

A Copula Approach to Modeling Extreme Values of Exchange Rates

Annika Johansson

March 6, 2014

Abstract

In this thesis we consider a general approach to modeling dependence in extreme values of exchange rates by using copulas. As specific examples the following pairs of currencies are analyzed: Swedish krona to U.S. dollar (SEK/USD), Swedish krona to Euro (SEK/EUR), Swedish krona to British pound (SEK/GBP), Swedish krona to Japanese Yen (SEK/JPY), Swedish krona to Danish krone (SEK/DKK) and Swedish krona to Norwegian krone (SEK/NOK). The daily log-return series are first modeled individually using ARMA-GARCH models. In some cases, when it is statistically significant, we also use three month Stockholm Interbank Offered Rate (STIBOR) as a covariate in our models. After the models have been fitted to the datasets the residuals are considered for further analysis. We fit several bivariate copula models to the residual series and use different measures of goodness-of fit to choose one between competing models. We demonstrate the flexibility of the approach by repeating our analysis both for the original residuals as well as the monthly and quarterly extreme values of the series.

Keywords. GARCH, EGARCH, copula, extreme values, exchange rates, modelling

Acknowledgements

I would like to thank my supervisor Nader Tajvidi for all his support and suggestions. I am also grateful for all the interesting courses that are given at the Centre for Mathematical Sciences at Lund University which have enabled me to write this thesis. At last I would like to express my gratitude towards the Technische Universität Berlin where I during my exchange got the opportunity to deepen my understanding of the effects of monetary policies, which has been very useful in this thesis.

Contents

1	Introduction	5
2	Information about the exchange rate series	6
3	Used models and their distributions	7
3.1	The ARMAX(r, m, v) - GARCH(p, q) model	8
3.2	The EGARCH(p, q) model	9
3.3	The copula function	9
3.3.1	The t-copula	10
3.3.2	The Plackett copula	10
3.3.3	The Gumbel copula	11
3.3.4	The Hustler-Reiss copula	11
3.3.5	The Galambos copula	11
3.3.6	The Tawn copula	11
3.4	The distributions of innovations and residuals	11
3.4.1	The Student's t distribution	12
3.4.2	The skewed Student's t distribution	13
3.4.3	The Logistic distribution	13
3.4.4	The Generalized Extreme Value distribution	13
4	Validation of the models	13
4.1	The autocorrelation function and the partial autocorrelation function	14
4.2	The Lagrange Multiplier Test of Engle	16
4.3	Q-Statistics	16
4.4	Likelihood Ratio Test	17
4.5	Information Criteria	17
4.6	Sign Bias Test	18
4.7	The GMM Orthogonality Test	18
4.8	The Non-Parametric Density Test of Hong and Li	18
4.9	Kupiec's and Christoffersen's test	19
4.10	Measures of dependence	21
4.11	Goodness-of-fit test for copulas	22
5	Method of finding suitable models	24
6	Result	25
6.1	Model choice - ARMA-GARCH models for the daily returns	25
6.2	Model choice - copula models for the residuals	28
7	Conclusions and fields of further studies	33
A	Appendix: plots and tables	37

1 Introduction

In this thesis several different exchange rates will be studied, namely SEK/USD (Swedish krona to U.S. dollar), SEK/EUR (Swedish krona to Euro), SEK/GBP (Swedish krona to British pound), SEK/JPY (Swedish krona to Japanese Yen), SEK/DKK (Swedish krona to Danish krone) and SEK/NOK (Swedish krona to Norwegian krone). The daily log-return series will first be modelled individually using ARMA-GARCH models. In cases when it adds information to the model an external regressor will be used as an input signal, in this case the three month Stockholm Interbank Offered Rate, $STIBOR$, which consists of daily observations of the reference rate that is based on the interest rate Swedish banks are willing to offer each other for lending money over a three-months period. After the data has been fitted to the models the residuals will be considered, since these might contain information about the series that the ARMA-GARCH models were not able to replicate. Since the exchange rate market in general, and exchange rates with one currency in common in particular, are affected by the same external factors, it seems probable that the dependence between them will be high. With this in mind the residual series will be bivariate fitted to copula models; the original residual series as well as the monthly and quarterly extreme values of the series. To find the ARMAX-GARCH models for the exchange rate returns the package *rugarch* in R will be used. For estimation of the copula models for the residuals the R-package *copula* will be used.

Before starting a few things should be noticed. Some aspects make financial time series special. By observing daily financial time series across a wide range of products like log-returns of exchange rates, equities and commodity prices it can be seen that these time series show special characteristics that are less common in time series within other fields. Based on these empirical observations a number of statistical properties have been recognized over the years and now go under the name of *stylized facts*. These are:

- Return series, although they show very little serial correlation, are not i.i.d.
- Absolute or squared returns show serial correlation.
- The conditional expected returns are close to zero.
- Volatility seems to vary over time.
- The return series are leptokurtic, i.e. heavy-tailed.
- Extreme values appear in clusters.

In general this is true for the daily returns, but for longer time horizons the return series start to show less signs of volatility clustering, become less heavy-tailed and more i.i.d. This can be explained through the fact that the central limit theorem applies to the GARCH models that in general can be used to model these time series.[11] These stylized facts explain the need and development of the GARCH models, which allow the volatility to change over time.

Next we will consider the dynamics between two currencies. To buy a currency also means selling another currency. Hence an exchange rate can both be expressed as *foreign currency/domestic currency* and *domestic currency/foreign currency*. The first system is called *British terms* and the second system is called *European terms*. The exchange rates used in this thesis are quoted in European terms. How the exchange rate changes is complex and depends on many factors, but to give a very rough and simplified picture of it, it works in the following way. It is assumed that there is no arbitrage in the exchange rate market, that is, it is not possible to e.g.

buy Euro by paying in Swedish krona, use the Euro to buy U.S. dollars and then use the U.S. dollars to buy Swedish krona again and thereby getting more money than what was originally used to buy currency. It is also assumed that if a product is much cheaper in one country than in another, the country with the expensive product will buy the cheap product from the other country, thereby increasing the import from that country as well as the demand of that currency. But since the demand of that currency increases, it will also become more expensive to buy the currency, until in the long run the exchange rate has changed so that the price of the product is the same in both countries. Another factor that affect the exchange rate is inflation. Assume that one country prints money to pay for its expenses. This will lead to inflation in this country, and if one country has a higher level of inflation than another, then the money supply in the first country will be larger than before, causing the exchange rate to change so that the currency of the first country costs less than previously, measured in the currency of the other country. A way for a country to control the inflation is by using the prime lending rate. This is the rate to which the central bank lends money to banks. Since the banks want to make a profit they can never lend money to their customers to a rate lower than this one. In the same way the banks will never borrow money from customers to a rate higher than this one. If the prime lending rate is increased, this means that both the banks' rate for borrowing and lending will increase. This makes it more attractive for foreign investors to invest their money in that market, since, simplified, it will give a higher return to put the money into a bank account in that country than in the home country. Likewise people will be more reluctant to take loans and make investments. This will at first increase the demand of that currency, making it more expensive to buy. With the new exchange rate, it will be profitable to import foreign goods, which again will change the exchange rate until a new equilibrium has been reached.[5] A lower prime lending rate would make borrowing more tempting, both for domestic and foreign investors. This would lead to a larger supply of money in the market, the inflation would increase and the money could be invested in production, leading to more goods to consume as well as to export. The increased export would then increase the value of the domestic currency again, and a new equilibrium would be found. In Sweden the prime lending rate is set by the central bank only a few times a year and is hence not a very good source of information when daily exchange rate observations should be modelled, but as mentioned the STIBOR rate gives daily information about what rate the largest banks in Sweden are willing to offer each other and follow the prime lending rate very closely. Since this is an indirect indicator of the state of the Swedish economy, the choice was made to see if this rate could be used as an external regressor.

A closer look at the properties of the exchange rate data series will be taken in section 2, followed by section 3 on theory about the models and section 4 about different tests and measures that are used for finding a suitable model. After that the procedure that was used to find the models will be stated in section 5, followed by the results, i.e. the models that were chosen and their parameter values, in section 6. At last, in section 7, conclusions and fields of further studies will be discussed.

2 Information about the exchange rate series

As explained in the introduction six different time series of exchange rates will be considered. These are *SEK/USD* (Swedish krona to U.S. dollar), *SEK/EUR* (Swedish krona to Euro), *SEK/GBP* (Swedish krona to British pound), *SEK/JPY* (Swedish krona to Japanese Yen), *SEK/DKK* (Swedish krona to Danish krone) and *SEK/NOK* (Swedish krona to Norwegian krone). For the Swedish krona and Danish krone and the Swedish krona and the Norwegian krone the exchange rate is for historical reasons given as SEK/100 DKK and SEK/100 NOK

whereas the other rates are directly quoted as SEK/*other rate*. The choice of studying the relationship between SEK and the currencies USD, EUR, GBP and JPY was made because these four currencies are the ones most frequently traded according to the Triennial Central Bank Survey from 2013. According to the same survey the Swedish krona is the 11th most traded currency.[1] The exchange rates SEK/DKK and SEK/NOK were chosen since Denmark and Norway are small countries neighbouring to Sweden with the difference that the Danish krone is bound to the euro[3] whereas the Norwegian krone has no such bound. It could be interesting both to see how alike the model for SEK/EUR and SEK/100 DKK are and how big the difference is between the model for SEK/100 DKK and the one for SEK/100 NOK.

The exchange rates have been obtained from the homepage of *Sveriges riksbank*, which is the central bank of Sweden. Each exchange rate series contain daily observations of the exchange rate over a period of 10 years, from the 4th of January 1993 to the 30th of August 2013. This makes a total of 5190 observations in each series. In the case of SEK/USD one value in the series is missing. This has been replaced by using linear interpolation. After this all the daily observations of the exchange rates have been transformed into daily returns. Information about the mean, standard deviation, maximum, minimum, skewness and kurtosis for the different return series is presented in table (1). Comparing these to the ones in e.g. the work of Roy[15] these values seem reasonable.

Table 1: Information about the daily returns of the different exchange rate series before modelling.

	SEK/USD	SEK/EUR	SEK/JPY	SEK/GBP	SEK/100 DKK	SEK/100 NOK
mean	-0.00001	0.00000	0.00003	-0.00001	0.00001	0.00001
std. dev	0.00772	0.00483	0.00968	0.00643	0.00485	0.00477
maximum	0.04667	0.03671	0.07755	0.03741	0.03607	0.02793
minimum	-0.04667	-0.02927	-0.05918	-0.04185	-0.04527	-0.03186
skewness	0.13663	0.16817	0.39537	-0.04582	0.02676	-0.13616
kurtosis	5.93726	6.86278	7.69107	5.52991	7.77955	5.94280

For notational convenience the exchange rates will in some of the figures in this report only be referred to as "USD", "EUR" etc. when "SEK/USD", "SEK/EUR" are meant. Similarly, when pairs of residuals are considered, combinations like "USD-EUR" represents the residuals coming from the SEK/USD series and the residuals from the SEK/EUR series combined as pairs.

3 Used models and their distributions

All the exchange rates that were presented in the introduction show clear signs of heteroscedasticity. To be able to model these rates a model will be needed which can handle volatility clustering. A group of models that are known for being able to model this are the *Generalized AutoRegressive Conditional Heteroskedasticity* (GARCH) models. There is a large amount of different versions of the standard GARCH model, where each version has been developed to better replicate a certain behaviour of time series which the standard model can not handle. Out of all of these models, only the two GARCH models which were chosen as the final models for the exchange rates are presented: the standard GARCH model and the EGARCH model. These can be found in sections 3.1 and 3.2.

Even after the heteroscedasticity has been removed from the time series with the help of a GARCH model serial dependence might remain. This is solved by introducing an *AutoRegressive-Moving-Average*(ARMA) part to the model. When an explanatory external signal is available, which in the case of the exchange rates will be represented by the STIBOR rate, this signal can be added to the ARMA model as an external regressor. The ARMA model then turns into an *ARMAX* model. The ARMA and the ARMAX model is presented in section 3.1.

Once suitable ARMA(X)-GARCH models have been found for the series of exchange rate returns, bivariate *copula* models will be found for all pairs of residuals resulting from the ARMA(X)-GARCH models. The bivariate copula model is presented in section 3.3.

As will be explained in section 3.1, the innovations in the GARCH model can come from different distributions. Which distribution is assumed to have generated the data series has great impact on how the values of the series may look, especially when it comes to extreme values. The distributions that have been used in the final models for the exchange rates are *the Student's t distribution* and *the skewed Student's t distribution*, which are presented in sections 3.4.1 and 3.4.2. In the process of estimating a copula model for the residuals resulting from the ARMA(X)-GARCH models the choice of which distribution could best represent the residuals lay on either *the logistic model* or the skewed Student's t distribution. For extreme value copulas it is assumed that the observations can be modelled by a *generalized extreme value distribution*. The logistic distribution and the generalized extreme value distribution are presented in sections 3.4.3 and 3.4.4.

