters *	AIC	BIC	Hannan-Quinn
GARCH(1,2)	No	-6.8096	-6.6316	-6.7374
GARCH(2,1)	No	-6.8092	-6.6311	-6.7370
GARCH(2,2)	No	-6.7985	-6.6026	-6.7190
EGARCH(1,2)	No	-6.9147	-6.7189	-6.8353
EGARCH(2,1)	No	-6.9085	-6.6948	-6.8218
EGARCH(2,2)	No	-6.8945	-6.6630	-6.8007
TGARCH(1,2)	**			
TGARCH(2,1)	No	-6.8716	-6.6580	-6.7850
TGARCH(2,2)	**			
APARCH(1,2)	No	-6.8116	-6.5979	-6.7249
APARCH(2,1)	No	-6.7894	-6.5579	-6.6955
APARCH(2,2)	No	-6.7837	-6.5344	-6.6826

* Not significant is defined as at least one parameter not significant

** Covariance stationary requirement not fulfilled.