3.1 The ARMAX(r, m, v) - GARCH(p, q) model

For the standard Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) model we will assume that $\{Z_t\}$ is a sequence of i.i.d. random variables with mean 0 and variance 1. The process (X_t) will be a GARCH(p,q) process if it is strictly stationary and satisfies

$$X_t = \sigma_t Z_t \quad (1)$$

$$\sigma_t^2 = \omega + \sum_{i=1}^q \alpha_i X_{t-i}^2 + \sum_{j=1}^p \beta_j \sigma_{t-j}^2 \quad (2)$$

where $\omega > 0$, $\alpha_i \geq 0$ for $i = 1, \dots, p$ and $\beta_j \geq 0$ for $j = 1, \dots, q$ to ensure positive variances and $\sum_{i=1}^p \alpha_i + \sum_{j=1}^q \beta_j < 1$ to preserve stability.[17]

The GARCH model can, as mentioned in the introduction of this section, take care of the heteroscedasticity in the series, but the series might still show signs of serial correlation. This can be handled by adding an autoregressive–moving-average (ARMA) part to the GARCH model above. In the ARMA(r,m) model below the X_t is the same X_t that is found in equation (1) and thus it connects the two parts of the model:

$$Y_t = \mu + \sum_{k=1}^r \phi_k Y_{t-k} + X_t + \sum_{l=1}^m \theta_l X_{t-l} \quad (3)$$

Here Y_t in this report represents the log return of the exchange rate series at time t . If external regressors are available the ARMA model can be transformed into an ARMAX(r,m,v) model by adding one sum to equation (3) in the following manner:

$$Y_t = \mu + \sum_{k=1}^r \phi_k Y_{t-k} + X_t + \sum_{l=1}^m \theta_l X_{t-l} + \sum_{u=1}^v \delta_u w_{u,t} \quad (4)$$

[9]. Here $w_{u,t}$ is an external regressor, which in this report will be the STIBOR rate.

3.2 The EGARCH(p,q) model

A problem with financial time series is often that a negative innovation has larger impact on the volatility than a positive innovation. This is mostly observed in negative and positive asset returns and is known as the leverage effect. The standard GARCH model is symmetric, but by using the *exponential GARCH* model, EGARCH, asymmetry can be introduced in the model. The EGARCH(p,q) model is given by

$$\log(\sigma_t^2) = \omega + \sum_{i=1}^q (\alpha_i Z_{t-i} + \gamma_i (|Z_{t-i}| - \mathbb{E}|Z_{t-i}|)) + \sum_{j=1}^p \beta_j \log(\sigma_{t-j}^2) \quad (5)$$

where α_i gives impact to the sign of Z_{t-i} whereas γ_i gives impact to the size of Z_{t-i} . It should be noted that it is the standardized innovation Z_t that is used here and not X_t . [9]

3.3 The copula function

In section 3.1 we assumed that the innovations $\{Z_t\}$ were a sequence of i.i.d. random variables with mean 0 and variance 1. Had the ARMA(X)-GARCH model found been perfect we should indeed only have residuals equalling $\{Z_t\}$ coming from the distribution assumed. The model is in practise, however, not perfect and as mentioned in the introduction there is reason to believe that the exchange rates are affected by the same factors and thus that observing one exchange rate would give a certain measure of information about the behaviour of another exchange rate. This is where copulas come into the picture, which are distribution functions that describe the dependence between random variables.

If $\mathbf{X} = (X_1, \dots, X_d)$ is a random vector with continuous marginal cumulative distribution functions F_1, \dots, F_d , then their joint cumulative distribution function $H_{X_1, \dots, X_d}(x_1, \dots, x_d)$ can be described by

$$H_{X_1, \dots, X_d}(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)) \quad (6)$$

where the copula C is a unique function $C : [0, 1]^d \rightarrow [0, 1]$. For the copula C the following properties have to hold:

1. $C(u_1, \dots, u_d)$ is increasing in each component u_i .
2. $C(1, \dots, 1, u_i, 1, \dots, 1) = u_i$ for all $i \in [1, \dots, d]$, $u_i \in [0, 1]$.
3. For all $(a_1, \dots, a_d), (b_1, \dots, b_d) \in [0, 1]^d$ with $a_i \leq b_i$ it holds that

$$\sum_{i_1=1}^2 \dots \sum_{i_d=1}^2 (-1)^{i_1 + \dots + i_d} C(u_{1i_1}, \dots, u_{di_d}) \geq 0$$

where $u_{j1} = a_j$ and $u_{j2} = b_j$ for all $j \in [1, \dots, d]$. [11]

Using these properties one can derive the Fréchet-Hoeffding bounds which in the bivariate case are $\max(u+v-1, 0) \leq C(u, v) \leq \min(u, v)$, where $\max(u+v-1, 0)$ is known as *Fréchet-Hoeffding lower bound* and $\min(u, v)$ is known as *Fréchet-Hoeffding upper bound*.

A special class of copulas are known as *Archimedean copulas*. An Archimedean copula can be written on the form

$$C(u_1, \dots, u_d) = \Phi^{-1}(\Phi(u_1) + \dots + \Phi(u_d))$$

where Φ is a strictly decreasing function in $[0, 1] \rightarrow [0, +\infty)$ with pseudo-inverse Φ^{-1} . Φ is called the *generator*.

There also exists a class of copulas known as *extreme value copulas*. A copula C_* is called an extreme value copula if there exists a copula C such that

$$C_*(u_1, \dots, u_d) = \lim_{n \rightarrow \infty} C^n(u_1^{1/n}, \dots, u_d^{1/n})$$

for all $i \in [1, \dots, d]$, $u_i \in [0, 1]$. The copula C is said to be in the domain of attraction of C_* . [12]

3.3.1 The t-copula

The Student's t copula belongs to the family of elliptical copulas, for which the dependency structure can be described by the covariance matrix Σ . For these distributions independence is equivalent to zero correlation.

The multivariate Student's t density function is defined as

$$f_{\mathbf{x}}(\mathbf{x}) = \frac{\Gamma(\frac{1}{2}(v+d))}{\Gamma(\frac{1}{2}v)(\pi v)^{d/2}|\Sigma|^{1/2}} \left(1 + \frac{(\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})}{v} \right)^{-\frac{v+d}{2}}$$

Here d is the dimension, v is the degree of freedom, $\boldsymbol{\mu}$ is the mean, Σ is the covariance matrix and Γ is the gamma distribution function. In the bivariate case the t-copula is then defined as

$$C_{v, \Sigma}^t(\mathbf{u}) = t_{v, \Sigma}(t_v^{-1}(u_1), t_v^{-1}(u_2)) \quad (7)$$

where t_v is the one dimensional cumulative distribution function and $t_{v, \Sigma}$ is the bivariate cumulative distribution function. [16]

3.3.2 The Plackett copula

The Plackett copula is defined as

$$C(u, v) = \begin{cases} \frac{(1+(\theta-1)(u+v)) - \sqrt{(1+(\theta-1)(u+v))^2 - 4uv\theta(\theta-1)}}{2(\theta-1)} & \text{if } \theta > 0 \text{ and } \theta \neq 1 \\ uv & \text{if } \theta = 1 \end{cases} \quad (8)$$

If $\theta \rightarrow 0$ then the copula becomes the Fréchet-Hoeffding lower bound, if $\theta \rightarrow \infty$ then the copula becomes the Fréchet-Hoeffding upper bound and for $\theta = 1$ we have the independence copula. In other words the Plackett copula is a comprehensive copula. [12]

3.3.3 The Gumbel copula

The Gumbel copula is both an archimedean copula and an extreme value copula. In the bivariate case it is defined as

$$C(u, v) = \exp\left(-\left((-\log u)^\theta + (-\log v)^\theta\right)^{1/\theta}\right) \quad (9)$$

where $\theta \in [1, \infty)$. When $\theta = 1$ the variables (u, v) are independent and when $\theta \rightarrow \infty$ we obtain perfect positive dependence between the variables. For $\theta > 1$ the Gumbel copula exhibits upper tail dependence, i.e. if u is large then v is also expected to be large. [16]

3.3.4 The Hustler-Reiss copula

The Hustler-Reiss copula is an extreme value copula and in the bivariate case it is defined as

$$G(u, v) = \exp\left(-\hat{u} \Phi\left[\frac{1}{\theta} + \frac{\theta}{2}\log\left(\frac{\hat{u}}{\hat{v}}\right)\right] - \hat{v} \Phi\left[\frac{1}{\theta} + \frac{\theta}{2}\log\left(\frac{\hat{v}}{\hat{u}}\right)\right]\right) \quad (10)$$

where $\hat{u} = -\log(u)$, $\hat{v} = -\log(v)$, $\Phi[\cdot]$ is the standard normal distribution function and $\theta \in (0, \infty)$. As θ approaches zero independence is obtained and when θ goes to infinity complete dependence is obtained. [14]

3.3.5 The Galambos copula

The Galambos copula is an extreme value copula which, for the bivariate case, is defined as

$$C(u, v) = uv \exp\left(\left((-\log u)^{-\theta} + (-\log v)^{-\theta}\right)^{-1/\theta}\right) \quad (11)$$

where $\theta \in (0, \infty)$. When $\theta \rightarrow \infty$ complete dependence is obtained and as $\theta \rightarrow 0$ dependence is obtained. [14]

3.3.6 The Tawn copula

The Tawn copula is an extreme value copula and is in the bivariate case defined as

$$C(u, v) = uv \exp\left(-\theta \frac{\log(u) \log(v)}{\log(uv)}\right) \quad (12)$$

where $\theta \in (0, 1)$. For $\theta = 0$ independence is achieved but it is not possible to obtain complete dependence. [14]

3.4 The distributions of innovations and residuals

The properties of the distribution which Z_t in equation (1) belongs to determines how extreme the values that the random variables can take on can be, and whether or not it is possible for negative (positive) innovations to take on higher absolute values than positive (negative) innovations. The last effect mentioned has to do with the *skewness* of the distribution, where a skewness not equal to zero means that the density function is not symmetric. Skewness is defined as

$$skewness = \frac{\mathbb{E}[(Z - \mu)^3]}{\sigma^3}$$

where μ is the mean and σ the standard deviation. A positive skew indicates that the tail on the left hand side is shorter than the tail on the right hand side, whereas a negative skew would indicate the opposite.

How rare an extreme observation is and what magnitude it has is dependant on the *kurtosis*, which is a measure of how heavy the tails of the distribution are. The normal distribution has a kurtosis of 3 and a distribution with a kurtosis larger than 3 is said to have *excess kurtosis* and will have more mass at the tails of its distribution than what the normal distribution has. The kurtosis is defined as

$$kurtosis = \frac{\mathbb{E}[(Z - \mu)^4]}{\sigma^4}$$

using the same notation as above.[4]

The starting point when creating a GARCH model as in equations (1-2) is often to assume that the innovations, Z_t , come from a normal distribution. The problem with the normal distribution is often that the data series one want to fit have a skewness or excess kurtosis that does not agree with the one of the normal distribution. The problem of excess kurtosis can often be solved by instead assuming that Z_t belongs to a Student's t distribution since this distribution has heavier tails than the normal distribution. To solve the problem of skewness a distribution that allows for skewness has to be found. There are a lot of possible distributions that will not be discussed here since they did not result in any improvements in the models fitted to the exchange rates. In the following the Student's t distribution, the skewed Student's t distribution, the logistic distribution and the generalized extreme value distribution will be presented.

3.4.1 The Student's t distribution

The Student's t distribution is a symmetric distribution with heavier tails than the normal distribution. The distribution can be described by only using the shape parameter, v , but since the innovations of the GARCH model need to be standardized the following representation of the density function also includes the location, α , and scale, β , parameters:

$$f_X(x) = \frac{\Gamma(\frac{v+1}{2})}{\sqrt{\beta v \pi} \Gamma(\frac{v}{2})} \left(1 + \frac{(x - \alpha)^2}{\beta v}\right)^{-\frac{v+1}{2}} \quad \text{where} \quad -\infty < x < \infty$$

Since we want to standardize the distribution we can put the following requirement on the scale parameter:

$$\begin{aligned} Var(x) &= \frac{\beta v}{v - 2} = 1 \\ \Rightarrow \beta &= \frac{v - 2}{v} \end{aligned}$$

Substituting β in this way leads to the standardized Student's t distribution:

$$f_X\left(\frac{x - \mu}{\sigma}\right) = \frac{1}{\sigma} f_Z(z) = \frac{\Gamma(\frac{v+1}{2})}{\sigma \sqrt{(v - 2)\pi} \Gamma(\frac{v}{2})} \left(1 + \frac{z^2}{v - 2}\right)^{-\frac{v+1}{2}} \quad \text{where} \quad -\infty < x < \infty \quad (13)$$

For $v > 4$ the Student's t distribution has zero skewness and an excess kurtosis equal to $6/(v-4)$. As the degrees of freedom, given by the shape parameter v , grows, the Student's t distribution approaches the standard normal distribution.[9]

3.4.2 The skewed Student's t distribution

By introducing inverse scale factors in the positive and negative real half lines, a unimodal and symmetric distribution can be made skew. The density function of a random variable X can be represented as

$$f_X(x|\xi) = \frac{2}{\xi + \xi^{-1}} [f(\xi x)H(-x) + f(\xi^{-1}x)H(x)] \quad \text{where} \quad -\infty < x < \infty \quad (14)$$

Here $\xi \in \mathbb{R}$ is the skew parameter, $H(\cdot)$ is the Heaviside function and $f_X(\cdot)$ is the Student's t distribution as in equation (13). The standardization of the distribution is made using the central moments.[9]

3.4.3 The Logistic distribution

The logistic distribution shows a lot of resemblance with the normal distribution but has higher kurtosis. It has the density function

$$f_X(x) = \frac{\exp\left(-\frac{x-m}{s}\right)}{s\left(1 + \exp\left(-\frac{x-m}{s}\right)\right)} \quad (15)$$

where m is the location parameter and s is the scale parameter. Its mean is m and its variance is $\frac{\pi^2 s^2}{3}$. [14]

3.4.4 The Generalized Extreme Value distribution

The Generalized Extreme Value distribution or the GEV distribution is a distribution to which extreme values converge, regardless of the distribution of the population. Let

$$M_n = \max(X_1, \dots, X_n)$$

where X_1, \dots, X_n are independent random variables coming from the same distribution. Then, if there exists constants $a_n > 0$ and $b_n > 0$ such that

$$\mathbb{P}\left(\frac{M_n - b_n}{a_n} \leq z\right) \rightarrow G(z)$$

as $n \rightarrow \infty$, where G is a non-degenerate distribution function, then G has a distribution of the form

$$G(z) = \exp\left\{-\left(1 + \gamma\left(\frac{z - \mu}{\sigma}\right)\right)^{-1/\gamma}\right\} \quad (16)$$

where $-\infty < \mu < \infty$ is the location parameter, $\sigma > 0$ is the scale parameter and $-\infty < \gamma < \infty$ is the shape parameter. This distribution is known as the Generalized Extreme Value distribution and is defined on $\{z : 1 + \gamma(z - \mu)/\sigma > 0\}$. [2]

4 Validation of the models

In the process of finding a suitable ARMA-GARCH model for the data a lot of tests can be preformed. To get an idea as to which kind of model could fit the data it is useful to look at the *autocorrelation function* and the *partial autocorrelation function* for the returns and the squared returns of the data series to discover serial dependence and serial dependence in the variance,

i.e. conditional heteroscedasticity or ARCH effects. Another way to test for ARCH effects is to use the *Lagrange multiplier test* of Engle.

After a model has been fitted to the data it is important to check if the model suits the data, that is to check if it seems probable that the residuals would look the way they do if the data truly came from the assumed model. This can be done in a number of ways. One way is to use the *sign bias test*, which uses the residuals to tests if the volatility is affected differently by positive and negative shocks and how well the model handles this. It is also possible to visually estimate the goodness of fit by looking for instance at the autocorrelation function and partial autocorrelation function, but this time for the standardized residuals and squared standardized residuals. Another way is to compare the histogram of the standardized residuals with the assumed probability density function of the chosen distribution, or to look at the *quantile-quantile plot*.

To compare different fitted models to each other it is possible to use visual methods for the comparison like the ones just mentioned, but this can be hard if the models are similar. It is also possible to use different *information criteria* or, if one model is just an extended version of the other model, one can use the *likelihood ratio test* to compare which of the models that gives the best fit to the data without adding an unnecessarily large amount of parameters.

There are several other tests that can be performed on the standardized residuals after the model has been fitted. One of them is the *GMM Orthogonality Test* which compares the moments of the standardized residuals with those expected, had they come from the chosen distribution. The *Hong and Li test* tests whether the conditional density fits the underlying data. By simulating a new data series from the chosen model and calculating the exceedances of Value at Risk the *Kupiec* and *Christoffersen* tests can be preformed.

For copulas it is important to know how large the dependence between two data series is and if this changes with time. To be able to measure dependence that is not linear *Kendall's τ* and *Spearman's ρ* is calculated. To check the goodness of fit of different types of copulas a goodness-of-fit test is used. The construction of the goodness-of-fit test depends on if the copula is an ordinary copula or an extreme value copula. The tests mentioned are described below.

4.1 The autocorrelation function and the partial autocorrelation function

The auto-covariance function for y_t is defined as $r_y(k) \equiv C[y_t, y_{t-k}]$ where k is the lag. The autocorrelation function for y_t is then defined as

$$\rho_y(k) = \frac{r_y(k)}{r_y(0)}, \quad k = 0, 1, 2, \dots \quad (17)$$

and hence $|\rho_y(k)| \leq 1$. For a white noise process with zero mean and constant variance the random variables are uncorrelated, why $r_y(k)$ will take on the value of σ^2 when $k = 0$ and zero otherwise, leading to $\rho_y(k) = 1$ when $k = 0$ and zero otherwise. In the following it will be explained how this can be used to find a suitable model order for a time series.

The biased estimator of the auto-covariance is

$$\hat{r}_y(k) = \frac{1}{N} \sum_{t=k+1}^N (y_t - \hat{m}_y)(y_{t-k} - \hat{m}_y)^T$$

for $0 \leq k \leq N - 1$, where $\hat{m}_y = \frac{1}{N} \sum_{t=1}^N y_t$ is the mean of the process. If e_t is the realization of a white noise process with zero mean and variance σ_e^2 , with $\hat{r}_e(k)$ estimated like above and the estimated autocorrelation function, $\hat{\rho}_e(k)$, is estimated using equation (17), then the following holds:

$$\begin{aligned}\mathbb{E}[\hat{\rho}_e(k)] &= 0 \\ \mathbb{V}[\hat{\rho}_e(k)] &= \frac{1}{N}\end{aligned}$$

for all $k \neq 0$, with N being the number of observations. For $k > 0$, $\hat{\rho}_e(k)$ will be asymptotically normal distributed.

Let us define a moving average process, *MA process*, as:

$$y_t = e_t + c_1 e_{t-1} + \dots + c_m e_{t-m}$$

where e_t is a zero mean white noise process with variance σ_e^2 . The autocorrelation function for the MA(m) process will be

$$r_y(k) = \begin{cases} \sigma_e^2(c_k + c_1 c_{k+1} + \dots + c_{m-l} c_m) & \text{if } |k| \leq m \\ 0 & \text{if } |k| > m \end{cases}$$

This means that the autocorrelation function should be zero for all lags greater than m . [10] This can be used in the process of finding a model for the data series. If the data was generated by an MA(m) process the autocorrelation function could be plotted and the model order can be found by finding when the autocorrelation function turns zero.

Similarly, the autoregressive process, *AR process*, is defined as

$$y_t = \phi_{k,1} y_{t-1} + \dots + \phi_{k,k} y_{t-k} + e_t$$

where, again, e_t is a white noise process with zero mean and variance σ_e^2 that is uncorrelated with y_{t-l} for $l > 0$ and $\phi_{k,l}$ is the l :th AR coefficient of an AR(k) model. The covariance function is

$$\mathbb{E}[e_t y_{t-k}] = r_y(k) - \phi_1 r_y(k-1) - \dots + \phi_p r_y(k-p) = \sigma_e^2 \delta_K(k)$$

where $\delta_K(k)$ is the Kronecker delta function and where the last equality comes from e_t being uncorrelated with y_{t-l} for $l > 0$. Re-writing this it is possible to form the Yule-Walker equations as

$$\begin{bmatrix} 1 & \rho_y(1) & \dots & \rho_y(k-1) \\ \rho_y(-1) & 1 & \dots & \rho_y(k-2) \\ \vdots & \vdots & \ddots & \vdots \\ \rho_y(k-1) & \rho_y(k-2) & \dots & 1 \end{bmatrix} \begin{bmatrix} \phi_{k,1} \\ \phi_{k,2} \\ \vdots \\ \phi_{k,k} \end{bmatrix} = \begin{bmatrix} \rho_y(1) \\ \rho_y(2) \\ \vdots \\ \rho_y(k) \end{bmatrix}$$

By solving this equation system one can obtain $\phi_{k,k}$ which is the *partial auto-correlation function*, PACF. For an AR(p) process, $\phi_{k,k}$ will be equal to zero for all $k > p$. Just like for the MA(m) process it is possible to determine the order of the AR process by determining for which k the partial autocorrelation function turns zero. If $\hat{\phi}_{k,k}$ is estimated using the equation above the following holds:

$$\begin{aligned}\mathbb{E}[\hat{\phi}_{k,k}] &= 0 \\ \mathbb{V}[\hat{\phi}_{k,k}] &= \frac{1}{N}\end{aligned}$$

for $k > p$ and for $k > p$, $\hat{\phi}_{k,k}$ will be asymptotically normal distributed.[10]

It has now been explained how the ACF and PACF can help with finding the model order of an MA and AR process, respectively. Further, if the data series has been fitted to a model that well describes the behaviour of the data, then the (standardized) residuals, Z_t , should behave like white noise and be asymptotically normal distributed. If the autocorrelation function and partial autocorrelation function is plotted, every lag, except for $\rho(0)$ which is equal to one, should be within the approximative confidence interval of $\pm\lambda_{\alpha/2}/\sqrt{N}$. [10] This is a good way to visually find out if there is still serial dependence in the residuals. To look for serial dependence in the variance, the (standardized) residuals can be squared or their absolute value can be used and then the ACF and PACF are calculated for these values. The correspondence to the parameters in the GARCH(p,q) model is however not as clear as the correspondence to the ARMA(r,m) parameters.

4.2 The Lagrange Multiplier Test of Engle

If the data series has no significant correlations it is possible to perform the Lagrange Multiplier Test of Engle directly. If not, we denote the residuals

$$\varepsilon_t = y_t - \mu - \sum_{k=1}^r \phi_k y_{t-k} - \sum_{l=1}^m \theta_l X_{t-l} \quad (18)$$

i.e. ε is the estimated value of X_t from equation (3). If the conditional mean model was defined properly the residuals ε_t should now have mean zero and be uncorrelated. They can however still be serially dependent. The Lagrange Multiplier Test of Engle checks this by testing the alternative hypothesis

$$H_1 : \quad \varepsilon_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \dots + \alpha_m \varepsilon_{t-m}^2 + e_t \quad \text{for} \quad t = m + 1, \dots, T$$

against the null hypothesis

$$H_0 : \quad \alpha_1 = \dots = \alpha_m = 0$$

where e_t is a white noise process, m is the number of lags tested and T is the sample size. The test statistics is the F statistics for the regression on the squared residuals ε^2 . Under the null hypothesis the F statistics is asymptotically chi-squared distributed with m degrees of freedom. The null hypothesis is rejected if $F > \chi_m^2(\alpha)$ with $\chi_m^2(\alpha)$ being the upper $100(1-\alpha)$ th percentile of χ_m^2 or, alternatively, if the p-value of F is less than α .

The F statistics is obtained by letting $SSR_0 = \sum_{t=m+1}^T (\varepsilon_t^2 - \bar{\varepsilon}^2)^2$, where $\bar{\varepsilon}^2$ is the mean of the series $\{\varepsilon^2\}$, and $SSR_1 = \sum_{t=m+1}^T \hat{e}_t^2$, where \hat{e}_t is the least squares residual of the linear regression above. At last

$$F = \frac{(SSR_0 - SSR_1)/m}{SSR_1/(T - 2m - 1)}$$

is calculated to conclude if H_0 should be rejected or not.[17]

4.3 Q-Statistics

This is the test statistics of Ljung and Box's test that checks for autocorrelation at multiple lags jointly. The null hypothesis is $H_0 : \rho_1 = \rho_2 = \dots = \rho_m = 0$ with ρ denoting the autocorrelation.

The test statistics is

$$Q(m) = N(N + 2) \sum_{i=1}^m \frac{\hat{\rho}_i^2}{N - i}$$

where N is the number of observations and m is the number of lags being tested.[17] This test is performed both on the standardized residuals and on the squared standardized residuals to make sure that the residuals have been reduced to white noise.

4.4 Likelihood Ratio Test

To perform a likelihood ratio test one of the two models that are to be compared has to be a restricted version of the other model. This means that there are m different restrictions on the value of the parameter vector $\boldsymbol{\theta}$ for the restricted model. The restricted model is called the null model whereas the unrestricted model is called the alternative model. When the models are fitted to the data the log-likelihood functions are maximized under the restricted and unrestricted model assumptions, respectively. If we let $l(\boldsymbol{\theta}_0)$ denote the maximized log-likelihood for the null model and $l(\boldsymbol{\theta}_1)$ denote the maximized log-likelihood for the alternative model, then

$$2(l(\boldsymbol{\theta}_1) - l(\boldsymbol{\theta}_0)) \sim \chi^2(m)$$

The null model is rejected in favour of the alternative model if $2(l(\boldsymbol{\theta}_1) - l(\boldsymbol{\theta}_0)) > c_\alpha$ where c_α is the $(1 - \alpha)$ quantile of the $\chi^2(m)$ distribution. [2]

4.5 Information Criteria

For models which are not nested the likelihood ratio test described above can not be used. To be able to choose between different models one can instead use information criteria for guidance. In the *rugarch* package four different types of information criteria are used. Common for all of them is that they weight the model's log likelihood value together with the number of parameters estimated. This means that models which are over-fitted and naturally have a higher log likelihood value get a penalty for the number of parameters that are estimated. This in turn makes the effects of the model comparable to models with less (more) parameters. The four criteria are *Akaike information criterion* (AIC), *Bayesian information criterion* (BIC), *Hannan-Quinn information criterion* (HQIC) and *Shibata information criterion* (SIC). These are defined as

$$\begin{aligned} AIC &= -\frac{2LL}{N} + \frac{2m}{N} \\ BIC &= -\frac{2LL}{N} + \frac{m \log(N)}{N} \\ HQIC &= -\frac{2LL}{N} + \frac{2m \log(\log(N))}{N} \\ SIC &= -\frac{2LL}{N} + \log\left(\frac{N + 2m}{N}\right) \end{aligned}$$

where LL is the log likelihood for the model, N is the number of observations and m is the number of parameters estimated.[9]

4.6 Sign Bias Test

A correctly specified GARCH model should be able to handle the changes in volatility. For some products, with stocks being among to most well known, the volatility reacts asymmetrically to the sign of the shock or innovation, and in the case with the stock, causes a larger increase in the volatility if the shock is negative than if it is positive. The sign bias test of Engle and Ng uses the squared standardized residuals to check for leverage effect by regressing the residuals on lagged negative and positive shocks:

$$\hat{z}_t^2 = c_0 + c_1 \mathbb{1}(\hat{z}_{t-1} < 0) + c_2 \mathbb{1}(\hat{z}_{t-1} < 0) \hat{z}_{t-1} + c_3 \mathbb{1}(\hat{z}_{t-1} \geq 0) \hat{z}_{t-1} + u_t$$

where $\mathbb{1}$ is the indicator function and \hat{z}_t are the estimated standardized residuals. The null hypothesis is that $c_i = 0$ for all i and that jointly $c_1 = c_2 = c_3$. [9]

4.7 The GMM Orthogonality Test

This test is based on the generalized method of moments, and checks if the standardized residuals fulfil a number of moment conditions which would hold if the model was correctly specified. The null hypothesis is that the model is correctly specified and the alternative hypothesis is that the moments are not equal to zero. In the *rugarch* package the following conditions are tested:

$$M_1 : \quad \mathbb{E}[z_t] \quad \quad \quad = 0 \quad \quad \quad (19)$$

$$M_2 : \quad \mathbb{E}[z_t^2 - 1] \quad \quad \quad = 0 \quad \quad \quad (20)$$

$$M_3 : \quad \mathbb{E}[z_t^3] \quad \quad \quad = 0 \quad \quad \quad (21)$$

$$M_4 : \quad \mathbb{E}[z_t^4 - 3] \quad \quad \quad = 0 \quad \quad \quad (22)$$

$$Q_2 : \quad \mathbb{E}[(z_t^2 - 1)(z_{t-j}^2 - 1)] \quad \quad \quad = 0 \quad \quad \quad (23)$$

$$Q_3 : \quad \mathbb{E}[(z_t^3)(z_{t-j}^3)] \quad \quad \quad = 0 \quad \quad \quad (24)$$

$$Q_4 : \quad \mathbb{E}[(z_t^4 - 3)(z_{t-j}^4 - 3)] \quad \quad \quad = 0 \quad \quad \quad (25)$$

where j is the lag, which in the *rugarch* package is set to 4. The first four moments tests the mean, variance, skewness and kurtosis whereas Q_2, Q_3 and Q_4 checks the joint conditional moment conditions for variance, skewness and kurtosis and are $\chi^2(j)$ distributed. A joint Wald test can also be preformed to test the moments jointly and is then $\chi^2(3j + 4)$ distributed. [9]

4.8 The Non-Parametric Density Test of Hong and Li

The test of Hong and Li jointly tests the hypothesis of i.i.d. and $\mathcal{U}(0,1)$ of data transformed by the probability integral transformation (PIT). The transformed standardized residuals are used for checking for misspecification in the conditional moments. Both individual tests and general misspecification of the model is tested and are both distributed as $\mathcal{N}(0,1)$ if the null hypothesis is correct. Only upper tail critical values are used since negative values only occur under the null hypothesis.

A kernel estimator $\hat{g}_j(x_1, x_2)$ for the joint density $g_j(x_1, x_2)$ of the pair $\{x_t, x_{t-j}\}$ is compared to the product of two $\mathcal{U}(0,1)$ distributions. Here j is the lag order and $j > 0$. If the sample size is n , the joint density is estimated as:

$$\hat{g}_j(x_1, x_2) = (n - j)^{-1} \sum_{t=j+1}^n K_h(x_1, \hat{X}_t) K_h(x_2, \hat{X}_{t-j})$$

with $\hat{X}_t = X_t(\hat{\theta})$, and $\hat{\theta}$ being a \sqrt{n} consistent estimator of θ_0 . K_h is the boundary modified kernel:

$$K_h(x, y) = \begin{cases} \frac{k(\frac{x-y}{h})h^{-1}}{\int_{-x/h}^1 k(u)du} & \text{if } x \in [0, h) \\ k(\frac{x-y}{h})h^{-1} & \text{if } x \in [h, 1-h) \\ \frac{k(\frac{x-y}{h})h^{-1}}{\int_{-1}^{(1-x)/h} k(u)du} & \text{if } x \in [1-h, 1) \end{cases}$$

where $h = h(n)$ is a bandwidth such that $h \rightarrow 0$ as $n \rightarrow \infty$ and the kernel $k(\cdot)$ is a symmetric probability density that in the *rugarch* package is implemented as

$$k(u) = \frac{15}{16}(1-u^2)^2 \mathbb{1}(|u| \leq 1)$$

where $\mathbb{1}$ is the indicator function. The test statistics is defined as

$$\hat{W}(p) = p^{-1/2} \sum_{j=1}^p \left[(n-j)h\hat{M}(j) - A_h^0 \right] / \sqrt{V_0} \quad (26)$$

where

$$\hat{M}(j) = \int_0^1 \int_0^1 [\hat{g}_j(x_1, x_2) - 1]^2 dx_1 dx_2$$

and the centering and scaling factors A_h^0 and V_0 are defined as

$$A_h^0 = \left[(h^{-1} - 2) \int_{-1}^1 k^2(u) du + 2 \int_0^1 \int_{-1}^b k_b^2(u) du db \right]^2 - 1$$

$$V_0 = 2 \left[\int_{-1}^1 \left[\int_{-1}^1 k(u+v)k(v) dv \right]^2 du \right]^2$$

where

$$k_b(\cdot) = k(\cdot) / \int_{-1}^b k(v) dv.$$

To be able to investigate the possible causes of misspecification when a model is rejected the following test statistics is also used:

$$M(m, l) = \frac{\left[\sum_{j=1}^{n-1} w^2(j/p)(n-j)\hat{\rho}_{ml}^2(j) - \sum_{j=1}^{n-1} w^2(j/p) \right]}{\left[2 \sum_{j=1}^{n-2} w^4(j/p) \right]^{1/2}} \quad (27)$$

where $\hat{\rho}_{ml}(j)$ is the sample cross-correlation between \hat{X}_t^m and $\hat{X}_{t-|j|}^l$ and $w(\cdot)$ is a weighting function of lag order j , in the *rugarch* package implemented as the Bartlett kernel.[9]

4.9 Kupiec's and Christoffersen's test

In risk management *Vaule at Risk*, VaR, is often used to calculate the amount of money a company would loose at the most in $(1 - \alpha)\%$ of the cases. In general, however, it calculates the α quantile, and hence the value is determined by the assumed distribution and the estimated

variance. Since the observed values should exceed VaR ($\alpha \times 100$) percent of the time it is possible to calculate the expected number of exceedances and compare it to the actual number of exceedances. If the model suits the data the number of observed exceedances should not deviate too much from the expected number of exceedances and this is where the tests of Kupiec and Christoffersen come into play.

In the following, the indicator function I_{t+1} will be defined as

$$I_{t+1}(\alpha) = \begin{cases} 1 & \text{if } y_{t,t+1} \leq -VaR_t(\alpha) \\ 0 & \text{if } y_{t,t+1} > -VaR_t(\alpha) \end{cases}$$

where $y_{t,t+1}$ is the observation over a fixed time interval and $VaR_t(\alpha)$ is the value at risk at time t for the α percentile. The indicator function will form a sequence, e.g. (0001010011), indicating when the VaR was violated. This sequence should be iid as Bernoulli random variables with probability α .

Kupiec's test was one of the first tests proposed for VaR backtesting and is only concerned with whether or not the number of times the VaR is violated differs too much from the expected α percent, i.e. unconditional coverage. If it does, this would indicate that the model is misspecified. Hence the null hypothesis is that the observed number of violations equals the expected number of violations. Letting $I(\alpha) = \sum_{t=1}^T I_t(\alpha)$ and $\hat{\alpha} = I(\alpha)/T$ where T is the number of observations, Kupiec's test statistics is

$$LR_{uc} = 2 \log \left(\left(\frac{1 - \hat{\alpha}}{1 - \alpha} \right)^{T - I(\alpha)} \left(\frac{\hat{\alpha}}{\alpha} \right)^{I(\alpha)} \right) \sim \chi^2(1) \quad (28)$$

Christoffersen's test does not only take the unconditional coverage into account, but also the conditional coverage, i.e. if the values of the indicator function are independent, since this is another way in which a model can be misspecified. The test statistics for the unconditional coverage is still given by equation (28). When it comes to the unconditional coverage, n_{ij} will denote the number of days state i on the first day was followed by state j on the next, and i and j takes on the values 1 or 0, where 1 represents a violation of VaR and 0 represents no violation. Further, π_{ij} will denote the probability that state j occurs if the state of the previous day was i , whereas the observed probabilities are

$$\hat{\pi}_{01} = \frac{n_{01}}{n_{00} + n_{01}}, \quad \hat{\pi}_{11} = \frac{n_{11}}{n_{10} + n_{11}}, \quad \hat{\pi}_2 = \frac{n_{01} + n_{11}}{n_{00} + n_{10} + n_{01} + n_{11}}$$

The test statistics for the conditional coverage is then

$$LR_{ind} = -2 \log((1 - \hat{\pi}_2)^{n_{00} + n_{10}} \hat{\pi}_2^{n_{01} + n_{11}}) + 2 \log((1 - \hat{\pi}_{01})^{n_{00}} \hat{\pi}_{01}^{n_{01}} (1 - \hat{\pi}_{11})^{n_{10}} \hat{\pi}_{11}^{n_{11}})$$

Under the null hypothesis that the unconditional coverage is correct and that the exceedances are independent the test statistics for the conditional coverage will be

$$LR_{cc} = LR_{uc} + LR_{ind} \quad (29)$$

with LR_{cc} being $\chi^2(2)$ distributed.[4]

Both Kupiec's and Christoffersen's tests should be performed out of sample to give a valid result of the hypothesis test, why in this thesis the 500 data points saved for model validation are used for this purpose.

4.10 Measures of dependence

In copula models the dependence between the data series it should represent is crucial. The common way of thinking of dependence is by correlation, but this is a measure of *linear* dependence. If one wants to account for other types of dependence other than the linear dependence, Kendall's τ and Spearman's ρ_S can be used. In this report these measures will be used to see how strong the dependence between different data series is and how this changes when going from daily observations to monthly and quarterly extreme values. Below the Pearson's correlation, which is the "standard" correlation, Kendall's τ and Spearman's ρ_S are stated together with their relationship to each other.

Pearson's correlation. Pearson's correlation is defined as

$$\rho(\mathbf{X}, \mathbf{Y}) = \frac{Cov(\mathbf{X}, \mathbf{Y})}{\sqrt{Var(\mathbf{X})}\sqrt{Var(\mathbf{Y})}}$$

where \mathbf{X} and \mathbf{Y} are vectors, $Cov(\mathbf{X}, \mathbf{Y}) = \mathbb{E}[(\mathbf{X} - \mathbb{E}[\mathbf{X}])(\mathbf{Y} - \mathbb{E}[\mathbf{Y}])]$ is the covariance between the vectors and $Var(\mathbf{X}) = Cov(\mathbf{X}, \mathbf{X})$ is the variance.

Kendall's τ . Kendall's τ is defined as

$$\begin{aligned} \tau(\mathbf{X}, \mathbf{Y}) &= \mathbb{P}[(\mathbf{X} - \tilde{\mathbf{X}})(\mathbf{Y} - \tilde{\mathbf{Y}}) > 0] - \mathbb{P}[(\mathbf{X} - \tilde{\mathbf{X}})(\mathbf{Y} - \tilde{\mathbf{Y}}) < 0] = \\ &= 4 \int \int C(u, v) dC(u, v) - 1 = 4\mathbb{E}[C(u, v)] - 1 \end{aligned}$$

where $\mathbb{P}()$ denotes the probability and both (\mathbf{X}, \mathbf{Y}) and $(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}})$ have the same copula C . The relationship between Kendall's τ and Pearson's ρ is

$$\tau(\mathbf{X}, \mathbf{Y}) = Corr[\text{Sign}(\mathbf{X} - \tilde{\mathbf{X}}), \text{Sign}(\mathbf{Y} - \tilde{\mathbf{Y}})] \quad \text{where } \text{Sign}(\mathbf{X}) = \begin{cases} 1 & \text{if } X > 0 \\ 0 & \text{if } X = 0 \\ -1 & \text{if } X < 0 \end{cases}$$

Spearman's ρ_S . Spearman's ρ_S is defined as

$$\rho_S(\mathbf{X}, \mathbf{Y}) = 3(\mathbb{P}[(\mathbf{X} - \tilde{\mathbf{X}})(\mathbf{Y} - \mathbf{Y}') > 0] - \mathbb{P}[(\mathbf{X} - \tilde{\mathbf{X}})(\mathbf{Y} - \mathbf{Y}') < 0]) = \diamond$$

with definitions as above. Spearman's ρ_S is related to Pearson's ρ by

$$\begin{aligned} \diamond &= 12 \int \int uv dC(u, v) - 3 = 12\mathbb{E}[uv] - 3 = \frac{\mathbb{E}[uv] - 1/4}{1/12} = \\ &= \rho(F_1(\mathbf{X}), F_2(\mathbf{Y})) \end{aligned}$$

where again C is the copula.[12]

4.11 Goodness-of-fit test for copulas

Goodness-of-fit test for copulas based on the empirical copula

The goodness-of-fit test used for the daily residual series tests the null hypothesis $H_0 : C \in C_0$, where C is the copula $C(F_X(x), F_Y(y))$ and C_0 is a specific parametric family of copulas. The test is based on ranks. The empirical distribution, defined as

$$C_n(u, v) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(U_i \leq u, V_i \leq v)$$

where (\mathbf{U}, \mathbf{V}) are pseudo-observations, is used to compare the "distance" between the empirical distribution C_n and the parametric estimation C_{θ_n} of C under the null hypothesis. θ_n is an estimate of θ where the estimation is based on the pseudo-observations. The test statistics is

$$S_n = \int_{[0,1]^2} n(C_n(u, v) - C_{\theta_n}(u, v))^2 dC_n(u, v) \quad (30)$$

In practice this is calculated by using a parametric bootstrap method in the following way: [8]

1. Calculate the empirical distribution C_n from the sample and estimate θ using an estimation method of choice.
2. (a) Generate a random sample from C_{θ_n} and compute the pseudo-observations $(\mathbf{U}_1^*, \dots, \mathbf{U}_m^*)$.
(b) Approximate C_{θ_n} by

$$B_m^*(\mathbf{u}) = \frac{1}{m} \sum_{i=1}^m \mathbf{1}(\mathbf{U}_i^* \leq \mathbf{u}), \quad \mathbf{u} \in [0, 1]^d$$

- (c) Approximate S_n by

$$S_n = \sum_{i=1}^n (C_n(\mathbf{U}_i) - B_m^*(\mathbf{U}_i))^2$$

3. Repeat for $k \in [1, \dots, N]$:
 - (a) Generate a random sample $(\mathbf{Y}_{1,k}^*, \dots, \mathbf{Y}_{n,k}^*)$ from C_{θ_n} and compute their rank vectors $(\mathbf{R}_{1,k}^*, \dots, \mathbf{R}_{n,k}^*)$
 - (b) Compute the pseudo-observations $\mathbf{U}_{i,k}^* = \mathbf{R}_{i,k}^*/(n+1)$ for $i \in [1, \dots, k]$. Then calculate

$$C_{n,k}^*(\mathbf{u}) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(\mathbf{U}_{i,k}^* \leq \mathbf{u}), \quad \mathbf{u} \in [0, 1]^d$$

and estimate θ .

- (c) Follow the procedure in step 2b and c to estimate $S_{n,k}^*$

4. The approximate p-value is then given by $\frac{1}{N} \sum_{k=1}^N \mathbf{1}(S_{n,k}^* > S_n)$.

In the computations for this report maximum pseudo-likelihood was used to estimate θ_n .

Goodness-of-fit for extreme value copulas using the dependence function

In this test for extreme value copulas a non-parametric estimator is compared to a parametric estimator to test the goodness-of-fit of the parametric model. It can be shown that a bivariate extreme value copula C can be described by

$$C_\theta(u, v) = \exp \left(\log(uv) A \left(\frac{\log(v)}{\log(uv)} \right) \right) \quad (31)$$

where $u, v \in (0, 1)$ and the function $A : [0, 1] \rightarrow [1/2, 1]$ is Pickands dependence function. For all $w \in [0, 1]$ we have that $\max(w, 1 - w) \leq A(w) \leq 1$ where $A(w) = 1$ would mean total independence between u and v and $A(w) = \max(w, 1 - w)$ would mean total dependence between u and v .

It is also possible to make a non-parametric estimation of A . Let $\xi_i(0) = -\log(U_i)$ and $\xi_i(1) = -\log(V_i)$ where U_i and V_i are pseudo-observations and

$$\xi_i(w) = \min \left(\frac{-\log(U_i)}{1 - w}, \frac{-\log(V_i)}{w} \right)$$

for all $w \in (0, 1)$. Then a non-parametric estimator of A_n is

$$A_n^{CFG}(w) = \exp \left\{ -\gamma - \frac{1}{n} \sum_{i=1}^n \log(\xi_i(w)) \right\} \quad (32)$$

where γ is Euler's constant and $A_n^{CFG}(w)$ is known as the CFG estimator after Caperaa, Fougères and Genest.

In the goodness-of-fit test for the extreme value copulas the pseudo-observations (\mathbf{U}, \mathbf{V}) will be used to test the null hypothesis

$$H_0 : A \in \mathcal{A} = \{A_\theta : \theta \in \mathcal{O}\}$$

where \mathcal{A} is a specific parametric class and \mathcal{O} is an open subset of \mathbb{R}^p for some integer p . This will be tested by comparing the non-parametric estimator A_n of A to the parametric estimator A_{θ_n} using the Cramér-von-Mises statistic

$$S_n = \int_0^1 n |A_n(w) - A_{\theta_n}(w)|^2 dw \quad (33)$$

Again a bootstrap procedure is used for the calculations of the goodness-of-fit test, this time with the following procedure: [7]

1. Compute A_n from the pseudo-observations (\mathbf{U}, \mathbf{V}) and estimate θ using a rank-based estimator.
2. Compute the test statistics in equation (33).
3. For $k = 1, \dots, N$, where N is some large integer, do:
 - (a) Generate a sample $(\mathbf{X}_k, \mathbf{Y}_k)$ from C_{θ_n} and compute the pseudo-observations $(\mathbf{U}_k, \mathbf{V}_k)$.
 - (b) Derive A_{nk} and θ_{nk} from $(\mathbf{U}_k, \mathbf{V}_k)$.
 - (c) Insert A_{nk} and θ_{nk} into equation (33) and compute S_{nk} .

4. The p-value of the test is given by

$$\frac{1}{N} \sum_{k=1}^N \mathbb{1}(S_{nk} \geq S_n)$$

In this report, inversion of Kendall's τ was used to estimate θ_{nk} .

5 Method of finding suitable models

As already mentioned, each of the six exchange rates consist of ten years of daily observations, resulting in 5190 observations per exchange rate. Since data will be needed to validate the models the full series is not used for modelling. Instead 500 data points are saved for validation and the models are thus based on the values of the first 4690 points.

As a first step the log returns of each series is created and the ACF and the PACF of the returns and the squared returns are plotted to get an idea of which type of model the series might have come from. The Lagrange Multiplier Test is also performed to make sure that there is serial dependence in the variance and thus that there is a need for a GARCH model. By varying the distribution of Z_t and the number of parameters in the ARMA(r, m) model as well as shifting between a standard GARCH and an EGARCH model the information criteria for all these combinations can be calculated and compared. By comparing these values it is often evident which type of model fits the return series the best. It is also checked if the model can be improved by adding the STIBOR rate as an external regressor.

The best models are examined more closely and it is checked if all estimated parameters are significant and if any of the lags can be removed while still maintaining a good model. The sign bias test is also used for checking if the model reflects the nature of the sign changes in the series. If the sign bias test does not signal that there is a problem with the sign changes for the standard GARCH model there is no need to consider the EGARCH model. The ACF and the PACF of the standardized residuals and the squared standardized residuals of the models should be looked upon to make sure that there is no longer any sign of serial correlation and the quantile-quantile-plot should be considered to make sure that it seems reasonable that the standardized residuals indeed came from the distribution that was assumed by the model. Once good models have been found they are compared using the likelihood ratio test or information criteria to conclude which model should be used. In cases when the value of the two information criteria tests are the same the rule of keeping the model as simple as possible while retaining a sufficiently good result is used.

The goodness of fit of the chosen model is further examined by using tests like Kupiec's and Christoffersen's test, the GMM orthogonality test and the test of Hong and Li. If these tests are not passed by the model it is a sign of model misspecification and one has to take a step back and try to find another model that suits the data better.

When suitable ARMA-GARCH models have been found, the residuals resulting from these models are next in line to be examined. First suitable distributions for the univariate residuals are to be found. This is done by fitting different distributions that graphically seems to be suitable for the series to the data. The fitted distributions are then evaluated with both a qq-plot and a χ^2 -test to find the best fit. For the monthly and quarterly maxima Generalized Extreme Value distributions, GEV, are considered since, given a large enough sample of extreme values, these

will converge in distribution to the GEV distribution.[2]

Due to the long computational time of the goodness-of-fit tests the residuals are next fitted pairwise to different probable copulas, using the distributions for the individual series as marginal distributions. This means that 15 different pairs of residuals are fitted to various copulas. For the fitted copulas with the lowest information criteria value a goodness of fit test is preformed. If the copula should be rejected by the test, other copulas are examined by the goodness-of-fit test in the hopes of finding a suitable copula.

6 Result

In this section the models that were chosen for the data set will be presented, first the ARMA-GARCH model for the daily returns and then the copula models for the residuals.

6.1 Model choice - ARMA-GARCH models for the daily returns

Before transformation the original data series consists of the exchange rate between the Swedish krona and another currency. A series like this might look as in figure (1) where the original series of the exchange rate SEK/EUR is plotted. Before the modelling can begin the series has to be transformed into daily log returns. As an example, figure (2) shows the daily log returns for SEK/EUR. Next the autocorrelation function and the partial autocorrelation function for the daily return series and the squared daily return series is studied to see if there are signs of heteroscedasticity and to see what type of model might be suitable for the series. Again taking the SEK/EUR as an example, the partial autocorrelation function for the daily return series and for the squared daily return series can be seen in figures (3) and (4). From figure (4) it can be seen that there are clear signs of heteroscedasticity. The same conclusion is reached by using the Lagrange multiplier test of Engle, in which the null hypothesis can be rejected for all of the exchange rates.




Figure 1: The original data series of the exchange rate SEK/EUR.




Figure 2: The series of SEK/EUR after it has been transformed into daily log returns.

After extensive testing of different combinations of models, parameters and distributions, suitable ARMA-GARCH models are found for the exchange rate series. The parameter values

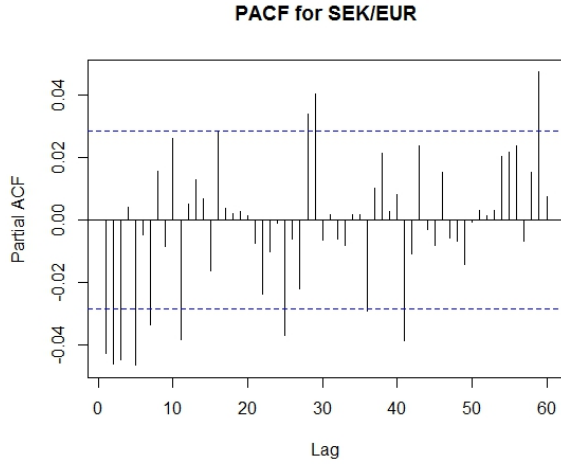


Figure 3: The partial autocorrelation function for the return series SEK/EUR together with the confidence interval where $\alpha = 0.05$.




Figure 4: The partial autocorrelation function for the squared return series SEK/EUR together with the confidence interval where $\alpha = 0.05$.

for the different models can be seen in table (2). As can be seen, the SEK/EUR, SEK/JPY, SEK/100 DKK and SEK/100 NOK are all versions of AR-EGARCH models, SEK/GBP is an ARX-EGARCH model which uses the STIBOR rate as external input and SEK/USD is a GARCH model. All of the models are assumed to have innovations coming from a Student's t distribution except the model for SEK/JPY in which it is assumed that the innovations come from a skewed Student's t distribution.

Table 2: Estimated parameter values for the chosen ARMA(X)-(E)GARCH models.

	ϕ_1	ϕ_2	ϕ_3	δ_1	ω	α_1	β_1	γ_1	df	$skew$
SEK/USD	-	-	-	-	0.0000	0.052	0.9374	-	12.95	-
SEK/EUR	-0.0392	-0.0578	-0.0528	-	-0.0954	0.0427	0.9913	0.1371	9.76	-
SEK/JPY	0	-0.0342	-0.0315	-	-0.1538	0.0485	0.9839	0.1436	10.24	1.11
SEK/GBP	0	-0.042	-0.0353	0.0464	-0.1346	0.0248	0.9869	0.1205	12.78	-
SEK/100 DKK	-0.0312	-0.0519	-0.0453	-	-0.0959	0.0437	0.9913	0.1338	9.86	-
SEK/100 NOK	0	-0.0427	-	-	0.0000	0.0590	0.9263	-	8.91	-

As already mentioned an important tool to see how well a model fits a data series is to look at the ACF and PACF for the residuals after a model has been fitted to the data. As an example of what this may look like, in figure (5) and (6) the autocorrelation function for the standardized residuals and the squared standardized residuals after the final model has been fitted can be seen. For the corresponding plots for the other data series, please see the appendix in appendix A. In the plots the 95% confidence interval has been marked. Another way to check if any serial dependence is left in the standardized residuals is to use the Q-Statistics on the standardized residuals and on the squared standardized residuals.

Next the sign bias test is considered. When this test signals that there is a problem with sign bias it is often a good idea to introduce an EGARCH model. Sometimes, however, it can be hard to get rid of the sign bias even with an EGARCH model or other versions of the GARCH model. As can be seen in table (3) which displays the probabilities, all values are insignificant at the 1% level except for SEK/JPY which shows a bad result and SEK/100 NOK where the is




Figure 5: The autocorrelation function for the standardized residuals of the SEK/EUR series after the final AR-EGARCH model has been fitted.




Figure 6: The autocorrelation function for the squared standardized residuals of the SEK/EUR series after the final AR-EGARCH model has been fitted.

a problem with the negative sign bias. It should be noted though, that several values are close to the 1% limit.

Table 3: Results from the sign bias test after the data has been fitted to the final models.

	SEK/USD	SEK/EUR	SEK/JPY	SEK/GBP	SEK/100 DKK	SEK/100 NOK
Sign Bias	0.4545	0.3539	0.6193	0.02	0.1884	0.0258
Negative Sign Bias	0.4041	0.0377	0.093	0.9572	0.1962	0.0063
Positive Sign Bias	0.1526	0.4671	4e-04	0.0963	0.0159	0.6964
Joint Effect	0.3706	0.1456	0.0000	0.0644	0.0427	0.0371

Apart from the visual tests of the ACF and PACF, the Q-Statistics and the sign bias test, three other tests are performed: the GMM orthogonality test, the non-parametric density test of Hong and Li and Kupiec's and Christoffersen's tests. The result of these tests for the different series after the data has been fitted to the models can be seen in table (4). In the table, Q_2 , Q_3 and Q_4 refer to the GMM test in equations (23), (24) and (25), and J refers to the joint test of these moments. Further, $M(1,1)$, $M(2,2)$, $M(3,3)$, $M(4,4)$, $M(1,2)$ and $M(2,1)$ in the table refer to Hong and Li's test in equation (27) and W refer to the test in equation (26). At last, LR_{uc} refer to the Kupiec's test in equation (28) and LR_{cc} refer to Christoffersen's test in equation (29). If the test failed to reject the null hypothesis H_0 this is marked in table (4) by *fail t.r.* and if the test could reject the null hypothesis then this is marked in the table by *reject*. As can be seen in the table, all models except the one for SEK/JPY pass the GMM test and the test of Hong and Li. The tests of Kupiec and Christoffersen are more troublesome: SEK/USD pass none of these tests and the other exchange rates pass the unconditional coverage test of Kupiec but fail the conditional coverage test. In all of the cases this is due to too few VaR exceedances, meaning that if the data had come from the chosen model it would have exceeded the VaR more times than the actual data did. To investigate this further, 5190 data points are simulated from the different models 100 times. The mean, standard deviation, maximum, minimum, skewness and kurtosis is then calculated for each series and the mean of these values

for the 100 simulations is calculated. These estimated values can be seen in table (5) and can be compared to the values of the original return series in table (1). In general this comparison looks good but the models seem to have a lower kurtosis than what the actual distributions have.

Table 4: Results from the GMM test, the test of Hong and Li and Kupiec's and Christoffersen's test of the ARMA-GARCH models. *Fail t.r.* means that the test failed to reject the null hypothesis and *reject* means that the test rejected the null hypothesis.

	SEK/USD	SEK/EUR	SEK/JPY	SEK/GBP	SEK/100 DKK	SEK/100 NOK
Q_2	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
Q_3	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
Q_4	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
J	fail t.r.	fail t.r.	reject	fail t.r.	fail t.r.	fail t.r.
$M(1, 1)$	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
$M(2, 2)$	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
$M(3, 3)$	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
$M(4, 4)$	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
$M(1, 2)$	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
$M(2, 1)$	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
W	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
LR_{uc}	reject	fail t.r.	fail t.r.	fail t.r.	fail t.r.	fail t.r.
LR_{cc}	reject	reject	reject	reject	reject	reject

Table 5: Mean, standard deviation, maxima, minima, skewness and kurtosis resulting from simulations from the ARMA-GARCH models chosen for each exchange rate. The values are the means of 100 simulations in which 5190 data points were generated each time.

	SEK/USD	SEK/EUR	SEK/JPY	SEK/GBP	SEK/100 DKK	SEK/100 NOK
mean	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
std. dev	0.0075	0.0048	0.0094	0.0064	0.0047	0.0048
maximum	0.0422	0.032	0.0643	0.0349	0.0314	0.0299
minimum	-0.0438	-0.0341	-0.0573	-0.0362	-0.0312	-0.0313
skewness	-0.0042	-0.0332	0.266	-0.0147	-0.0232	-0.0296
kurtosis	4.9948	7.0213	6.0863	4.6677	6.736	5.8045

As can be concluded from the results of the various tests presented above it has not always been easy to find suitable models that are not too complicated for the data series. For the data series for which the fit of the model is not very good the hope is that the remaining dependence will be removed when copulas are introduced with the purpose of modelling the residuals.

6.2 Model choice - copula models for the residuals

Once the models for the individual daily return series have been found, the search for copula models that can describe the standardized residuals of the ARMA-GARCH models begins, hereafter referred to only as "the residuals". For each residual series the daily residuals are used as well as the monthly maxima of the residuals and the quarterly maxima of the residuals. As expected, the daily residuals are all correlated with the other exchange rate residuals. The same thing is valid for the monthly and quarterly maxima of the residuals. This can be observed in table (6) where the values of the dependency measures Pearson's ρ , Kendall's τ and Spearman's ρ_S for the residuals can be seen. The values for the daily residual series are marked with

"D", the values for the monthly maxima are marked with "M" and the values for the quarterly maxima are marked with "Q". All measures show positive dependence between the exchange rates, but the dependence seems to lessen for the monthly and quarterly maxima. This could be interpreted as a sign of decreasing dependence when the time steps grow larger and that major changes in the exchange rates are less dependent on movements of other currencies than the daily fluctuations are.

Table 6: Values of Pearson's ρ , Kendall's τ and Spearman's ρ_S for the different pairs of exchange rates. The values for the daily residuals are marked with "D", the monthly maxima are marked with "M" and the quarterly maxima are marked with "Q".

	Pear D	Pear M	Pear Q	Kend D	Kend M	Kend Q	Spear D	Spear M	Spear Q
USD-EUR	0.4896	0.3818	0.3859	0.3314	0.245	0.2246	0.4694	0.3620	0.3145
USD-GBP	0.6530	0.4870	0.4320	0.4630	0.3456	0.2653	0.6416	0.4922	0.3715
USD-JPY	0.6294	0.4886	0.3314	0.4505	0.3386	0.2426	0.6220	0.4924	0.3480
USD-DKK	0.4695	0.3397	0.3676	0.3177	0.2286	0.2074	0.4507	0.3324	0.2975
USD-NOK	0.3573	0.3080	0.3355	0.2466	0.2070	0.1862	0.3564	0.3075	0.2904
EUR-GBP	0.5977	0.5298	0.5539	0.4132	0.3101	0.3427	0.5754	0.4407	0.4693
EUR-JPY	0.5217	0.2967	0.2775	0.3588	0.1810	0.1416	0.5077	0.2585	0.1984
EUR-DKK	0.9751	0.9144	0.9272	0.8952	0.7893	0.7872	0.9771	0.9013	0.8976
EUR-NOK	0.5794	0.5500	0.5335	0.4263	0.3396	0.3732	0.5818	0.4747	0.5144
GBP-JPY	0.5015	0.3937	0.2765	0.3533	0.2709	0.2387	0.5022	0.3998	0.3483
GBP-DKK	0.5689	0.4544	0.4938	0.3934	0.2534	0.2911	0.5505	0.3660	0.4019
GBP-NOK	0.4370	0.4400	0.4157	0.3020	0.2395	0.2465	0.4301	0.3483	0.3509
JPY-DKK	0.5163	0.2888	0.2529	0.3572	0.1837	0.1088	0.5054	0.2636	0.1687
JPY-NOK	0.3649	0.2416	0.1054	0.2486	0.1659	0.0923	0.3598	0.2363	0.1466
DKK-NOK	0.5787	0.4793	0.4904	0.4269	0.2990	0.3560	0.5828	0.4178	0.4889

The strength of the dependence between two data series is reflected in the copula by the copula parameters, θ . But for θ to be constant over time, and thus the copula, the dependence between two data series has to be constant. To check if this is the case the value of Kendall's τ is calculated once more, but this time a rolling estimation is used. Kendall's τ for the different residual pairs is thus calculated with a window of 250 days for the original residuals, 12 months for the monthly maxima and 3 years for the quarterly maxima. The result for the daily residuals can be seen in figures (7)-(10). As can be seen some of the estimated values of Kendall's τ seem fairly stationary whereas the estimation for some of the other residuals, for example daily JPY/DKK in figure (10), vary a lot over time. The corresponding plots for the monthly and quarterly maxima can be found in appendix A. There are several interesting ways to handle time-varying copulas, but they often suffer from problems with parameter estimation and no developed goodness-of-fit tests, see [13] for further details. Because of these problems, it will in the following be assumed that the dependence between the time series is stationary *enough* for the copula to be constant.

The first step in finding a copula is to find suitable marginal distribution for the individual residual series. For the extreme value series the marginal distribution should according to theory belong to a GEV-distribution as in equation (16). For the daily residuals the logistic distribution in equation (15) and the skewed Student's t distribution in equation (14) are found to be plausible choices of marginal distributions. An example can be seen in figure (11) where the histogram of the daily residuals from the SEK/EUR series are plotted together with the empirical cumulative distribution function (CDF) and the CDFs of the logistic and the skewed Student's t distribution. The choice between these distributions is determined by a χ^2 test. The χ^2 test

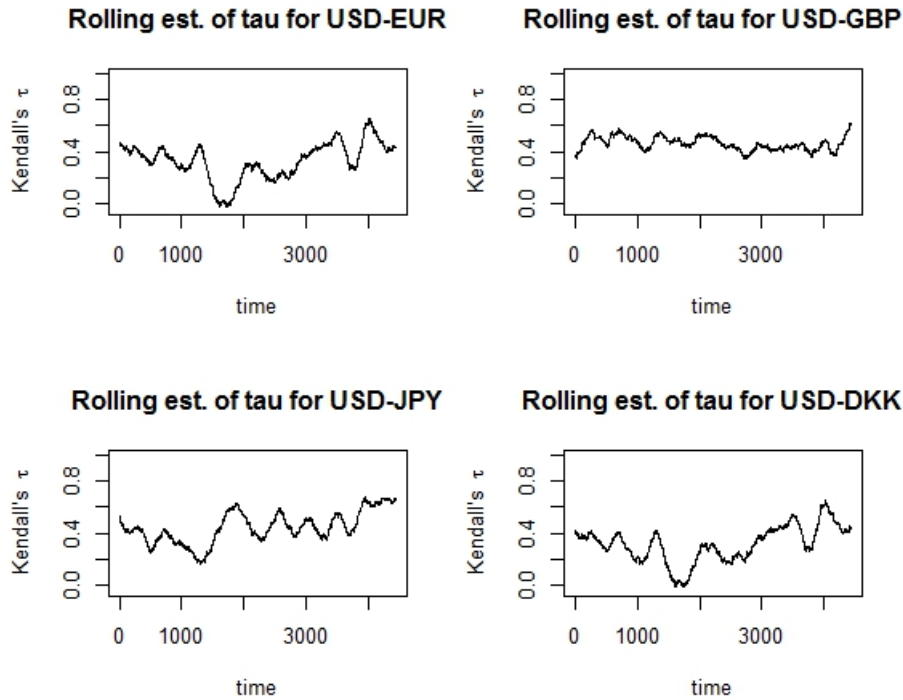


Figure 7: A rolling estimation of Kendall's τ with a window size of 250 days for the daily residual pairs USD-EUR, USD-GBP, USD-JPY and USD-DKK.

rejects the logistic distribution in several of the cases and in the other cases the p-value is close to the 5% significance level. That the residuals come from a skewed Student's t distribution can, on the other hand, not be rejected for any of the data series. Thus, the skewed Student's t distribution is the one used as marginal distribution for the daily residual series.

The scatterplot for the daily residuals can be seen in figure (12) and the corresponding plots for the monthly and quarterly maxima can be found in appendix A. Looking at this figure it seems as a good idea to investigate the fit of a normal copula and a t-copula for the different pairs. The first of the goodness-of-fit tests described in section 4.11 is used to get an estimate of how well the copulas fit the data. The normal copula is rejected for all of the residual pairs, but the t-copula seems to be a better alternative. The daily residual pairs USD-EUR, USD-JPY, USD-DKK, USD-NOK, EUR-GBP, EUR-JPY, GBP-DKK, GBP-NOK, JPY-DKK and JPY-NOK pass the goodness-of-fit test, i.e. 10 out of 15 pairs pass the test.

For the remaining daily residual pairs that did not pass the goodness-of-fit test other types of copulas are examined. It turns out however, that no copula is found for which the goodness-of-fit test is passed. In cases when it is not possible to calculate the p-value it is recommended that the copula with the lowest S_n statistics is chosen. [6] Thus, in these cases the copula with the lowest statistics value is presented as the final model, even though these have not passed the goodness-of-fit test. For all daily residual pairs except for one this means that the t-copula is chosen with skewed Student's t margins. For GBP-JPY the Plackett copula with skewed Student's t margins is chosen instead. The type of copula, margins and their parameter values can be seen in table (7).




Figure 8: A rolling estimation of Kendall's τ with a window size of 250 days for the daily residual pairs USD-NOK, EUR-GBP, EUR-JPY and EUR-DKK.

Table 7: The type of copula, its parameter θ and when applicable the degrees of freedom of the copula df_{cop} , the parameter values for the skewed Student's t marginal distributions df_1 , $skew_1$, df_2 and $skew_2$ and the p-value from the goodness-of-fit test for the copula.

	copula	θ	df_{cop}	df_1	$skew_1$	df_2	$skew_2$	$p - value$
USD-EUR	t	0.4938	5.5429	11.6446	1.0308	8.8849	1.052	0.5030
USD-GBP	t	0.6558	9.0738	11.5891	1.0134	10.3915	0.9939	0.0015
USD-JPY	t	0.6429	5.5304	11.1041	1.0148	8.9754	1.0920	0.0584
USD-DKK	t	0.4750	5.4263	11.6021	1.0325	8.949	1.0521	0.4920
USD-NOK	t	0.3727	7.0187	11.7533	1.0285	8.4801	0.9844	0.1234
EUR-GBP	t	0.5999	5.4924	9.0273	1.0330	11.2646	0.9952	0.2682
EUR-JPY	t	0.5291	6.6795	8.9340	1.0543	9.4998	1.1174	0.7507
EUR-DKK	t	0.9841	0.8975	5.9202	1.0411	6.0205	1.0411	0.0005
EUR-NOK	t	0.6162	3.0664	7.9585	1.0242	7.4518	0.9856	0.0065
GBP-JPY	Plackett	9.2808	-	12.0487	1.0466	10.3138	0.9847	0.0005
GBP-DKK	t	0.5738	5.6128	11.0116	0.9977	8.9931	1.0333	0.1314
GBP-NOK	t	0.4528	5.3248	11.1323	1.0019	8.5275	0.9865	0.3711
JPY-DKK	t	0.5267	6.4191	9.4599	1.1163	8.9735	1.0569	0.7168
JPY-NOK	t	0.3761	8.0647	9.7964	1.1124	8.6071	0.9914	0.1334
DKK-NOK	t	0.6161	3.1024	7.8611	1.0283	7.3078	0.9876	0.0035

Next, extreme value copulas should be found for the monthly and quarterly residual maxima. Since, assuming that the copulas pass the goodness-of-fit test, the copula for which the information criteria is the lowest should be chosen, the residual pairs are first fitted to the different extreme value copulas of Husler-Reiss, Gumbel, Galambos and Tawn and Akaike's and the




Figure 9: A rolling estimation of Kendall's τ with a window size of 250 days for the daily residual pairs EUR-NOK, GBP-JPY, GBP-DKK and GBP-NOK.

Bayesian information criteria are calculated and compared. Thereafter the goodness-of-fit test for extreme value copulas is performed for the copula with the lowest value of the information criteria. It is done in this order since the goodness-of-fit test is computationally expensive whereas the fitting of the copula to the data is relatively cheap. As long as the copula passes the goodness-of-fit test, this is the copula that would have been chosen and if it does not pass the test, the search would have continued in any case.

For the monthly extreme value residuals it turns out that all pairs except for EUR-DKK pass the goodness-of-fit test. For the EUR-DKK pair the remaining three copulas, i.e. Husler-Reiss, Galambos and Tawn copula, are examined with the goodness-of-fit test. Unfortunately, none of these copulas pass the goodness-of-fit test. As all models have the same value of the test statistics, the Gumbel copula is chosen as the final model for the EUR-DKK pair, even though this is *not* a good model. All of the chosen copulas, their parameters, the parameters of the marginal GEV-distributions and the p-value for the goodness-of-fit test for the monthly maxima pairs can be found in table (8).

For the quarterly extreme value residuals all pairs except for EUR-DKK pass the goodness-of-fit test. When making the goodness-of-fit test for the remaining copulas for the EUR-DKK pair the Husler-Reiss copula and the Galambos copula turn out to pass the test at the 1% significance level. Out of these the Galambos copula has the lowest information criteria value, why this copula is chosen for the EUR-DKK pair. The chosen copulas, their parameters, the parameters of the marginal GEV-distributions and the p-value for the goodness-of-fit test for the quarterly maxima pairs can be found in table (9).

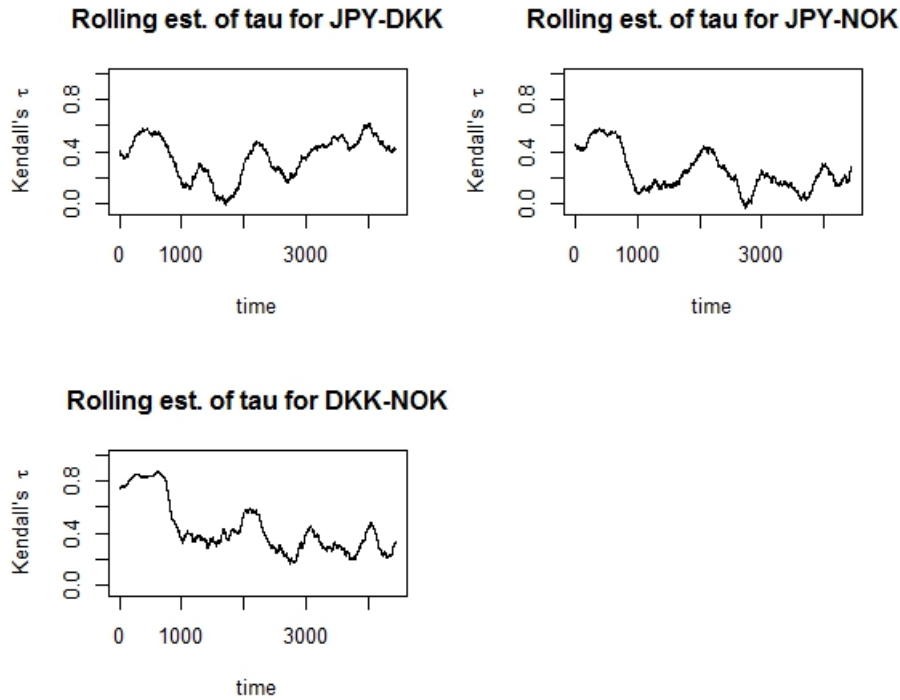


Figure 10: A rolling estimation of Kendall's τ with a window size of 250 days for the daily residual pairs JPY-DKK, JPY-NOK and DKK-NOK.

Table 8: The chosen copula, its parameter θ , the parameters of the marginal GEV-distribution $\mu_1, \sigma_1, \gamma_1, \mu_2, \sigma_2$ and γ_2 and the p-value for the goodness-of-fit test for the monthly maxima residual pairs.

	copula	θ	μ_1	σ_1	γ_1	μ_2	σ_2	γ_2	p-value
USD-EUR	HR	0.9475	1.7048	0.5527	-0.1594	1.6303	0.5566	0.0806	0.9476
USD-GBP	Tawn	0.7894	1.7036	0.5485	-0.1596	1.6582	0.4752	0.0198	0.4171
USD-JPY	HR	1.1423	1.699	0.5524	-0.1376	1.7019	0.5896	-0.0138	0.1663
USD-DKK	Gumbel	1.2764	1.7056	0.5544	-0.1562	1.6360	0.5553	0.0648	0.8736
USD-NOK	HR	0.8169	1.7034	0.552	-0.1589	1.6289	0.5826	-0.0094	0.7847
EUR-GBP	Galambos	0.7160	1.6374	0.5591	0.0601	1.6549	0.4696	0.0076	0.6648
EUR-JPY	Gumbel	1.1951	1.6397	0.5613	0.0536	1.7071	0.5893	-0.0495	0.6938
EUR-DKK	Gumbel	5.2441	1.6074	0.5441	0.1461	1.6147	0.5449	0.1270	0.0005
EUR-NOK	Tawn	0.8267	1.6240	0.5470	0.0529	1.6422	0.5783	-0.0377	0.1334
GBP-JPY	HR	0.9760	1.6524	0.4739	0.0286	1.7025	0.5892	-0.0188	0.2592
GBP-DKK	HR	0.9991	1.6569	0.4723	0.003	1.6406	0.5567	0.0512	0.8556
GBP-NOK	HR	0.9594	1.6558	0.4715	-0.0013	1.6313	0.5819	-0.0247	0.5020
JPY-DKK	Galambos	0.4360	1.7097	0.5907	-0.0491	1.6420	0.5594	0.0503	0.7468
JPY-NOK	Galambos	0.3998	1.7044	0.5880	-0.0398	1.6327	0.5853	-0.0172	0.2383
DKK-NOK	Tawn	0.7471	1.6264	0.5455	0.0468	1.6436	0.5787	-0.0417	0.1104

7 Conclusions and fields of further studies

In this report we have seen how it is possible to model different exchange rates with ARMA-GARCH models and that even though it is normally assumed that all of the dependence in the data has been modelled when the ARMA-GARCH modelling is finished, it is still possible to find




Figure 11: Histogram of the residuals from the SEK/EUR series together with the empirical CDF (in black) and the CDFs of the logistic (in dotted green) and the skewed Student's t distribution (in dotted red).

Table 9: The chosen copula, its parameter θ and the parameters of the marginal GEV-distribution $\mu_1, \sigma_1, \gamma_1, \mu_2, \sigma_2$ and γ_2 and the p-value for the goodness-of-fit test for the quarterly maxima residual pairs.

	copula	θ	μ_1	σ_1	γ_1	μ_2	σ_2	γ_2	p-value
USD-EUR	Tawn	0.5962	2.241	0.4578	-0.1768	2.2732	0.6176	0.0851	0.5380
USD-GBP	Tawn	0.7228	2.2412	0.4608	-0.1696	2.2013	0.5054	0	0.4091
USD-JPY	HR	0.9266	2.2494	0.4684	-0.1454	2.2713	0.5326	0.0748	0.8786
USD-DKK	Galambos	0.526	2.2522	0.4674	-0.1792	2.2519	0.6088	0.0398	0.4141
USD-NOK	HR	0.8738	2.2489	0.4675	-0.1541	2.2263	0.5777	0.0187	0.7727
EUR-GBP	Tawn	0.8285	2.2796	0.6121	0.0369	2.1739	0.4912	-0.0248	0.5539
EUR-JPY	HR	0.7802	2.2625	0.6056	0.0636	2.283	0.5386	0.0251	0.6998
EUR-DKK	Galambos	0.9475	1.7048	0.5527	-0.1594	1.6303	0.5566	0.0806	0.0115
EUR-NOK	Tawn	0.8479	2.2407	0.5913	0.0851	2.2407	0.5788	-0.0051	0.5679
GBP-JPY	HR	0.9248	2.1779	0.4971	0.0223	2.2693	0.5340	0.1003	0.3511
GBP-DKK	Galambos	1.5004	2.2333	0.6027	0.0778	2.2333	0.5782	0.0077	0.0115
GBP-NOK	HR	0.9599	2.1842	0.4962	-0.0272	2.2303	0.5782	0.0057	0.7847
JPY-DKK	HR	0.7386	2.2838	0.5384	0.0217	2.2516	0.6078	0.0383	0.6269
JPY-NOK	HR	0.5751	2.2789	0.5374	0.0411	2.2315	0.5797	0.0027	0.7647
DKK-NOK	Gumbel	1.5004	2.2333	0.6027	0.0778	2.2333	0.5782	0.0077	0.7268

a bivariate copula model for the residuals resulting from the ARMA-GARCH model. This is a clear sign that the exchange rates are closely connected to each other and that there is reason to take other exchange rates into account if an exchange rate should be modelled. This is especially so when risk management is concerned, as we have seen that it was easiest to find good copula models for the extreme values of the residuals. The implication of this is that if dealing with positively dependant exchange rates, you would be extra vulnerable should an extreme event occur. On the contrary, if it was possible to find two exchange rates with negative dependence, one could be used to hedge a position in the other.

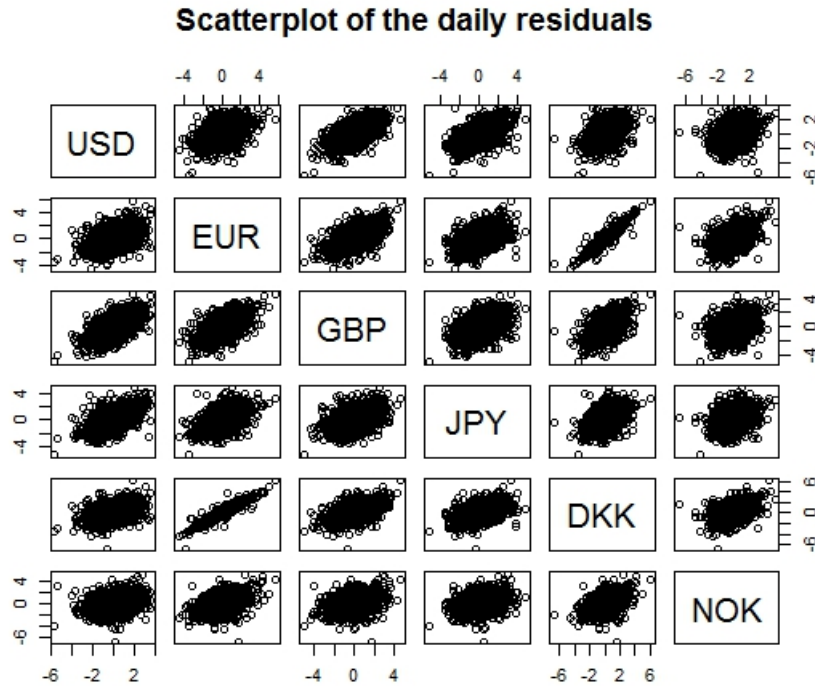


Figure 12: Scatterplot for the daily residual pairs.

As mentioned in section 2 it could be interesting to compare the models found for the data series SEK/EUR, SEK/100 DKK and SEK/100 NOK. It comes as no surprise that the AR-EGARCH models for SEK/EUR and SEK/100 DKK are close to identical as the Danish krone is bound to the euro. When it comes to the model for SEK/100 NOK there are however large differences. This is an AR-GARCH model with parameters clearly different from the other two models. This difference is a sign of the autonomy of the Norwegian krone and shows the difference in the relations between the Swedish krona and the Danish krone and that of the Swedish krona and the Norwegian krone. Comparing the copula models for the daily residuals of EUR-NOK and DKK-NOK the parameter of the t copula and the degrees of freedom are close to the same. Also the estimated parameters for the marginal distributions are very similar. For the monthly maxima models both EUR-NOK and DKK-NOK are modelled with Tawn copulas with parameters fairly close to each other. But the similarities end by the quarterly maxima models where EUR-NOK is modelled with a Tawn copula whereas DKK-NOK is modelled by a Gumbel copula.

When it comes to the usefulness of the STIBOR rate, it turned out that it could only be used in one of the cases, namely to model SEK/GBP. In the cases of the other rates, it was hard to tell if it actually made a difference, why it was left out of the models in order to keep them as simple as possible. But since the STIBOR was not deemed to be entirely useless, it could be interesting to investigate if other rates of this type could be used as external signals, perhaps in combination, e.g. the Euribor and the LIBOR, which are the Eurozone's and British counterparts to STIBOR.

It could be discussed whether or not it is sensible to use an EGARCH model for exchange rates. The reason is that there is always two sides of an exchange rate: Currency X/Y quoted in the first currency is equal to Y/X, then quoted in the second currency. As a result, that negative

innovations of an exchange rate have larger impact on the volatility than the positive innovations must mean that, quoted in the other currency, the positive innovations have larger impact than the negative innovations. That a negative innovation has larger impact on the volatility could seem natural, but that a positive innovation would have the same effect seems less so. A symmetrical model would appear to better catch this two-way behaviour. As it is, however, the EGARCH model proved a better fit for several of the exchange rates and helped removing some of the sign bias that was there while only using a GARCH model.

For the monthly residual maxima no good copula model was found for the EUR-DKK series. To find a better model more types of copulas should be tried, e.g. the models implemented in the *evd* package in R. In that case a goodness-of-fit test for extreme value copulas would however need to be implemented, since there is none in this package. It should also be noticed that there was also a problem with the EUR-DKK series for the quarterly maxima, and that the EUR-DKK series has a dependence far higher than any of the other residual pairs. This makes it more probable that there is something peculiar with this series that the other series lack and that this might be solved with another type of model. In the same way more models would have to be tried for the daily residuals, since there were several of these for which no good copula model was found. It could also be interesting to investigate why it was so much harder to find convincing models for the daily residuals than for the monthly and quarterly maxima. Could it be that there is less information to model in the daily returns since this is better caught by the ARMA-GARCH model and that more information remains in the extreme values, or is it simply depending on what copula models were tested for the different series?

A Appendix: plots and tables

The auto correlation function for the standardized residuals and the squared standardized residuals after fitting the data series to the final ARMA-GARCH models in table (2) can be found in figures (13)-(14) for SEK/USD, (5)-(6) for SEK/EUR, (15)-(16) for SEK/JPY, (17)-(18) for SEK/GBP, (19)-(20) for SEK/100 DKK and (21)-(22) for SEK/100 NOK.




Figure 13: The autocorrelation function for the standardized residuals of the SEK/USD series after the final GARCH model has been fitted.




Figure 14: The autocorrelation function for the squared standardized residuals of the SEK/USD series after the final GARCH model has been fitted.

The rolling estimation of Kendall's τ with a window size of 12 months for the monthly residual maxima pairs can be seen in figures (23)-(26). The corresponding plots for the quarterly maxima with a window size of 12 quarters can be seen in figures (27)-(30).

The scatterplots for the monthly maxima pairs and the quarterly maxima pairs can be found in figures (31) and (32).




Figure 15: The autocorrelation function for the standardized residuals of the SEK/JPY series after the final AR-EGARCH model has been fitted.




Figure 16: The autocorrelation function for the squared standardized residuals of the SEK/JPY series after the final AR-EGARCH model has been fitted.




Figure 17: The autocorrelation function for the standardized residuals of the SEK/GBP series after the final ARX-EGARCH model has been fitted.




Figure 18: The autocorrelation function for the squared standardized residuals of the SEK/GBP series after the final ARX-EGARCH model has been fitted.

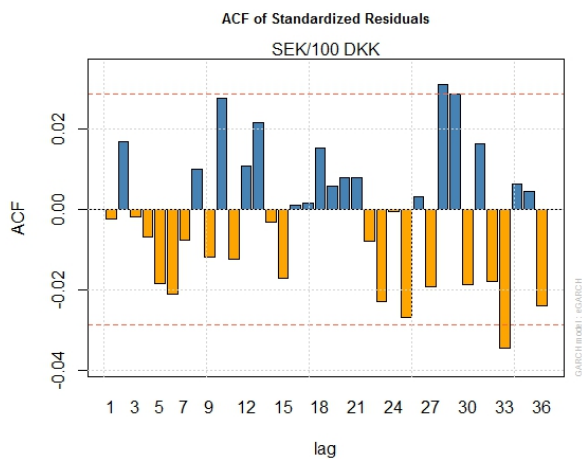


Figure 19: The autocorrelation function for the standardized residuals of the SEK/100 DKK series after the final AR-EGARCH model has been fitted.




Figure 20: The autocorrelation function for the squared standardized residuals of the SEK/100 DKK series after the final AR-EGARCH model has been fitted.




Figure 21: The autocorrelation function for the standardized residuals of the SEK/100 NOK series after the final AR-EGARCH model has been fitted.




Figure 22: The autocorrelation function for the squared standardized residuals of the SEK/100 NOK series after the final AR-EGARCH model has been fitted.




Figure 23: A rolling estimation of Kendall's τ with a window size of 12 months for the monthly residual maxima pairs USD-EUR, USD-GBP, USD-JPY and USD-DKK.




Figure 24: A rolling estimation of Kendall's τ with a window size of 12 months for the monthly residual maxima pairs USD-NOK, EUR-GBP, EUR-JPY and EUR-DKK.




Figure 25: A rolling estimation of Kendall's τ with a window size of 12 months for the monthly residual maxima pairs EUR-NOK, GBP-JPY, GBP-DKK and GBP-NOK.

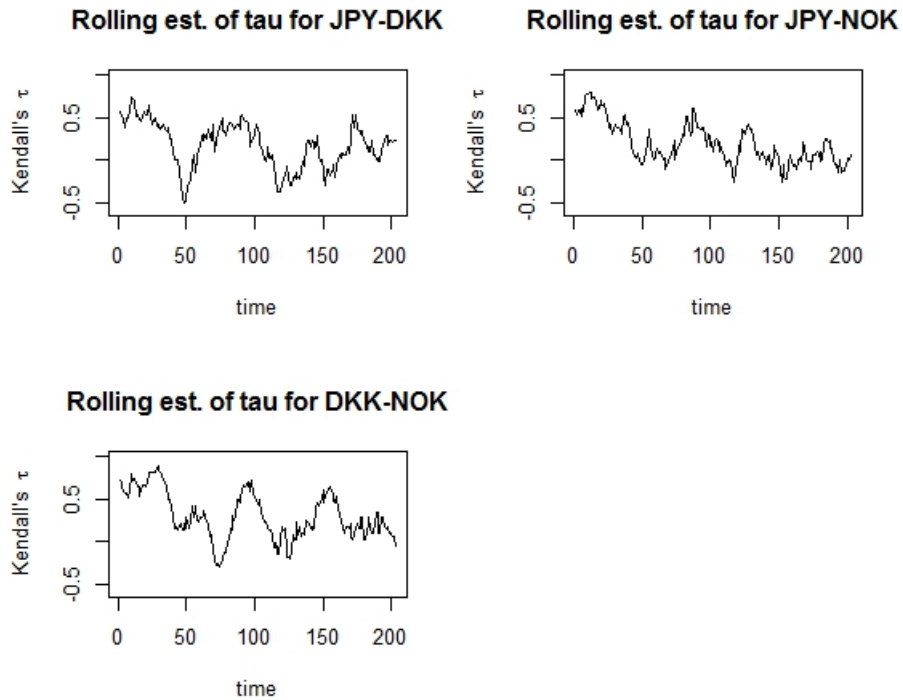


Figure 26: A rolling estimation of Kendall's τ with a window size of 12 months for the monthly residual maxima pairs JPY-DKK, JPY-NOK and DKK-NOK.

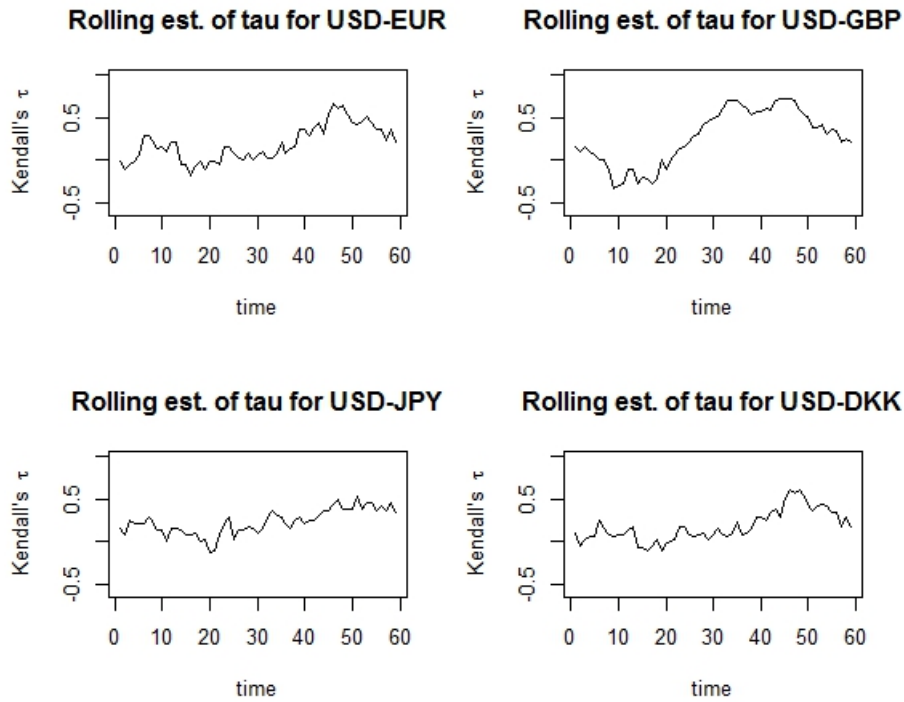


Figure 27: A rolling estimation of Kendall's τ with a window size of 12 quarters for the quarterly residual maxima pairs USD-EUR, USD-GBP, USD-JPY and USD-DKK.

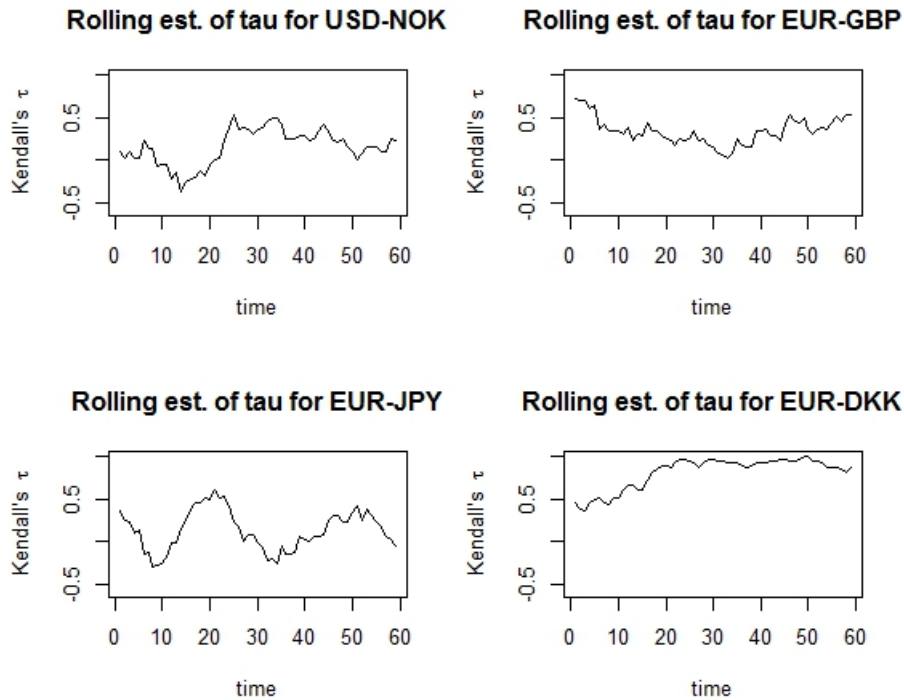


Figure 28: A rolling estimation of Kendall's τ with a window size of 12 quarters for the quarterly residual maxima pairs USD-NOK, EUR-GBP, EUR-JPY and EUR-DKK.




Figure 29: A rolling estimation of Kendall's τ with a window size of 12 quarters for the quarterly residual maxima pairs EUR-NOK, GBP-JPY, GBP-DKK and GBP-NOK.




Figure 30: A rolling estimation of Kendall's τ with a window size of 12 quarters for the quarterly residual maxima pairs JPY-DKK, JPY-NOK and DKK-NOK.





Figure 31: Scatterplot for the monthly maxima of the residual pairs. Figure 32: Scatterplot for the quarterly maxima of the residual pairs.

References

- [1] Bank for International Settlements. 2013. *Foreign exchange turnover in April 2013: preliminary global results*. Triennial Central Bank Survey, Monetary and Economic Department. <http://www.bis.org/publ/rpfx13.htm>
- [2] Coles, S. 2001. *An Introduction to Statistical Modeling of Extreme Values*. London: Springer London Ltd.
- [3] Danmarks Nationalbank. 2014. *Penge- og valutapolitik*. Danmarks Nationalbank. Accessed Jan 31. http://coolcosmos.ipac.caltech.edu/cosmic_classroom/ir_tutorial/
- [4] Dowd, D. 2005. *Measuring Market Risk*. Chichester: John Wiley & Sons.
- [5] Fregert, K., Jonung, L. 2008. *Makroekonomi. Teori, politik och institutioner*. Lund: Studentlitteratur.
- [6] Genest, C., Favre, A. 2007. *Everything you always wanted to know about copula modeling but were afraid to ask*. Journal of Hydrologic Engineering JULY/AUGUST: 347-368.
- [7] Genest, C., Kojadinovic, I., Nešlehová, J., Yan, J. 2011. *A goodness-of-fit test for bivariate extreme-value copulas*. Bernoulli 17(1): 253–275.
- [8] Genest, C., Remillard, B., Beaudoin, D. 2009. *Goodness-of-fit tests for copulas: A review and a power study*. Insurance: Mathematics and Economics 44 Issue 2: 199-213.
- [9] Ghalanos, A. 2013. *Introduction to the rugarch package*. Version 1.2-3. CRAN. http://cran.r-project.org/web/packages/rugarch/vignettes/Introduction_to_the_rugarch_package.pdf
- [10] Jakobsson, A. 2012. *Time Series Analysis and Signal Modelling*. Version 121015. Lund: Lund University.
- [11] McNeil, A. J., Frey, R. and Embrechts, P. 2005 *Quantitative Risk Management*. New Jersey: Princeton University Press.
- [12] Nelson, R.B. 2006. *An introduction to copulas*. New York: Springer.
- [13] Patton, A. 2007. *Copula-Based Models for Financial Time Series*. Oxford: The Oxford-Man Institute of Quantitative Finance, University of Oxford.
- [14] The R Core Team. 2013. *R: A Language and Environment for Statistical Computing, Reference Index*. Version 3.0.2 (2013-09-25). <http://cran.r-project.org/doc/manuals/r-release/fullrefman.pdf>
- [15] Roy, I. 2011. *Estimation of Portfolio Value at Risk using Copula*. RBI Working Paper Series 4/2011.
- [16] Schmidt, T. 2006. *Coping with copulas*. Leipzig: Department of Mathematics, University of Leipzig.
- [17] Tsay, Ruey S. 2002. *Analysis of Financial Time Series*. New Jersey: John Wiley & Sons.